

GVP: Generative Volumetric Primitives

Mallikarjun B R Xingang Pan Mohamed Elgharib Christian Theobalt
Max Planck Insitute for Informatics, Saarland Informatics Campus

Abstract

Advances in 3D-aware generative models have pushed the boundary of image synthesis with explicit camera control. To achieve high-resolution image synthesis, several attempts have been made to design efficient generators, such as hybrid architectures with both 3D and 2D components. However, such a design compromises multiview consistency, and the design of a pure 3D generator with high resolution is still an open problem. In this work, we present Generative Volumetric Primitives (GVP), the first pure 3D generative model that can sample and render 512-resolution images in real-time. GVP jointly models a number of volumetric primitives and their spatial information, both of which can be efficiently generated via a 2D convolutional network. The mixture of these primitives naturally captures the sparsity and correspondence in the 3D volume. The training of such a generator with high degree of freedom is made possible through a knowledge distillation technique. Experiments on several datasets demonstrate superior efficiency and 3D consistency of GVP over the state-of-the-art.

1. Introduction

Synthesizing photorealistic objects with high-resolution and multi-view consistency is a long-term goal in computer vision and graphics. Recently, advances in 3D-aware generative adversarial networks (GANs) have made inspiring progress towards this goal [26, 4, 20, 12, 5]. By learning from unstructured 2D images, 3D-aware GANs can synthesize different views of generated objects with explicit control over camera pose. This is made possible through 3D generator architectures that naturally possess a 3D inductive bias.

The key challenge towards a high-quality 3D-aware GAN lies in the design of an efficient 3D generator. For example, voxel-based 3D representation [13] suffers from cubic memory growth, and thus can only afford a limited resolution. Recently, a more widely used representation for 3D-aware GAN is neural radiance fields (NeRF) [19, 26, 4], which integrates a coordinate-based MLP and volume ren-

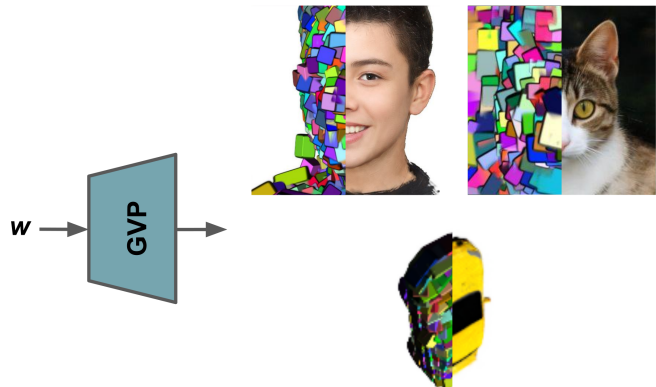


Figure 1. We propose the first 3D-aware generative model that relies on pure 3D representation and can render images at a resolution of 512×512 in real-time. At the core of our method, a scene is represented by a number of volumetric primitives that can be rendered very efficiently. Our method can generalize to different object categories such as human and cat portraits and cars.

dering. As NeRF requires querying the MLP densely across 3D space, it is prohibitively computationally expensive for high-resolution image rendering. Several works propose to improve efficiency via a compromised hybrid generator that models a 3D feature space via NeRF followed by a 2D CNN-based renderer that upsamples the feature to a high-resolution 2D image [20, 12, 5, 22]. However, the use of a 2D CNN compromises 3D consistency and is vulnerable to artifacts for camera views unseen during training. Although a few recent works attempt to improve the efficiency of a pure 3D generative model [28, 27], the synthesis of high-resolution images in real-time remains an open problem.

In this work, we present Generative Volumetric Primitives (GVP), which, to the best of our knowledge, is the first 3D-aware generative model that is based on a pure 3D representation and can sample and render 512-resolution images in real-time. To achieve high rendering efficiency, we take the first attempt to build a 3D generative model based on the mixture of volumetric primitives (MVP) representation [17]. Each primitive models the color and density of a local volume and a number of primitives are composed together to form the entire 3D volume, as shown in Fig. 1. Unlike the original MVP that heavily relies on mesh track-

ing to determine the spatial information (*i.e.*, position, orientation, and scale) of the primitives, our generator learns to model the volumetric primitives and their spatial information jointly. Such a representation can be efficiently modeled via a 2D convolutional network architecture and naturally captures spatial sparsity, thus is extremely efficient to sample and render compared to previous 3D-aware GANs.

The use of multiple movable primitives also introduces a high degree of freedom to our generator, making it unstable to train via a conventional adversarial loss on the raw training images. To address this issue, we propose to bypass the adversarial training via knowledge distillation of a pretrained 3D-aware GAN, EG3D [5]. Such knowledge distillation significantly stabilizes training and inherits the generative capability of EG3D. Although EG3D has inconsistent rendering for out-of-distribution camera views, we found the in-distribution views sufficient to train our model with high quality, which would then be able to perform consistent rendering for any camera views.

