# Generative Multiplane Images: Making a 2D GAN 3D-Aware

Xiaoming Zhao<sup>1,2\*</sup>, Fangchang Ma<sup>1</sup>, David Güera<sup>1</sup>, Zhile Ren<sup>1</sup>, Alexander G. Schwing<sup>2</sup>, Alex Colburn<sup>1</sup>

> <sup>1</sup> Apple
> <sup>2</sup> University of Illinois, Urbana-Champaign https://github.com/apple/ml-gmpi

**Abstract.** What is really needed to make an existing 2D GAN 3D-aware? To answer this question, we modify a classical GAN, *i.e.*, Style-GANv2, as little as possible. We find that only two modifications are absolutely necessary: 1) a multiplane image style generator branch which produces a set of alpha maps conditioned on their depth; 2) a pose-conditioned discriminator. We refer to the generated output as a 'generative multiplane image' (GMPI) and emphasize that its renderings are not only high-quality but also guaranteed to be view-consistent, which makes GMPIs different from many prior works. Importantly, the number of alpha maps can be dynamically adjusted and can differ between training and inference, alleviating memory concerns and enabling fast training of GMPIs in less than half a day at a resolution of  $1024^2$ . Our findings are consistent across three challenging and common high-resolution datasets, including FFHQ, AFHQv2 and MetFaces.

Keywords: GANs, 3D-aware generation, multiplane images

# 1 Introduction

Generative adversarial networks (GANs) [22] have been remarkably successful at sampling novel images which look 'similar' to those from a given training dataset. Notably, impressive advances have been reported in recent years which improve quality and resolution of the generated images. Most of these advances focus on the setting where the output space of the generator and the given dataset are identical and, often, these outputs are either images or occasionally 3D volumes.

The latest literature, however, has focused on generating novel outputs that differ from the available training data. This includes methods that generate 3D geometry and the corresponding texture for one class of objects, *e.g.*, faces, while the given dataset only contains widely available single-view images [56, 51, 10, 53, 16, 9, 65, 64]. No multi-view images or 3D geometry are used to supervise the training of these 3D-aware GANs. To learn the 3D geometry from such a limited form of supervision, prior work typically combines 3D-aware inductive

<sup>\*</sup> Work done as part of an internship at Apple.



Fig. 1: Making a 2D GAN 3D-aware with a minimal set of changes: on three datasets (FFHQ, AFHQv2 and MetFaces) we observe that it suffices to augment StyleGANv2 with an additional branch that generates alpha maps conditioned on depths, and to modify training by using a discriminator conditioned on pose.

biases such as a 3D voxel grid or an implicit representation with a rendering engine [56, 51, 10, 53, 16, 9, 65, 64].

Nevertheless, improving the quality of the results of these methods remains challenging: 3D-aware inductive biases are often memory-intensive explicit or implicit 3D volumes, and/or rendering is often computationally demanding, *e.g.*, involving a two-pass importance sampling in a 3D volume and a subsequent decoding of the obtained features. Moreover, lessons learned from 2D GANs are often not directly transferable because the generator output or even its entire structure has to be adjusted. This poses the question:

"What is really needed to make an existing 2D GAN 3D-aware?"

To answer this question, we aim to modify an existing 2D GAN as little as possible. Further, we aim for an efficient inference and training procedure. As a

starting point, we chose the widely used StyleGANv2 [35], which has the added benefit that many training checkpoints are publicly available.

More specifically, we develop a new generator branch for StyleGANv2 which yields a set of fronto-parallel alpha maps, in spirit similar to multiplane images (MPIs) [66]. As far as we know, we are the first to demonstrate that MPIs can be used as a scene representation for *unconditional* 3D-aware generative models. This new alpha branch is trained from scratch while the regular StyleGANv2 generator and discriminator are simultaneously fine-tuned. Combining the generated alpha maps with the single standard image output of StyleGANv2 in an end-to-end differentiable multiplane style rendering, we obtain a 3D-aware generation from different views while *guaranteeing view-consistency*. Although alpha maps have limited ability to handle occlusions, rendering is very efficient. Moreover, the number of alpha maps can be dynamically adjusted and can even differ between training and inference, alleviating memory concerns. We refer to the generated output of this method as a 'generative multiplane image' (GMPI).

To obtain alpha maps that exhibit an expected 3D structure, we find that only two adjustments of StyleGANv2 are really necessary: (a) the alpha map prediction of any plane in the MPI has to be conditioned on the plane's depth or a learnable token; (b) the discriminator has to be conditioned on camera poses. While these two adjustments seem intuitive in hindsight, it is still surprising that an alpha map with planes conditioned on their depth and use of camera poses information in the discriminator are sufficient inductive biases for 3D-awareness.

An additional inductive bias that improves the alpha maps is a 3D rendering that incorporates shading. Albeit useful, we didn't find this inductive bias to be necessary to obtain 3D-awareness. Moreover, we find that metrics designed for classic 2D GAN evaluation, *e.g.*, the Fréchet Inception Distance (FID) [28] and the Kernel Inception Distance (KID) [6] may lead to misleading results since they do not consider geometry.

In summary, our contribution is two-fold: 1) We are the first to study an MPI-like 3D-aware generative model trained with standard single-view 2D image datasets; 2) We find that conditioning the alpha planes on depth or a learnable token and the discriminator on camera pose are sufficient to make a 2D GAN 3D-aware. Other information provides improvements but is not strictly necessary. We study the aforementioned ways to encode 3D-aware inductive biases on three high-resolution datasets: FFHQ [34], AFHQv2 [12], and MetFaces [32]. Across all three datasets, as illustrated in Fig. 1, our findings are consistent.

# 2 Related Work and Background

In the following, we briefly review recent advances in classical and neural scene rendering, as well as the generation of 2D and 3D data with generative models. We then discuss the generation of 3D data using only 2D supervision. We also provide a brief review of single image reconstruction techniques before we review StyleGANv2 in greater detail. 4

Scene Rendering. Image-Based Rendering (IBR) is well-studied [11]. IBR 1) models a scene given a set of images; and 2) uses the scene model to render novel views. Methods can be grouped based on their use of scene geometry: explicit, implicit, or not using geometry. Texture-mapping methods use explicit geometry, whereas layered depth images (LDIs) [58], flow-based [11], lumigraphs [8], and tensor-based methods [4] use geometry implicitly. In contrast, light-field methods [39] do not rely on geometry. Hybrid methods [14] have also been studied. More recently, neural representations have been used in IBR, for instance, neural radiance fields (NeRFs) [45] and multiplane images (MPIs) [66, 60, 62, 21, 41]. In common to both is the goal to extract from a given set of images a volumetric representation of the observed scene. The volumetric representation in MPIs is discrete and permits fast rendering, whereas NeRFs use a continuous spatial representation.

These works differ from the proposed method in two main aspects. 1) The proposed method is unconditionally generative, *i.e.*, novel, never-before-seen scenes can be synthesized without requiring any color [62], depth, or semantic images [25]. In contrast, IBR methods and related recent advances focus on reconstructing a scene representation from a set of images. 2) The proposed method uses a collection of 'single-view images' from different scenes during training. In contrast, IBR methods use multiple viewpoints of a single scene for highly accurate reconstruction.

**2D** Generative Models. Generative adversarial networks (GANs) [22] and variational auto-encoders (VAEs) [36] significantly advanced modeling of probability distributions. In the early days, GANs were notably difficult to train whereas VAEs often produced blurry images. However, in the last decade, theoretical and practical understanding of these methods has significantly improved. New loss functions and other techniques have been proposed [40, 24, 37, 18, 13, 48, 5, 47, 43, 29, 55, 44, 3, 19, 2, 42, 61] to improve the stability of GAN optimization and to address mode-collapse, some theoretically founded and others empirically motivated. We follow the architectural design and techniques in StyleGANv2 [35], including exponential moving average (EMA), gradient penalties [44], and minibatch standard deviation [31].

**3D** Generative Models from 2D Data. 3D-aware image synthesis has gained attention recently. Sparse volume representations are used to generate photorealistic images based on given geometry input [26]. Many early approaches use voxel-based representations [67, 63, 50, 49, 27, 20], where scaling to higher resolutions is prohibitive due to the high memory footprint. Rendering at low-res followed by 2D CNN-based upsampling [52] has been proposed as a workaround, but it leads to view inconsistency. As an alternative, methods built on implicit functions, *e.g.*, NeRF, have been proposed [10, 56, 54]. However, their costly querying and sampling operations limit training efficiency and image resolutions. To generate high-resolution images, concurrently, EG3D [9], StyleNeRF [23], CIPS-3D [65], VolumeGAN [64], and StyleSDF [53] have been developed. Our work differs primarily in the choice of scene representation: EG3D uses a hybrid tri-plane representation while the others follow a NeRF-style implicit representa-

tion. In contrast, we study an MPI-like representation. In our experience, MPIs provide extremely fast rendering speed without incurring quality degradation. Most related to our work are GRAM [16] and LiftedGAN [59]. GRAM uses a NeRF-style scene representation and learns scene-specific isosurfaces. Queries of RGB and density happen on those isosurfaces. Although isosurfaces are conceptually similar to MPIs, the queries of this NeRF-style representation are expensive, limiting the image synthesis to low resolution at  $256 \times 256$ . LiftedGAN reconstructs the geometry of an image by distilling intermediate representations from a fixed StyleGANv2 to a separate 3D generator which produces a depth map and a transformation map in addition to an image. Different from our proposed approach, because of the transformation map, LiftedGAN is not strictly view-consistent. Moreover, the use of a single depth-map is less flexible.

**StyleGANv2 Revisited.** Since our method is built on StyleGANv2, we provide some background on its architecture. StyleGANv2 generates a square 2D image  $C \in \mathbb{R}^{H \times H \times 3}$  by upsampling and accumulating intermediate GAN generation results from various resolutions. Formally, the image  $C \triangleq C^H$  is obtained via

$$C^{h} = \begin{cases} \tilde{C}^{h} + \mathsf{UpSample}_{\frac{h}{2} \to h}(C^{\frac{h}{2}}), & \text{if } h \in \mathcal{R} \setminus \{4\}, \\ \tilde{C}^{4}, & \text{if } h = 4, \end{cases}$$
(1)

where  $C^h \in \mathbb{R}^{h \times h \times 3}$  is the GAN image generation at resolution h and  $\tilde{C}^h$  is the generated residual at the same resolution.  $\mathcal{R} = \{4, 8, \dots, H\}$  is the set of available resolutions whose values are powers of 2. UpSample\_{\frac{h}{2} \to h} refers to the operation that upsamples from  $\mathbb{R}^{\frac{h}{2} \times \frac{h}{2} \times 3}$  to  $\mathbb{R}^{h \times h \times 3}$ . The residual generation  $\tilde{C}^h$ at resolution h is generated with a single convolutional layer  $f_{\text{TorGB}}^h$ , *i.e.*,

$$\tilde{C}^{h} = f^{h}_{\text{ToRGB}}(\mathcal{F}^{h}), \qquad (2)$$

where  $\mathcal{F}^h \in \mathbb{R}^{h \times h \times dim_h}$  is the intermediate GAN feature representation at resolution h with  $\dim_h$  channels. These intermediate GAN feature representations at all resolutions  $\mathcal{R}$  are computed with a synthesis network  $f_{\text{Syn}}$ , *i.e.*,

$$\{\mathcal{F}^h : h \in \mathcal{R}\} = f_{\mathsf{Syn}}(\boldsymbol{\omega}),\tag{3}$$

where  $\boldsymbol{\omega}$  is the style embedding.  $\boldsymbol{\omega}$  is computed via the mapping network  $f_{\text{Mapping}}$  which operates on the latent variable  $\boldsymbol{z}$ , *i.e.*,  $\boldsymbol{\omega} = f_{\text{Mapping}}(\boldsymbol{z})$ .

# 3 Generative Multiplane Images (GMPI)

Our goal is to adjust an existing GAN such that it generates images which are 3D-aware and view-consistent, *i.e.*, the image  $I_{v_{tgt}}$  can illustrate the *exactly identical* generated object from different camera poses  $v_{tgt}$ . In order to achieve 3D-awareness and guaranteed view-consistency, different from existing prior work, we aim to augment an existing generative adversarial network, in our case StyleGANv2, as little as possible. For this, we modify the classical



Fig. 2: Overview of the developed 3D-aware generative multiplane images (GMPI). We find that two components are necessary to make a 2D StyleGANv2 3D-aware: accompanying StyleGANv2 with an MPI branch (Sec. 3.2) and conditioning the discriminator on pose during training (Sec. 3.4). We find adding shading during training (Sec. 3.5) improves the generated geometry. Please also refer to Sec. 3.3 for details about the employed MPI rendering. Green blocks denote trainable components.

generator by adding an 'alpha branch' and incorporate a simple and efficient alpha composition rendering. Specifically, the 'alpha branch' produces alpha maps of a multiplane image while the alpha composition rendering step transforms generated alpha maps and generated image into the target view  $I_{v_{tgt}}$  given a user-specified pose  $v_{tgt}$ . We refer to the output of the generator and the renderer as a 'generative multiplane image' (GMPI). To achieve 3D-awareness, we also find the pose conditioning of the discriminator to be absolutely necessary. Moreover, additional miscellaneous adjustments like the use of shading help to improve results. We first discuss the generator, specifically our alpha branch (Sec. 3.2) and our rendering (Sec. 3.3). Subsequently we detail the discriminator pose conditioning (Sec. 3.4) and the miscellaneous adjustments (Sec. 3.5).