We conduct an extensive evaluation of GVP on several datasets including human face, cat face, and cars. Thanks to the highly efficient design, GVP achieves much faster rendering than previous pure 3D GANs and preserves better multiview consistency compared with those with hybrid architectures. Visualization of the learned primitives demonstrates that GVP effectively captures the sparsity of the 3D volume, avoiding redundant computation and memory costs. The predicted positions of the volumetric primitives also adapt to different samples, which to some extent captures correspondence between samples.

2. Related works

Building generative models is a long-standing computer vision problem. Early efforts include works such as learning 3D Morphable models of the human face [2, 24]. These approaches, however, lack photorealism and are trained using controlled data captured with multi-view camera rigs. With the advances in deep learning, several generative models have been proposed [14, 15]. These models are learned purely from in-the-wild 2D images and produce highly photorealistic results. However, they are 2D-based and hence produce results that are not multi-view consistent. The past couple of years witnessed a rapidly growing interest in 3D generative models. This is largely fueled with the recent advances in implicit scene representations and Neural Radiance Fields (NeRF) [23, 19] which quickly revolutionized 3D scene understanding and rendering in an unprecedented manner. While there are several 3D generative methods designed for a specific class of objects such as the human face [3], body [21] and hands [9], several other approaches are generic [4, 27, 10, 5, 1, 29]. These approaches combine volumetric rendering with adversarial training to achieve highly photorealistic renderings that are multi-view consistent.

tent.

The works of Chan *et al.* [4] and Schwarz *et al.* [26] were one of the first to show how to marry NeRF representation with adversarial training. They use co-ordinate based MLP to parameterize their generator and train in an adversarial manner. As they need to densely sample space to render the scene, it is quite slow. VoxGraf [27] acknowledges the fact that volumetric rendering is usually very slow due to querying multiple points along rays. To address this limitation, VoxGraf utilizes sparse voxel grid representations. This improves computational efficiency, but still, their method is not real-time even at 256×256 resolution. Deng *et al.* [10] proposed an alternative solution to handle the large computational requirements of volumetric rendering. Their approach uses 2D manifolds to guide point sampling and radiance field learning. These manifolds are embedded in the 3D volume and help in reducing the computational complexity while still producing multi-view consistent results. However, such manifolds can lead to clear rendering artifacts, especially in side-view angles, where the manifolds are almost parallel to camera rays. Rebain *et al.* [25] showed it is possible to learn a generative model using single-view in-the-wild 2D images without adversarial training. They follow an autoencoder architecture that learns a shared latent representation. The method is trained in a self-supervised manner, where an off-the-shelf 2D landmark detector is used to determine the camera poses. While most of these methods [4, 27] and more [30, 28] produce photorealistic renderings that are 3D-consistent, they suffer from expensive computational requirements that prohibit them from rendering in real-time.

To address the expensive computational requirement of the previously discussed generative models, another class of methods performs volumetric rendering in a low-resolution instead [20, 5, 12, 22, 1, 32]. This is then followed by a super-resolution network that upscales the output to the full resolution. An example is the method EG3D by Chan *et al.* [5]. This method extracts 2D StyleGAN2 features [15] and utilizes a 3D tri-planar representation together with a lightweight feature decoder to extract 3D features. A volumetric renderer operating in a low resolution produces an image output, which is then passed through the super-resolution network. The method is trained in an adversarial manner using in-the-wild 2D images. EG3D [5] produces impressive multi-view consistent photorealistic results and is significantly more computationally efficient than previously discussed methods [4, 27, 30, 28]. This super-resolution-based framework has also been utilized in several other works [12, 22, 1, 32], including methods that handle articulated objects [1] and others that use an SDF representation [22]. However, the framework suffers from one main limitation, the fact that it is not real-3D. While the super-resolution network helps in reducing computational com-