## 3.1 Overview

An overview of our method to generate GMPIs is shown in Fig. 2. Its generator and the subsequent alpha-composition renderer produce an image  $I_{v_{tgt}}$  illustrating the generated object from a user-specified pose  $v_{tgt}$ . Images produced for different poses are guaranteed to be view-consistent. The generator and rendering achieve 3D-awareness and guaranteed view consistency in two steps. First, a novel 'alpha branch' uses intermediate representations to produce a multiplane image representation  $\mathcal{M}$  which contains alpha maps at various depths in addition to a single image. Importantly, to obtain proper 3D-awareness we find that it is necessary to condition alpha maps on their depth. We discuss the architecture of this alpha branch and its use of the plane depth in Sec. 3.2. Second, to guarantee view consistency, we employ a rendering step (Sec. 3.3). It converts the representation  $\mathcal{M}$  obtained from the alpha branch into the view  $I_{v_{tgt}}$ , which shows the generated object from a user-specified pose  $v_{tgt}$ . During training, the generated image  $I_{v_{tgt}}$  is compared to real images from a single-view dataset via a pose conditioned discriminator. We discuss the discriminator details in Sec. 3.4.



Fig. 3: The alpha branch proposed in Sec. 3.2. Here we show the generation of the alpha map of plane *i*. At each intermediate resolution  $h \in \mathcal{R}_{\alpha} = \{4, 8, \ldots, H_{\alpha}\}$ , we utilize the plane's depth  $d_i$  to transform the feature from  $\mathcal{F}^h$  (Eq. (3)) to  $\mathcal{F}^h_{\alpha_i}$  (Eq. (8)). The final alpha map  $\alpha_i$  is obtained by accumulating all intermediate results, which are generated by the single convolutional layer  $f^h_{\text{ToAlpha}}$ .

Finally, in Sec. 3.5, we highlight some miscellaneous adjustments which we found to help improve 3D-awareness while not being strictly necessary.

# 3.2 GMPI: StyleGANv2 with an Alpha Branch

Given an input latent code z, our goal is to synthesize a multiplane image inspired representation  $\mathcal{M}$  from which realistic and view-consistent 2D images can be rendered at different viewing angles. A classical multiplane image refers to a set of tuples  $(C_i, \alpha_i, d_i)$  for L fronto-parallel planes  $i \in \{1, \ldots, L\}$ . Within each tuple,  $C_i \in \mathbb{R}^{H \times H \times 3}$  denotes the color texture for the  $i^{\text{th}}$  plane and is assumed to be a square image of size  $H \times H$ , where H is independent of the plane index i. Similarly,  $\alpha_i \in [0, 1]^{H \times H \times 1}$  and  $d_i \in \mathbb{R}$  denote the alpha map and the depth of the corresponding plane, *i.e.*, its distance from a camera. i = 1 and i = L denote the planes closest and farthest from the camera.

As a significant simplification, we choose to reuse the same color texture image  $C_i \triangleq C \forall i$  across all planes. *C* is synthesized by the original StyleGANv2 structure as specified in Eq. (1). Consequently, the task of the generator  $f_{\mathsf{G}}$ has been reduced from predicting an entire multiplane image representation to predicting a single color image and a set of alpha planes, *i.e.*,

$$\mathcal{M} \triangleq \{C, \{\alpha_1, \alpha_2, \dots, \alpha_L\}\} = f_{\mathsf{G}}(\boldsymbol{z}, \{d_1, d_2, \dots, d_L\}). \tag{4}$$

For this, we propose a simple modification of the original StyleGANv2 network by adding an additional alpha branch, as illustrated in Fig. 3. For consistency we follow StyleGANv2's design: the alpha map  $\alpha_i \triangleq \alpha_i^H$  of plane *i* is obtained by upsampling and accumulating alpha maps at different resolutions. Notably, we do not generate all the way up to the highest resolution *H*, but instead use the final upsampling step

$$\alpha_i \triangleq \alpha_i^H = \mathsf{UpSample}_{H_\alpha \to H}(\alpha_i^{H_\alpha}), \tag{5}$$

if  $H_{\alpha} < H$ . Here  $H_{\alpha} \leq H$  refers to a possibly lower resolution. We will explain the reason for this design below. Following Eq. (1), we have

$$\alpha_i^h = \begin{cases} \tilde{\alpha}_i^h + \mathsf{UpSample}_{\frac{h}{2} \to h}(\alpha_i^{\frac{h}{2}}), & \text{if } h \in \mathcal{R}_\alpha \setminus \{4\}, \\ \tilde{\alpha}_i^4, & \text{if } h = 4. \end{cases}$$
(6)

Here,  $\alpha_i^h \in [0,1]^{h \times h \times 1}$  and  $\tilde{\alpha}_i^h$  are the alpha map and the residual at resolution h respectively.  $\mathcal{R}_{\alpha} \triangleq \{4, 8, \dots, H_{\alpha}\} \subseteq \mathcal{R}$  denotes the set of alpha maps' intermediate resolutions. Inspired by Eq. (2),  $\tilde{\alpha}_i^h$  is generated from an intermediate feature representation  $\mathcal{F}_{\alpha_i}^h$  through a single convolutional layer  $f_{\mathsf{ToAlpha}}^h$ :

$$\tilde{\alpha}_i^h = f_{\text{ToAlpha}}^h(\mathcal{F}_{\alpha_i}^h). \tag{7}$$

Note,  $f_{\mathsf{ToAlpha}}^h$  is shared across all planes *i*, while the input feature  $\mathcal{F}_{\alpha_i}^h$  is plane-specific. Inspired by AdaIn [30], we construct this plane-aware feature as follows:

$$\mathcal{F}^{h}_{\alpha_{i}} = \frac{\mathcal{F}^{h} - \mu(\mathcal{F}^{h})}{\sigma(\mathcal{F}^{h})} + f_{\text{Embed}}(d_{i}, \boldsymbol{\omega}), \tag{8}$$

where  $\mu(\mathcal{F}^h) \in \mathbb{R}^{\dim_h}$  and  $\sigma(\mathcal{F}^h) \in \mathbb{R}^{\dim_h}$  are mean and standard deviation of the feature  $\mathcal{F}^h \in \mathbb{R}^{h \times h \times \dim_h}$  from Eq. (3). Meanwhile,  $f_{\mathsf{Embed}}$  uses the depth  $d_i$ of plane *i* and the style embedding  $\boldsymbol{\omega}$  from Eq. (3) to compute a plane-specific embedding in the space of  $\mathbb{R}^{\dim_h}$ . Note, our design of  $\mathcal{F}^h_{\alpha_i}$  disentangles alpha map generation from a pre-determined number of planes as  $f^h_{\mathsf{ToAlpha}}$  operates on each plane individually. This provides the ability to use an arbitrary number of planes which helps to reduce artifacts during inference as we will show later.

Note, the plane specific feature  $\mathcal{F}^{h}_{\alpha_{i}} \forall i$  will in total occupy L times the memory used by the feature  $\mathcal{F}^{h}$ . This might be prohibitive if we were to generate these intermediate features up to a resolution of H. Therefore, we only use this feature until some lower resolution  $H_{\alpha}$  in Eq. (5). We will show later that this design choice works well on real-world data.

### 3.3 Differentiable Rendering in GMPI

We obtain the desired image  $I_{v_{tgt}}$  which illustrates the generated MPI representation  $\mathcal{M} = \{C, \{\alpha_1, \ldots, \alpha_L\}\}$  from the user-specified target view  $v_{tgt}$  with an MPI renderer in two steps: 1) a warping step transforms the representation  $\mathcal{M}$ from its canonical pose  $v_{cano}$  to the target pose  $v_{tgt}$ ; 2) a compositing step combines the planes into the desired image  $I_{v_{tgt}}$ . Importantly, both steps entail easy computations which are end-to-end differentiable such that they can be included into any generator. Please see Appendix ?? for details.

#### 3.4 Discriminator Pose Conditioning

For the discriminator  $f_{D}$  to encourage 3D-awareness of the generator, we find the conditioning of the discriminator on camera poses to be essential. Formally,

9

inspired by Miyato and Koyama [46], the final prediction of the discriminator is

$$\log P(y = \operatorname{real}|I_{v_{\operatorname{tgt}}}, v_{\operatorname{tgt}}) \propto \operatorname{Normalize}(f_{\operatorname{Embed}}^{\mathrm{D}}(v_{\operatorname{tgt}})) \cdot f_{\mathrm{D}}(I_{v_{\operatorname{tgt}}})^{\mathrm{T}}, \qquad (9)$$

where  $P(y = \text{real}|I_{v_{tgt}}, v_{tgt})$  denotes the probability that image  $I_{v_{tgt}}$  from the camera pose  $v_{tgt}$  is real.  $v_{tgt} \in \mathbb{R}^{16}$  denotes the flattened extrinsic matrix of the camera pose.  $f_{\text{Embed}}^{D} : \mathbb{R}^{16} \mapsto \mathbb{R}^{16}$  denotes an embedding function, while Normalize(·) results in a zero-mean, unit-variance embedding.  $f_{D}(\cdot) \in \mathbb{R}^{16}$  denotes the feature computed by the discriminator D. For real images of faces, *e.g.*, humans and cats,  $v_{tgt}$  can be estimated via off-the-shelf tools [17, 1].<sup>3</sup>

#### 3.5 Miscellaneous Adjustment: Shading-guided Training

Inspired by [54], we incorporate shading into the rendering process introduced in Sec. 3.3. Conceptually, shading will amplify artifacts of the generated geometry that might be hidden by texture, encouraging the alpha branch to produce better results. To achieve this, we adjust the RGB component  $C \in \mathbb{R}^{H \times W \times 3}$  via

$$\widehat{C} = C \cdot \left(k_a + k_d \, \boldsymbol{l} \cdot \mathsf{N}\left(D_{v_{\mathsf{cano}}}\right)\right),\tag{10}$$

where  $k_a$  and  $k_d$  are coefficients for ambient and diffuse shading,  $l \in \mathbb{R}^3$  indicates lighting direction, and  $N(D_{v_{cano}}) \in \mathbb{R}^{H \times W \times 3}$  denotes the normal map computed from the depth map  $D_{v_{cano}}$  (obtained by using the canonical alpha maps  $\alpha_i$ , see Eq. (??) in Appendix ??). We find shading to slightly improve results while not being required for 3D-awareness. Implementation details are in Appendix ??.

#### 3.6 Training

**Model structure.** The trainable components of GMPI are  $f_{\text{ToRGB}}^{h}$  (Eq. (2)),  $f_{\text{Syn}}$  and  $f_{\text{Mapping}}$  (Eq. (3)),  $f_{\text{ToAlpha}}^{h}$  (Eq. (7)),  $f_{\text{Embed}}$  (Eq. (8)), and  $f_{\text{D}}$  and  $f_{\text{Embed}}^{\text{D}}$  (Eq. (9)). Please check the appendix for implementation details. Our alpha maps (Eq. (4)) are equally placed in the disparity (inverse depth) space during training and we set  $H_{\alpha} = 256$  (Eq. (5)).

**Initialization.** For any training, we initialize weights from the officially-released StyleGANv2 checkpoints.<sup>4</sup> This enables a much faster training as we will show. **Loss.** We use  $\theta$  to subsume all trainable parameters. Our training loss consists of a non-saturating GAN loss with R1 penalties [44], *i.e.*,

$$\mathcal{L}_{\theta} = \mathbb{E}_{I_{v_{\text{tgt}}}, v_{\text{tgt}}} \left[ f(\log P(y = \text{real}|I_{v_{\text{tgt}}}, v_{\text{tgt}})) \right] \\ + \mathbb{E}_{I, v_{\text{tgt}}} \left[ f(\log P(y = \text{real}|I, v_{\text{tgt}})) + \lambda |\nabla_I \log P(y = \text{real}|I, v_{\text{tgt}})|^2 \right].$$
(11)

Here  $f(x) = -\log(1 + \exp(-x))$ ,  $\lambda = 10.0$  in all studies, *I* refers to real images and  $v_{tgt}$  denotes the corresponding observer's pose information.

<sup>&</sup>lt;sup>3</sup> Concurrently, EG3D [9] also finds that pose conditioning of the discriminator is required for their tri-plane representation to produce 3D-aware results, corroborating that this form of inductive bias is indeed necessary.

<sup>4</sup> https://github.com/NVlabs/stylegan2-ada-pytorch

### 4 Experiments

We analyze GMPI on three datasets (FFHQ, AFHQv2 and MetFaces) and across a variety of resolutions. We first provide details regarding the three datasets before discussing evaluation metrics and quantitative as well as qualitative results.

### 4.1 Datasets

**FFHQ.** The FFHQ dataset [34] consists of 70,000 high-quality images showing real-world human faces from different angles at a resolution of  $1024 \times 1024$ . To obtain the pose of a face we use an off-the-shelf pose estimator [17].

**AFHQv2.** The AFHQv2-Cats dataset [12, 33] consists of 5,065 images showing faces of cats from different views at a resolution of  $512 \times 512$ . We augment the dataset by horizontal flips and obtain the pose of a cat's face via an off-the-shelf cat face landmark predictor [1] and OpenCV's perspective-n-point algorithm. **MetFaces.** The MetFaces dataset [32] consists of 1,336 images showing high-

quality faces extracted from the collection of the Metropolitan Museum of Art.<sup>5</sup> To augment the dataset we again use horizontal flips. To obtain the pose of a face we use an off-the-shelf pose estimator [17].