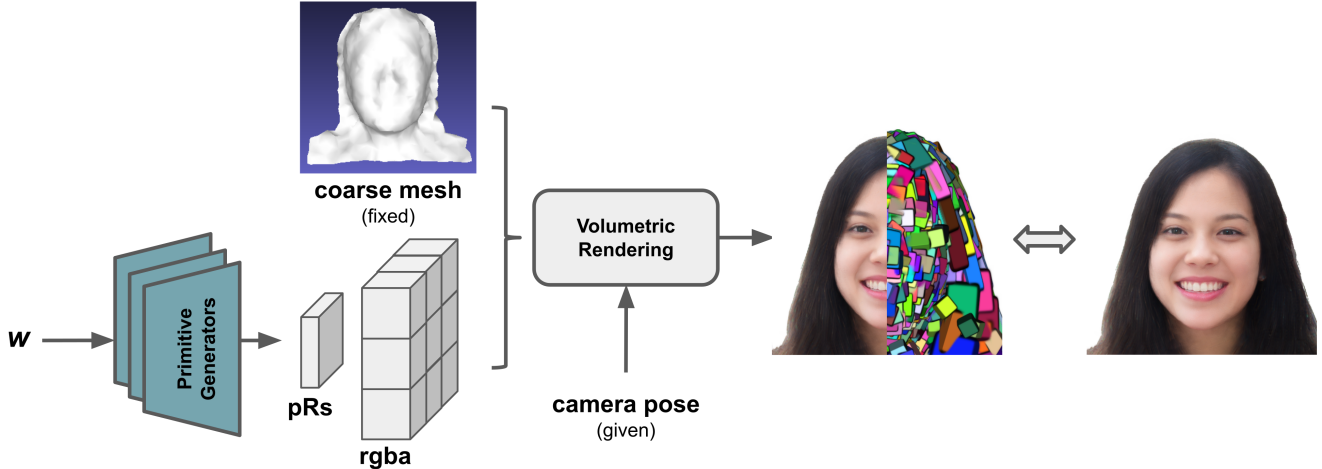


Figure 2. **Method overview:** We propose a generative model that represents each sample using a number of volumetric primitives (voxel grid) that can be rendered very efficiently in real-time. Our method consists of 3 2D-CNN-based generators. The geometric generator takes latent vector w and outputs the delta position, delta orientation, and delta scales of primitives with respect to initialized position, orientation, and scale using a coarse mesh. The output of the geometric generator is represented as pRs. The second generator takes the latent vector w as input and outputs the density values of all the primitives at once. And the last generator takes latent vector w and viewing direction as input and outputs color content of the primitives. The density and color outputs are jointly represented as rgba. Note we use pre-trained mapping network of EG3D [5] to map the gaussian noise to our latent vector w . The scene can then be volumetrically rendered efficiently to synthesize samples under any camera pose.

plexity significantly, it doesn't guarantee multi-view consistent results. Hence, the framework fails when rendering from viewpoints that are out of the training distribution.

3. Method

The goal of our method is to build efficient generative volumetric models. We take inspiration from the recent method, MVP [17] to represent the 3D scene. In specific, MVP [17] proposes an efficient 3D scene representation with a mixture of volumetric primitives. But to train MVP, they need registered meshes of an identity performing diverse expressions. Our goal is to build and generalize the training across different categories, where sophisticated 3D morphable models don't exist. Here we show, it's possible to build a generative model without the need for sophisticated morphable models for different categories. One way to train this framework is by learning the whole model end-to-end in an adversarial manner using only loss from an image-based discriminator similar to existing methods [4]. Our initial experiments in that direction proved to be unstable because of multiple moving components. The overview of our method can be found in Fig. 2 To address this issue, we trained our model in a knowledge-distillation fashion using pre-trained EG3D [5] models. In the following, we describe the representation of our model in detail in Sec. 3.1 and also provide information about training in Sec. 3.2.

3.1. Representation

Volumetric Primitives: As mentioned before, we take inspiration from MVP [17] to model our scene efficiently. The efficiency arises because of mainly 2 reasons. One, by having only sparsely distributed volumetric primitives in the scene, they can avoid sampling unnecessary regions that do not contribute to volumetric rendering. Second, the content of the primitives can be generated at once by an efficient CNN network instead of multiple queries required by a coordinate based MLP methods [4]. As we want to learn a generative model for a category, let $\mathbf{w} \in \mathbb{R}^{512}$ be the latent vector that defines a sample. Let N_{prim} be the number of primitives in the scene and each of the primitives is a 3D voxel grid representing only a small region in the 3D space. The position, orientation, and scale of k^{th} primitive in 3D space are defined by $\mathbf{t}_k(\mathbf{w}) \in \mathbb{R}^3$, $\mathbf{R}_k(\mathbf{w}) \in \text{SO}(3)$, and $\mathbf{s}_k(\mathbf{w}) \in \mathbb{R}^3$ respectively. The content of each primitive is defined by a dense voxel grid $V_k(\mathbf{w}) \in \mathbb{R}^{4 \times M_x \times M_y \times M_z}$ that contains color and opacity information. In all our experiments, we use $M_x = M_y = M_z = 32$. The payload of the volumetric primitive contains color and opacity and is defined as a function of sample \mathbf{w} .