## 4.2 Evaluation Metrics

We follow prior work and assess the obtained results using common metrics: **2D GAN metrics.** We report Fréchet Inception Distance (FID) [28] and Kernel Inception Distance (KID) [6], computed by using 50k artificially generated images that were rendered from random poses and 1) 50k real images for FFHQ; 2) all 10,130 real images in the *x*-flip-augmented dataset for AFHQv2-Cat [12]. **Identity similarity (ID).** Following EG3D [9], we also evaluate the level of facial identity preservation. Concretely, we first generate 1024 MPI-like representations  $\mathcal{M}$ . For each representation we then compute the identity cosine similarity between two views rendered from random poses using ArcFace [15, 57].

**Depth accuracy (Depth).** Similar to [9, 59], we also assess geometry and depth accuracy. For this we utilize a pre-trained face reconstruction model [17] to provide facial area mask and pseudo ground-truth depth map  $\hat{D}_{v_{tgt}}$ . We report the MSE error between our rendered depth  $D_{v_{tgt}}$  (Eq. (??)) and  $\hat{D}_{v_{tgt}}$  on areas constrained by the face mask. Note, following prior work, we normalize both depth maps to zero-mean and unit-variance. The result is obtained by averaging over 1024 representations  $\mathcal{M}$ .

**Pose accuracy (Pose).** Following [9, 59], we also study the 3D geometry's pose accuracy. Specifically, for each MPI, we utilize a pose predictor [17] to estimate the yaw, pitch, and roll of a rendered image. The predicted pose is then compared to the pose used for rendering via the MSE. The reported result is averaged over 1024 representations  $\mathcal{M}$ . Notably but in hindsight expected, 2D metrics, as well as ID, lack the ability to capture 3D errors which we will show next.

<sup>&</sup>lt;sup>5</sup> https://metmuseum.github.io/

Table 1: Speed comparison. '-' indicates that corresponding papers do not report this result. 1) Training: EG3D, GRAM, StyleNeRF and GMPI report training time when using 8 Tesla V100 GPUs. pi-GAN uses two RTX 6000 GPUs or a single RTX 8000 GPU. For GMPI, the results reported in this paper come from 3/5/11-hour training for a resolution of  $256^2/512^2/1024^2$  with initialization from official pretrained-checkpoints. 2) Inference: we measure frames-per-second (FPS) for each model. GMPI uses 96 planes. GRAM reports speed by utilizing a specified mesh rasterizer [38] while the others use pure forward passes. EG3D uses an RTX 3090 GPU. GRAM and StyleNeRF use a Tesla V100 GPU, which we also utilize to run GIRAFFE, pi-GAN, LiftedGAN, and our GMPI. We observe: GMPI is quick to train and renders the fastest among all approaches.

	Res.	Unit	GIRAFFE [	52] p	i-GAN [10]	LiftedGAN [59]	EG3D <sup>†</sup> [9]	$\mathrm{GRAM}^{\dagger}$ [16]	StyleNeRF <sup><math>\dagger</math></sup> [23]	GMPI
Train	-	$\mathrm{Time}{\downarrow}$		-	56h	_	8.5d	. 3-7d	3d	3/5/11h
ICe	$256^{2}$	$\text{FPS}\uparrow$	2	50	1.63	25	36	180	16	$328^{*}$
erei	$512^{2}$	$\mathrm{FPS}\uparrow$		-	0.41	-	35	-	14	$83.5^{*}$
Inf	$1024^{2}$	$\mathrm{FPS}\uparrow$		-	0.10	_	_		11	$19.4^{*}$

<sup>†</sup> We quote results from their papers.

\* We report the rendering speed. GMPI, different from StyleNeRF, pi-GAN, and LiftedGAN, only needs a single forward pass to generate the scene representation. Further rendering doesn't involve the generator. We hence follow EG3D to report inference FPS without forward pass time that is 82.34 ms, 99.97 ms, and 115.10 ms for 96 planes at 256<sup>2</sup>, 512<sup>2</sup>, and 1024<sup>2</sup> on a V100 GPU.

#### 4.3 Results

We provide speed comparison in Tab. 1 and a quantitative evaluation in Tab. 2. With faster training, GMPI achieves on-par or better performance than start-ofthe-art when evaluating on 256<sup>2</sup> images and can generate high-resolution results up to 1024<sup>2</sup> which most baselines fail to produce. Specifically, GMPI results on resolutions 256<sup>2</sup>, 512<sup>2</sup>, and 1024<sup>2</sup> are reported after 3/5/11-hours of training. Note, the pretrained StyleGANv2 initialization for FFHQ (see Sec. 3.6) requires training for 1d 11h (256<sup>2</sup>), 2d 22h (512<sup>2</sup>), and 6d 03h (1024<sup>2</sup>) with 8 Tesla V100 GPUs respectively, as reported in the official repo.<sup>4</sup> In contrast, EG3D, GRAM, and StyleNeRF require training of at least three days. At a resolution of 256<sup>2</sup>, 1) GMPI outperforms GIRAFFE, pi-GAN, LiftedGAN, and GRAM on FID/KID while outperforming StyleSDF on FID; 2) GMPI demonstrates better identity similarity (ID) than GIRAFFE, pi-GAN, and LiftedGAN; 3) GMPI outperforms GIRAFFE regarding depth; 4) GMPI performs best among all baselines on pose accuracy. Overall, GMPI demonstrates that it is a flexible architecture which achieves 3D-awareness with an affordable training time.

#### 4.4 Ablation Studies

In order to show the effects of various design choices, we run ablation studies and selectively drop out the discriminator conditioned on pose (DPC), the plane-specific feature  $\mathcal{F}_{\alpha_i}^h$ , and shading. The following quantitative (Tab. 3) and qualitative (Fig. 4 and Fig. 5) studies were run using a resolution of 512<sup>2</sup> to set baseline metrics, and answer the following questions:

Table 2: Quantitative results. GMPI uses 96 planes and no truncation trick [34, 7] is applied. '-' indicates that the corresponding paper does not report this result. KID is reported in KID×100. For GIRAFFE, pi-GAN, and LiftedGAN, we quote results from [9]. Compared to baselines, GMPI achieves on-par or even better performance despite much faster training as reported in Tab. 1. Moreover, GMPI can produce high-resolution images which most baselines fail to achieve, demonstrating the flexibility.