Guide Mesh: MVP [17] utilizes tracked mesh to guide the position and orientation of primitives. Although it is possible to obtain tracked mesh for faces, obtaining the same for other categories like cats and cars becomes challenging. Our surprising finding is that, using our training




Figure 3. Here we show samples of our method trained on the FFHQ dataset. We also provide renderings of primitives. The color coding of primitives is consistent across different samples. Please note the semantically similar parts are occupied by the same colored primitives across different samples, which provides correspondence information.

strategy, we do not need registered meshes across different samples of the category to train the primitives. We make use of a fixed mesh as the guide mesh, \mathcal{M} for all samples, and learn the delta position, orientation, and scales for all the primitives. In order to get the fixed guide mesh \mathcal{M} , we extract a mesh from a manually chosen EG3D [5] generated sample and build its UV map using Blender [8]. We observe that the performance of our model does not depend much on the type of sample chosen. We obtain the guide mesh’s contribution of primitive position and orientation similar to MVP and let it be $\hat{\mathbf{t}}_k(\mathcal{M})$ and $\hat{\mathbf{R}}_k(\mathcal{M})$. We refer readers to MVP [17] for the exact details. And we also learn delta position, orientation, and scales as a function of latent vector \mathbf{w} to generalize to different samples. The final position, orientation, and scale are defined by:

$$\begin{aligned} \mathbf{t}_k(\mathbf{w}) &= \hat{\mathbf{t}}_k(\mathcal{M}) + \delta_{\mathbf{t}_k}(\mathbf{w}) \\ \mathbf{R}_k(\mathbf{w}) &= \hat{\mathbf{R}}_k(\mathcal{M}) \cdot \delta_{\mathbf{R}_k}(\mathbf{w}) \\ \mathbf{s}_k(\mathbf{w}) &= \hat{\mathbf{s}}_k + \delta_{\mathbf{s}_k}(\mathbf{w}) \end{aligned}$$

We set $\hat{\mathbf{s}}_k$ to a fixed value for all primitives which was found empirically.

Network Architecture: Our model mainly consists of 3 generators. The 3 generators are geometric generator \mathcal{G}_{geo} and two payload generators $\mathcal{G}_a, \mathcal{G}_{rgb}$. The geometric generator $\mathcal{G}_{geo} : \mathbb{R}^{512} \rightarrow \mathbb{R}^{9 \times N_{prim}}$ is for obtaining the

delta position, orientation, and scale of the primitives. And the payload generators $\mathcal{G}_a : \mathbb{R}^{512} \rightarrow \mathbb{R}^{1 \times M^3 \times N_{prim}}$ and $\mathcal{G}_{rgb} : \mathbb{R}^{512+3} \rightarrow \mathbb{R}^{3 \times M^3 \times N_{prim}}$ are used to obtain the opacity and color values.

Efficient Differentiable Rendering Once we have the primitives and their attributes, we make use of efficient differentiable rendering to render our scene. For a given ray $r_{\mathbf{p}}(t) = \mathbf{o}_{\mathbf{p}} + t\mathbf{d}_{\mathbf{p}}$ with starting position $\mathbf{o}_{\mathbf{p}}$ and ray direction $\mathbf{d}_{\mathbf{p}}$,

$$\mathcal{I}_{\mathbf{p}} = \int_{t_{\min}}^{t_{\max}} \mathbf{V}_{\text{col}}(r_{\mathbf{p}}(t)) \cdot \frac{dT(t)}{dt} \cdot dt .$$

$$T(t) = \min \left(\int_{t_{\min}}^t \mathbf{V}_{\alpha}(r_{\mathbf{p}}(t)) \cdot dt, 1 \right) .$$

Here, \mathbf{V}_{col} and \mathbf{V}_{α} are the color and opacity values at a given point on the primitive.

3.2. Training

One way to train this framework is by learning the whole model end-to-end in an adversarial manner using only loss from an image-based discriminator similar to existing methods [4]. Our initial experiments in that direction proved to be unstable because of multiple moving components. To stabilize training, we propose to train our model in a



Figure 4. Here we show samples of our method trained on the AFHQ dataset. We also provide renderings of primitives. The color coding of primitives is consistent across different samples.

knowledge-distillation fashion using pre-trained EG3D [5] models.

The generators are trained in a supervised manner. As EG3D can synthesize fairly consistent multi-view data for in-distribution camera views, we make use of it as the supervision. Since the latent space of EG3D is known to be well structured, we mimic its properties by using their latent space \mathbf{w} as \mathbf{w} in our model. Using EG3D intermediate latent space is critical in building a high-quality model as we will show later. Given multi-view renderings \mathbf{X} and their corresponding latent space \mathbf{w} , we employ the following loss function to train our generators in a supervised manner.

$$\mathcal{L} = \mathcal{L}_{\text{rec}} + \mathcal{L}_{\text{disc}} + \lambda_{\text{perc}} \mathcal{L}_{\text{perc}} \quad (1)$$

where \mathcal{L}_{rec} is an L_1 reconstruction loss between rendered images from our model and \mathbf{X} . As reconstruction loss alone results in blurry results, we also employ a discriminator loss, which helps in getting sharper results. The discriminator loss is defined as follows,

$$\mathcal{L}_{\text{disc}} = f(D(\mathcal{I}(\mathbf{w}))) + f(-D(\mathbf{X})) + \lambda_{\text{reg}} \|\nabla D(\mathbf{X})\|^2. \quad (2)$$

where D is an image-based discriminator, $f(u) = -\log(1 + \exp(-u))$, and λ_{reg} is the coefficient for R_1 regularization.

Although discriminator loss increases the sharpness of the results, it also introduces grid-like artifacts. We found

having an additional perceptual loss helps in mitigating such effects. The perceptual loss $\mathcal{L}_{\text{perc}}$ is an LPIPS-based [31] loss between rendered images from our model and \mathbf{X} .

4. Results

Datasets We demonstrate the results of our method on three datasets: FFHQ [15], AFHQv2-Cats [7], and Shapenet Cars [6]. FFHQ [15] is a large-scale portrait dataset of human faces under diverse camera poses. AFHQ-Cats [7] is a similar dataset for cat portraits. Shapenet Cars [6] is a synthetically rendered image dataset of different cars.

Baselines We compare our model with 3 baselines. The first one is EG3D [5] as this is the state-of-the-art method that can synthesize high-resolution multi-view images in real-time. We also train another model with the same architecture as our model, but with a learnable latent code for each sample. This is similar to auto-decoder training followed in LoIneRF [25] but trained with the dataset described above. We call this model MVP++. This experiment shows the importance of leveraging the \mathbf{w} latent space in EG3D to train our models in a stable manner. Finally, we qualitatively compare to Authentic Volumetric Avatars [3]. Although it [3] is not a generative model, since they also use MVP [17] representation, we provide a qualitative compar-




Figure 5. Here we show samples of our method trained on the Shapenet Cars dataset. We also provide renderings of primitives. The color coding of primitives is consistent across different samples. Please note that our method is robust to different car shapes, even though we use a coarse fixed mesh as an initialization.

ison. Please note editing expressions is not a goal of our method, unlike [3].

Training Details We create a dataset of around 600k multi-view images and their latent spaces by sampling the pre-trained EG3D model on the FFHQ dataset. Similarly, we sample 300k images for the training model on AFHQ cats and Shapenet cars datasets. Similar to MVP [17], we also use opacity fade factor and volume minimization prior. The opacity fade factor encourages primitives to explain the scene content by the movement of primitives than to have primitives just modify their payload content. And the volume minimization prior makes the maximum usage of voxels to explain the opaque part of the scene rather than also including transparent parts. We refer readers to MVP [17] for more details. We set $N_{prim} = 1024$, $\lambda_{perc} = 20$ and $\lambda_{reg} = 0.0001$ and use mesh \mathcal{M} with 1024 vertices in all our experiments. The discriminator used for the discriminator loss is similar to the one used in π -GAN [4] and is learned with a learning rate of $1e-5$. We train our models at 512×512 resolution with a batch size of 32 images for around 5 days on 4 Quadro RTX 8000 GPUs. We use Adam optimizer [16] with a learning rate of 0.001 for the parameters of the generator. Please refer to the supplemental document for more details.

Qualitative Results We present the qualitative results of our method in Fig. 3, Fig. 4 and Fig. 5 with models trained

	MVP++	EG3D	Ours
FID (FFHQ) ↓	92.25	7.53	31.93
Speed ↓	22.1	46.1	22.6

Table 1. Quantitative comparisons using the FID score metric (a lower value is better) on the FFHQ dataset and computation time (in milliseconds). Here we see that ours perform better than MVP++ while EG3D performs the best in FID score. Our method is considerably faster than EG3D. Our method has other advantages of having pure 3D representation, rendering speed, and can provide correspondences across different samples.

	EG3D	Ours
FID (AFHQ) ↓	3.65	28.44

Table 2. Quantitative comparisons using the FID score metric (a lower value is better) on the AFHQ dataset. EG3D performs better than our method on FID score, while our method has other advantages of having pure 3D representation, rendering speed, and can provide correspondences across different samples.

	F1 (Out-Out) ↑	F1 (Out-In) ↑
EG3D	0.2448	0.1832
Ours	0.3333	0.2906

Table 3. Quantitative comparisons using the F1 score of detected Facial Action Units (a higher value is better). Out-Out denotes the metric computed using renderings of the same sample under 2 views, which are both outside training-camera-pose-distribution. Out-In denotes the metric computed using renderings of the same sample under 2 views, with 1 view in training-camera-pose-distribution and the other being outside. Here we can see that our method clearly outperforms EG3D in both metrics.

on FFHQ, AFHQ, and Shapenet car datasets respectively. Our method can synthesize multi-view consistent results under different camera positions. It can also be observed that our method can generalize well to different categories of objects. We also provide visualization of primitives rendering with consistent color coding across various samples. One can observe the correspondence between semantically similar parts of samples. In Fig. 6, we provide interpolation results between 2 randomly sampled faces. One can observe the smooth transition between them where the primitive move accordingly, which provides correspondence information that might be useful for downstream tasks. This is better appreciated in the video. Please refer to the supplemental video for the same.