					AFHQv2-Cat				
			$\mathrm{FID}\downarrow$	KID↓	ID↑	Depth↓	$\mathrm{Pose}{\downarrow}$	$\mathrm{FID}\downarrow$	KID↓
$256^{2}$	1	GIRAFFE [52]	31.5	1.992	0.64	0.94	0.089	16.1	2.723
	2	pi-GAN $128^2$ [10]	29.9	3.573	0.67	0.44	0.021	16.0	1.492
	3	LiftedGAN [59]	29.8	_	0.58	0.40	0.023	—	_
	4	$\text{GRAM}^{\dagger}$ [16]	29.8	1.160	-	_	—	—	_
	5	$StyleSDF^{\dagger}$ [53]	11.5	0.265	-	_	_	$12.8^{*}$	$0.447^{*}$
	6	StyleNeRF <sup><math>\dagger</math></sup> [23]	8.00	0.370	-	-	-	$14.0^{*}$	$0.350^{*}$
	7	$CIPS-3D^{\dagger}$ [65]	6.97	0.287	-	_	_	—	-
	8	$EG3D^{\dagger}$ [9]	4.80	0.149	0.76	0.31	0.005	3.88	0.091
	9	GMPI	11.4	0.738	0.70	0.53	0.004	$n/a^{\S}$	$n/a^{\S}$
$512^{2}$	10	$EG3D^{\dagger}$ [9]	4.70	0.132	0.77	0.39	0.005	2.77	0.041
	11	StyleNeRF <sup><math>\dagger</math></sup> [23]	7.80	0.220	-	_	_	$13.2^{*}$	$0.360^{*}$
	12	GMPI	8.29	0.454	0.74	0.46	0.006	7.79	0.474
$1024^{2}$	13	$CIPS-3D^{\dagger}$ [65]	12.3	0.774	-	_	_	_	-
	14	StyleNeRF <sup><math>\dagger</math></sup> [23]	8.10	0.240	-	_	—	—	_
	15	GMPI	7.50	0.407	0.75	0.54	0.007	$n/a^{\S}$	$n/a^{\S}$

<sup>†</sup> We quote results from their papers.

<sup>8</sup> Performance is reported on the whole AFHQv2 dataset instead of only cats.
 <sup>8</sup> No GMPI results on AFHQv2 for Row 9 and 13 since there are no available pre-

trained StyleGANv2 checkpoints for AFHQv2 with a resolution of  $256^2$  or  $1024^2$ .

**Baseline condition** (Tab. 3 row a) – a pre-trained StyleGANv2 texture without transparency ( $\alpha_i \forall i$  is set to 1). This representation is not truly 3D, as it is rendered as a textured block in space. It is important to note that the commonly used metrics for 2D GAN evaluation, *i.e.*, FID, KID, as well as ID yield good results. They are hence not sensitive to 3D structure. In contrast, the depth and pose metrics do capture the lack of 3D structure in the scene.

Can a naïve MPI generator learn 3D? (Tab. 3 row b) – a generator which only uses  $\mathcal{F}^{h}$  (Eq. (2)) without being trained with pose-conditioned discriminator, plane-specific features, or a shading loss. The Depth (2.190 vs. 2.043) and Pose (0.062 vs. 0.060) metrics show that this design fails, yielding results similar to the baseline (row a).

**Does DPC alone enable 3D-awareness?** (Tab. 3 row c) – we now use discriminator pose conditioning (DPC) while ignoring plane-specific features  $\mathcal{F}_{\alpha_i}^h$ . This design improves the Depth (0.501 vs. 2.190) and Pose (0.006 vs. 0.062) metrics. However upon closer inspection of the alpha maps shown in Fig. 5a we observe that the generator fails to produce realistic structures. The reason that this design performs well on Depth and Pose metrics is primarily due to these two metrics only evaluating the surface while not considering the whole volume. Table 3: Ablation studies. We evaluate at a resolution of  $512^2$ . DPC refers to discriminator pose conditioning (Sec. 3.4),  $\mathcal{F}_{\alpha_i}^h$  refers to the plane-specific feature introduced in Eq. (8), and Shading indicates the shading-guided training discussed in Sec. 3.5. #planes denotes the number of planes we used during evaluation. Note, rows (b) and (c) can only use 32 planes during training and inference. Therefore, to make the comparison fair, all ablations use 32 planes. We provide an additional 96-plane result for our full model, which is used in Tab. 2. Please check Sec. 4.4 and Fig. 4 for a detailed discussion regarding this table as we find that 2D metrics, *e.g.*, FID/KID as well as ID can be misleading when evaluating 3D generative models.

	#planes	DPC	$\mathcal{F}^h_{\alpha_i}$	Shading	$\rm FFHQ$					AFHQv2-Cat	
					FID↓	KID↓	ID↑	$\mathrm{Depth}{\downarrow}$	Pose↓	$FID\downarrow$	KID↓
(a)	32				6.64	0.368	0.91	2.043	0.060	4.29	0.199
(b)	32				4.61	0.167	0.89	2.190	0.062	4.00	0.166
(c)	32	1			7.98	0.347	0.75	0.501	0.006	7.13	0.385
(d)	32		1		4.35	0.150	0.89	2.140	0.061	3.70	0.132
(e)	32	1	1		7.19	0.313	0.73	0.462	0.006	7.54	0.433
(f)	32	1	1	1	7.40	0.337	0.74	0.457	0.006	7.93	0.489
(g)	96	1	1	1	8.29	0.454	0.74	0.457	0.006	7.79	0.474



Fig. 4: The images in this figure correspond the the ablation studies in Tab. 3. Panels (a)-(g) correspond to Tab. 3's rows (a)-(g). All results are generated from the same latent code z. The face is rendered with a camera positioned to the right of the subject, *i.e.*, the rendered face should look to the left of the viewer. (a) is the 2D image produced by the pre-trained StyleGANv2. Note how GMPI becomes 3D-aware in (e), and generates geometry and texture that is occluded in the pre-trained StyleGANv2 image. (e) vs. (f): shading-guided training (Sec. 3.5) alleviates geometric artifacts such as concavities in the forehead. (f) vs. (g): the ability to use more planes during inference (Sec. 3.2) reduces "stair step" artifacts, visible on the cheek and the ear.

**Does**  $\mathcal{F}_{\alpha_i}^h$  **alone enable 3D-awareness?** (Tab. 3 row d) – we use only the plane-specific feature  $\mathcal{F}_{\alpha_i}^h$  without DPC. Geometry related metrics (Depth and Pose) perform poorly, which is corroborated by Fig. 4 (d), indicating an inability to model 3D information.

**Do DPC and**  $\mathcal{F}_{\alpha_i}^h$  enable 3D-awareness? (Tab. 3 row e) – we combine both plane-specific features  $\mathcal{F}_{\alpha_i}^h$  and DPC. This produce good values for Depth and Pose metrics. Visualizations in Fig. 4e, and Fig. 5b verify that GMPI successfully generates 3D-aware content. However, the FID/KID, as well as ID values are generally worse than non-3D-aware generators (row a-e).



Fig. 5: Qualitative results accompanying Tab. 3. For alpha maps, from top to bottom, left to right, we show  $\alpha_1$  to  $\alpha_L$  respectively (Eq. (4)): the whiter, the denser the occupancy. Grey boundaries are added for illustration.

**Does shading improve 3D-awareness?** (Tab. 3 row f) – we use plane-specific features  $\mathcal{F}_{\alpha_i}^h$ , DPC and shading loss. As discussed in Sec. 3.1, DPC and  $\mathcal{F}_{\alpha_i}^h$  are sufficient to make a 2D GAN 3D-aware. However, inspecting the geometry reveals artifacts such as the concave forehead in Fig. 4e. Shading-guided rendering tends to alleviate these issues (Fig. 4e vs. f) while not harming the quantitative results (Tab. 3's row e vs. f).

Can we reduce aliasing artifacts? (Tab. 3 row g) – we use plane-specific features  $\mathcal{F}_{\alpha_i}^h$ , DPC, shading loss and 96 planes. Due to the formulation of  $\mathcal{F}_{\alpha_i}^h$ , we can generate an arbitrary number of planes during inference, which helps avoid "stair step" artifacts that can be observed in Fig. 4f vs. g.

Qualitative results are shown in Fig. 1. Please see the appendix for more.

# 5 Conclusion

To identify what is really needed to make a 2D GAN 3D-aware we develop generative multiplane images (GMPIs) which guarantee view-consistency. GMPIs show that a StyleGANv2 can be made 3D-aware by 1) adding a multiplane image style branch which generates a set of alpha maps conditioned on their depth in addition to a single image, both of which are used for rendering via an end-toend differentiable warping and alpha compositing; and by 2) ensuring that the discriminator is conditioned on the pose. We also identify shortcomings of classical evaluation metrics used for 2D image generation. We hope that the simplicity of GMPIs inspires future work to fix limitations such as occlusion reasoning.

Acknowledgements: We thank Eric Ryan Chan for discussion and providing processed AFHQv2-Cats dataset. Supported in part by NSF grants 1718221, 2008387, 2045586, 2106825, MRI #1725729, NIFA award 2020-67021-32799.

# References

- Cat hipsterizer. https://github.com/kairess/cat\_hipsterizer (2022), accessed: 2022-03-06 9, 10
- 2. Aneja, J., Schwing, A.G., Kautz, J., Vahdat, A.: A contrastive learning approach for training variational autoencoder priors. ICLR (2020) 4
- 3. Arjovsky, M., Chintala, S., Bottou, L.: Wasserstein GAN. ArXiv (2017) 4
- 4. Avidan, S., Shashua, A.: Novel view synthesis in tensor space. In: CVPR (1997) 4
- Berthelot, D., Schumm, T., Metz, L.: Began: Boundary equilibrium generative adversarial networks. In: arXiv preprint arXiv:1703.10717 (2017) 4
- Binkowski, M., Sutherland, D.J., Arbel, M., Gretton, A.: Demystifying mmd GANs. ICLR (2018) 3, 10
- Brock, A., Donahue, J., Simonyan, K.: Large scale gan training for high fidelity natural image synthesis. ICLR (2019) 12
- 8. Buehler, C., Bosse, M., McMillan, L., Gortler, S., Cohen, M.: Unstructured Lumigraph Rendering. SIGGRAPH (2001) 4
- Chan, E.R., Lin, C.Z., Chan, M.A., Nagano, K., Pan, B., Mello, S.D., Gallo, O., Guibas, L., Tremblay, J., Khamis, S., Karras, T., Wetzstein, G.: Efficient Geometryaware 3D Generative Adversarial Networks. In: CVPR (2022) 1, 2, 4, 9, 10, 11, 12
- Chan, E.R., Monteiro, M., Kellnhofer, P., Wu, J., Wetzstein, G.: pi-GAN: Periodic implicit generative adversarial networks for 3D-aware image synthesis. In: CVPR (2021) 1, 2, 4, 11, 12
- Chen, S.E., Williams, L.: View Interpolation for Image Synthesis. SIGGRAPH (1993) 4
- Choi, Y., Uh, Y., Yoo, J., Ha, J.W.: Stargan v2: Diverse image synthesis for multiple domains. In: CVPR (2020) 3, 10
- 13. Cully, R.W.A., Chang, H.J., Demiris, Y.: MAGAN: Margin adaptation for generative adversarial networks. ArXiv (2017) 4
- Debevec, P.E., Taylor, C.J., Malik, J.: Modeling and Rendering Architecture from Photographs: A hybrid geometry- and image-based approach. SIGGRAPH (1996) 4
- Deng, J., Guo, J., Zafeiriou, S.: ArcFace: Additive Angular Margin Loss for Deep Face Recognition. CVPR (2019) 10
- Deng, Y., Yang, J., Xiang, J., Tong, X.: GRAM: Generative Radiance Manifolds for 3D-Aware Image Generation. CVPR (2022) 1, 2, 5, 11, 12
- Deng, Y., Yang, J., Xu, S., Chen, D., Jia, Y., Tong, X.: Accurate 3D Face Reconstruction With Weakly-Supervised Learning: From Single Image to Image Set. CVPRW (2019) 9, 10
- Deshpande, I., Zhang, Z., Schwing, A.G.: Generative Modeling using the Sliced Wasserstein Distance. In: CVPR (2018) 4
- Deshpande, I., Hu, Y.T., Sun, R., Pyrros, A., Siddiqui, N., Koyejo, O., Zhao, Z., Forsyth, D.A., Schwing, A.G.: Max-Sliced Wasserstein Distance and Its Use for GANs. CVPR (2019) 4
- Gadelha, M., Maji, S., Wang, R.: 3D shape induction from 2D views of multiple objects. In: 3DV (2017) 4
- Ghosh, S., Lv, Z., Matsuda, N., Xiao, L., Berkovich, A., Cossairt, O.: Liveview: Dynamic target-centered mpi for view synthesis. arXiv preprint arXiv:2107.05113 (2021) 4
- Goodfellow, I.J., Pouget-Abadie, J., Mirza, M., Xu, B., Warde-Farley, D., Ozair, S., Courville, A.C., Bengio, Y.: Generative adversarial nets. In: NeurIPS (2014) 1, 4