Next, we provide a qualitative comparison of our method with 3 baselines. In Fig. 7, we compare our method with EG3D [5] and MVP++. MVP++ has the same architecture and supervision as our method and the only difference is that we use the output of the mapping network of EG3D [5] as our latent space and MVP++ learns the latent vectors while training similar to LoLNeRF [25]. One can observe




Figure 6. Here we show interpolation results of 2 random samples in the latent space. Even though our scene representation is represented by discrete primitives, we can observe how the primitives smoothly adapt themselves to synthesize high-quality results. For example, notice the primitive represented by an orange block covering a portion of long hair on the right-hand side adapts itself to change its position, orientation, and scale to finally represent a portion of short hair. This shows that our method captures correspondences across different identities.




Figure 7. Here, we compare our method to 2 baselines. We provide 2 novel view renderings of a random sample for all methods. MVP++ is trained in an auto-decoder manner, and as a result, has very poor random samples in its latent space. EG3D can synthesize high-quality samples, but with the help of a 2D-based super-resolution module. In contrast, our method is synthesized with pure 3D representation and is faster than EG3D.

the random samples of MVP++ clearly suffer from severe artifacts. It is also observed numerically in Tab. 1. The training strategy used in the Auto-Decoders framework results in sparse supervision of limited samples, equal to the number of training samples in latent space. In contrast, a model learned in an adversarial manner gets more supervision across dense points in latent space. As a result, the models trained in an adversarial manner are known to have well-behaved latent space. Even though our model is not trained in a purely adversarial manner, as we use the latent space of EG3D, we clearly transfer good characteristics of the latent space of the adversarially trained model.

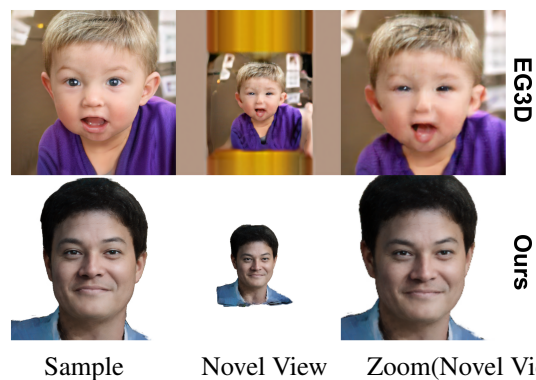


Figure 8. Here we compare our method (second row) with EG3D (first row) in out-of-training-distribution of camera pose. The first column shows samples for both our and EG3D in training distribution and the second column shows rendering with the camera placed a bit far and the third column shows the zoomed-in version of the second column. As can be seen, EG3D quality is not consistent.

We also provide a qualitative comparison with EG3D [5] in Fig. 7 8. The samples from EG3D have slightly better perceptual quality than that of our method in capturing high-frequency details. This quality comes because of their 2D super-resolution module, which is known to capture high-frequency details better. In contrast, our representation is purely 3D without any 2D components. This ensures multi-view consistency even for extreme out-of-distribution camera views. In Fig. 8, we provide renderings of both ours and EG3D in training camera distribution and a sample of rendering with a camera placed far from training distribution. As can be observed, EG3D suffers from degraded quality

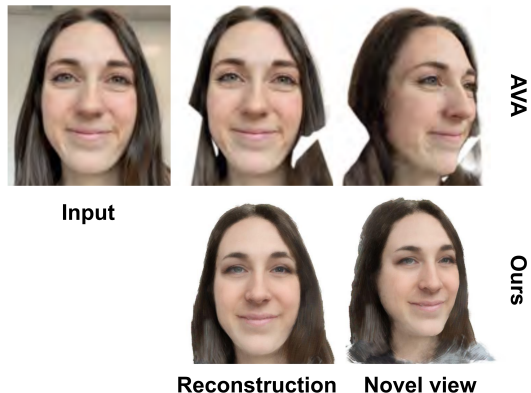


Figure 9. Here we show a comparison with Authentic Volumetric Avatar (AVA) [3]. The first column shows the input image and the second column shows the reconstruction using AVA and ours. The third column shows the novel view rendering of both methods. As can be seen, AVA suffers from truncated hair reconstruction when reconstructing identity with long hair, while ours faithfully reconstructs long hair.

because of inconsistency in the 2D super-resolution module. In contrast, our method retains the quality in out-of-training camera distribution.

In Fig. 9, we provide the monocular fitting results to a given input image. Given an input image, we use our trained model to fit to the input image and show reconstruction and novel view renderings. We obtain the camera pose for the input image using off-the-shelf face pose estimator [11]. Then we optimize the latent vector w for about 1200 iterations using reconstruction loss and simultaneously optimize the generator parameter and latent vector for another 800 iterations to get the better fitting. In the same Fig. 9, we provide a comparison with Authentic Volumetric Avatars(AVA) [3]. While AVA is not a generative model, it can generate a volumetric avatar using a phone scan, we thus qualitatively compare with this method as well. Please note AVA is not a main comparison, as this method targets expression retargeting, while we don't. [3] heavily relies on the tracked mesh to place the primitives and suffers in the case of long hair, which results in truncated results. In contrast, our method is robust to long hair, illumination, and spectacles.

Quantitative Results We provide commonly used FID scores to evaluate generative models in Tab. 1, 2. We sample 10k images for all the methods at 512×512 resolution. We can observe that EG3D [5] performs better than our model. Although perceptually samples from our model look good, we believe it suffers a higher score because of multiple reasons. Since our model has purely 3D representation without any 2D components, it is hard to achieve a very high-quality 2D synthesis for a generic model. Our

model is also discontinuous in 3D space because of our efficient sparse primitive representation. This sometimes leads to piecewise linear regular patterns. But our model is robust to out-of-training-distribution camera renderings. To evaluate this scenario, we sample 1024 identities for both our model and EG3D and render them under 2 camera distributions. One distribution is similar to training data, as we can easily obtain that by using off-the-shelf face model estimators [11]. And the other is out of this distribution, by placing the camera a bit far from the face. To evaluate the multi-view consistency of these renderings, we obtain the active Facial Active Units using state-of-the-art detectors [18] and compute the mean F1 score of all pairs of renderings for both methods. Our method clearly outperforms EG3D in this evaluation, proving the advantage of having pure 3D representation over a hybrid generator that relies on 2D super-resolution blocks.

In spite of having pure 3D representation, our method can render images faster than EG3D. We report the average time taken to render 512×512 images for both methods in Tab. 1. Please note that the evaluation was done on an NVIDIA A40 GPU. The computation time for the EG3D model is calculated using their version which volumetrically renders at 64×64 and upsamples to 512×512 . The version with 128×128 volumetric rendering is twice as slow as the one with 64×64 .

5. Limitation

Although our method can generate high-quality samples, it struggles in cases of curly hair. This is because of the discontinuous representation of the 3D scene. This is also shown numerically in the quantitative tables. We also note that this representation couldn't be trained stably in a purely adversarial manner because of the discontinuous nature of the representation.

6. Conclusion

We presented the first 3D-aware generative model with pure 3D scene representation that can be rendered at more than 43 FPS at 512×512 resolution. In addition, our method can also provide dense correspondence between samples by tracking the primitive position, orientation, and scale. Although our FID score is slightly worse than that of methods that rely on 2D super-resolution modules, our method has the advantage of robustness to camera views because of pure 3D representation, and computation speed and can provide correspondences. We believe our method takes a solid step further in building high-quality, pure 3D, and efficient 3D-aware generative models and can inspire future work in this direction.

Acknowledgements: This work was supported by the ERC Consolidator Grant 4DRely (770784).