- Gu, J., Liu, L., Wang, P., Theobalt, C.: StyleNeRF: A style-based 3D-aware generator for high-resolution image synthesis. ICLR (2022) 4, 11, 12
- Gulrajani, I., Ahmed, F., Arjovsky, M., Dumoulin, V., Courville, A.: Improved training of wasserstein GANs. In: NeurIPS (2017) 4
- 25. Habtegebrial, T., Jampani, V., Gallo, O., Stricker, D.: Generative view synthesis: From single-view semantics to novel-view images. NeurIPS (2020) 4
- Hao, Z., Mallya, A., Belongie, S., Liu, M.Y.: GANcraft: Unsupervised 3D neural rendering of minecraft worlds. In: ICCV (2021) 4
- 27. Henzler, P., Mitra, N.J., Ritschel, T.: Escaping Plato's cave: 3D shape from adversarial rendering. In: ICCV (2019) 4
- Heusel, M., Ramsauer, H., Unterthiner, T., Nessler, B., Hochreiter, S.: GANs trained by a two time-scale update rule converge to a local nash equilibrium. In: NeurIPS (2017) 3, 10
- 29. Heusel, M., Ramsauer, H., Unterthiner, T., Nessler, B., Hochreiter, S.: GANs trained by a two time-scale update rule converge to a local nash equilibrium. In: NeurIPS (2017) 4
- Huang, X., Belongie, S.J.: Arbitrary style transfer in real-time with adaptive instance normalization. ICCV (2017) 8
- Karras, T., Aila, T., Laine, S., Lehtinen, J.: Progressive growing of GANs for improved quality, stability, and variation. ICLR (2018) 4
- 32. Karras, T., Aittala, M., Hellsten, J., Laine, S., Lehtinen, J., Aila, T.: Training generative adversarial networks with limited data. NeurIPS (2020) 3, 10
- Karras, T., Aittala, M., Laine, S., Härkönen, E., Hellsten, J., Lehtinen, J., Aila, T.: Alias-free generative adversarial networks. In: NeurIPS (2021) 10
- Karras, T., Laine, S., Aila, T.: A style-based generator architecture for generative adversarial networks. CVPR (2019) 3, 10, 12
- Karras, T., Laine, S., Aittala, M., Hellsten, J., Lehtinen, J., Aila, T.: Analyzing and improving the image quality of stylegan. CVPR (2020) 3, 4
- 36. Kingma, D.P., Welling, M.: Auto-encoding variational bayes. ICLR (2014) 4
- Kolouri, S., Rohde, G.K., Hoffman, H.: Sliced wasserstein distance for learning gaussian mixture models. In: CVPR (2018) 4
- Laine, S., Hellsten, J., Karras, T., Seol, Y., Lehtinen, J., Aila, T.: Modular primitives for high-performance differentiable rendering. TOG (2020) 11
- 39. Levoy, M., Hanrahan, P.: Light field rendering. SIGGRAPH (1996) 4
- 40. Li, C.L., Chang, W.C., Cheng, Y., Yang, Y., Póczos, B.: MMD GAN: Towards deeper understanding of moment matching network. In: NeurIPS (2017) 4
- 41. Li, J., Feng, Z., She, Q., Ding, H., Wang, C., Lee, G.H.: Mine: Towards continuous depth mpi with nerf for novel view synthesis. In: ICCV (2021) 4
- 42. Li, Y., Schwing, A.G., Wang, K.C., Zemel, R.: Dualing GANs. In: NeurIPS (2017) 4
- 43. Lin, Z., Khetan, A., Fanti, G., Oh, S.: Pacgan: The power of two samples in generative adversarial networks. In: NeurIPS (2018) 4
- 44. Mescheder, L.M., Geiger, A., Nowozin, S.: Which training methods for GANs do actually converge? In: ICML (2018) 4, 9
- Mildenhall, B., Srinivasan, P.P., Tancik, M., Barron, J.T., Ramamoorthi, R., Ng, R.: NeRF: Representing Scenes as Neural Radiance Fields for View Synthesis. In: ECCV (2020) 4
- 46. Miyato, T., Koyama, M.: cgans with projection discriminator. ICLR (2018) 9
- 47. Mroueh, Y., Sercu, T.: Fisher gan. In: NeurIPS (2017) 4
- 48. Mroueh, Y., Sercu, T., Goel, V.: McGan: Mean and covariance feature matching GAN. ArXiv (2017) 4

- 49. Nguyen-Phuoc, T., Li, C., Theis, L., Richardt, C., Yang, Y.L.: HoloGAN: Unsupervised learning of 3D representations from natural images. In: ICCV (2019) 4
- Nguyen-Phuoc, T., Richardt, C., Mai, L., Yang, Y.L., Mitra, N.J.: BlockGAN: Learning 3D object-aware scene representations from unlabelled images. In: NeurIPS (2020) 4
- 51. Niemeyer, M., Geiger, A.: Giraffe: Representing scenes as compositional generative neural feature fields. CVPR (2021) 1, 2
- 52. Niemeyer, M., Geiger, A.: GIRAFFE: Representing scenes as compositional generative neural feature fields. In: CVPR (2021) 4, 11, 12
- Or-El, R., Luo, X., Shan, M., Shechtman, E., Park, J.J., Kemelmacher-Shlizerman, I.: StyleSDF: High-resolution 3d-consistent image and geometry generation. CVPR (2022) 1, 2, 4, 12
- 54. Pan, X., Xu, X., Loy, C.C., Theobalt, C., Dai, B.: A shading-guided generative implicit model for shape-accurate 3d-aware image synthesis. NeurIPS (2021) 4, 9
- 55. Salimans, T., Zhang, H., Radford, A., Metaxas, D.: Improving GANs using optimal transport. In: ICLR (2018) 4
- 56. Schwarz, K., Liao, Y., Niemeyer, M., Geiger, A.: GRAF: Generative radiance fields for 3D-aware image synthesis. In: NeurIPS (2020) 1, 2, 4
- 57. Serengil, S.I., Ozpinar, A.: Lightface: A hybrid deep face recognition framework. In: 2020 Innovations in Intelligent Systems and Applications Conference (ASYU). IEEE (2020) 10
- Shade, J., Gortler, S., Hey, L.W., Szeliski, R.: Layered Depth Images. SIGGRAPH (1998) 4
- Shi, Y., Aggarwal, D., Jain, A.K.: Lifting 2d StyleGAN for 3D-aware face generation. CVPR (2021) 5, 10, 11, 12
- 60. Srinivasan, P.P., Tucker, R., Barron, J.T., Ramamoorthi, R., Ng, R., Snavely, N.: Pushing the boundaries of view extrapolation with multiplane images. In: CVPR (2019) 4
- Sun, R., Fang, T., Schwing, A.G.: Towards a Better Global Loss Landscape of GANs. In: Proc. NeurIPS (2020) 4
- 62. Tucker, R., Snavely, N.: Single-view view synthesis with multiplane images. In: CVPR (2020) 4
- Wu, J., Zhang, C., Xue, T., Freeman, W.T., Tenenbaum, J.B.: Learning a probabilistic latent space of object shapes via 3D generative-adversarial modeling. In: NeurIPS (2016) 4
- 64. Xu, Y., Peng, S., Yang, C., Shen, Y., Zhou, B.: 3D-aware Image Synthesis via Learning Structural and Textural Representations. CVPR (2022) 1, 2, 4
- Zhou, P., Xie, L., Ni, B., Tian, Q.: CIPS-3D: A 3D-Aware Generator of GANs Based on Conditionally-Independent Pixel Synthesis. ArXiv (2021) 1, 2, 4, 12
- Zhou, T., Tucker, R., Flynn, J., Fyffe, G., Snavely, N.: Stereo magnification: Learning view synthesis using multiplane images. TOG (2018) 3, 4
- Zhu, J.Y., Zhang, Z., Zhang, C., Wu, J., Torralba, A., Tenenbaum, J.B., Freeman, W.T.: Visual object networks: Image generation with disentangled 3D representations. In: NeurIPS (2018) 4