References

- [1] Alexander W. Bergman, Petr Kellnhofer, Wang Yifan, Eric R. Chan, David B. Lindell, and Gordon Wetzstein. Generative neural articulated radiance fields. In *NeurIPS*, 2022.
- [2] Volker Blanz and Thomas Vetter. A morphable model for the synthesis of 3d faces. In *Proceedings of the 26th Annual Conference on Computer Graphics and Interactive Techniques*, SIGGRAPH '99, page 187–194, USA, 1999. ACM Press/Addison-Wesley Publishing Co.
- [3] Chen Cao, Tomas Simon, Jin Kyu Kim, Gabe Schwartz, Michael Zollhofer, Shun-Suke Saito, Stephen Lombardi, Shih-En Wei, Danielle Belko, Shou-I Yu, Yaser Sheikh, and Jason Saragih. Authentic volumetric avatars from a phone scan. *ACM Trans. Graph.*, 41(4), jul 2022.
- [4] Eric Chan, Marco Monteiro, Petr Kellnhofer, Jiajun Wu, and Gordon Wetzstein. pi-gan: Periodic implicit generative adversarial networks for 3d-aware image synthesis. In *Proc. CVPR*, 2021.
- [5] Eric R. Chan, Connor Z. Lin, Matthew A. Chan, Koki Nagano, Boxiao Pan, Shalini De Mello, Orazio Gallo, Leonidas Guibas, Jonathan Tremblay, Sameh Khamis, Tero Karras, and Gordon Wetzstein. Efficient geometry-aware 3D generative adversarial networks. In *CVPR*, 2022.
- [6] Angel X. Chang, Thomas Funkhouser, Leonidas Guibas, Pat Hanrahan, Qixing Huang, Zimo Li, Silvio Savarese, Manolis Savva, Shuran Song, Hao Su, Jianxiong Xiao, Li Yi, and Fisher Yu. ShapeNet: An Information-Rich 3D Model Repository. Technical Report arXiv:1512.03012 [cs.GR], Stanford University — Princeton University — Toyota Technological Institute at Chicago, 2015.
- [7] 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.
- [8] Blender Online Community. *Blender - a 3D modelling and rendering package*. Blender Foundation, Stichting Blender Foundation, Amsterdam, 2018.
- [9] Enric Corona, Tomas Hodan, Minh Vo, Francesc Moreno-Noguer, Chris Sweeney, Richard Newcombe, and Lingni Ma. Lisa: Learning implicit shape and appearance of hands. In *CVPR*, 2022.
- [10] Yu Deng, Jiaolong Yang, Jianfeng Xiang, and Xin Tong. Gram: Generative radiance manifolds for 3d-aware image generation. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2022.
- [11] 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.
- [12] 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.
- [13] Philipp Henzler, Niloy J Mitra, and Tobias Ritschel. Escaping plato's cave: 3d shape from adversarial rendering. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 9984–9993, 2019.
- [14] Tero Karras, Samuli Laine, and Timo Aila. A style-based generator architecture for generative adversarial networks. In *2019 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 4396–4405, 2019.
- [15] 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.
- [16] Diederik P Kingma and Jimmy Ba. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*, 2014.
- [17] Stephen Lombardi, Tomas Simon, Gabriel Schwartz, Michael Zollhofer, Yaser Sheikh, and Jason Saragih. Mixture of volumetric primitives for efficient neural rendering. *ACM Transactions on Graphics (ToG)*, 40(4):1–13, 2021.
- [18] Cheng Luo, Siyang Song, Weicheng Xie, Linlin Shen, and Hatice Gunes. Learning multi-dimensional edge feature-based au relation graph for facial action unit recognition. In *Proceedings of the Thirty-First International Joint Conference on Artificial Intelligence, IJCAI-22*, pages 1239–1246, 2022.
- [19] 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.
- [20] Michael Niemeyer and Andreas Geiger. Giraffe: Representing scenes as compositional generative neural feature fields. In *Proc. IEEE Conf. on Computer Vision and Pattern Recognition (CVPR)*, 2021.
- [21] Atsuhiko Noguchi, Xiao Sun, Stephen Lin, and Tatsuya Harada. Neural articulated radiance field. In *International Conference on Computer Vision*, 2021.
- [22] 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.
- [23] Jeong Joon Park, Peter Florence, Julian Straub, Richard Newcombe, and Steven Lovegrove. Deepsdf: Learning continuous signed distance functions for shape representation. In *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2019.
- [24] Pascal Paysan, Reinhard Knothe, Brian Amberg, Sami Romdhani, and Thomas Vetter. A 3d face model for pose and illumination invariant face recognition. In *Proceedings of the 2009 Sixth IEEE International Conference on Advanced Video and Signal Based Surveillance, AVSS '09*, page 296–301, USA, 2009. IEEE Computer Society.
- [25] Daniel Rebain, Mark J. Matthews, Kwang Yi, Dmitry Lagun, and Andrea Tagliasacchi. Lolnerf: Learn from one look. In *Computer Vision Pattern Recognition (CVPR)*, 2022.
- [26] 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.

- [27] Katja Schwarz, Axel Sauer, Michael Niemeyer, Yiyi Liao, and Andreas Geiger. Voxgraf: Fast 3d-aware image synthesis with sparse voxel grids. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2022.
- [28] Ivan Skorokhodov, Sergey Tulyakov, Yiqun Wang, and Peter Wonka. EpiGRAF: Rethinking training of 3d GANs. In Alice H. Oh, Alekh Agarwal, Danielle Belgrave, and Kyunghyun Cho, editors, *Advances in Neural Information Processing Systems*, 2022.
- [29] Jianfeng Xiang, Jiaolong Yang, Yu Deng, and Xin Tong. Gram-hd: 3d-consistent image generation at high resolution with generative radiance manifolds, 2022.
- [30] Yinghao Xu, Sida Peng, Ceyuan Yang, Yujun Shen, and Bolei Zhou. 3d-aware image synthesis via learning structural and textural representations. *arXiv preprint arXiv:2112.10759*, 2021.
- [31] Richard Zhang, Phillip Isola, Alexei A Efros, Eli Shechtman, and Oliver Wang. The unreasonable effectiveness of deep features as a perceptual metric. In *CVPR*, 2018.
- [32] Xuanmeng Zhang, Zhedong Zheng, Daiheng Gao, Bang Zhang, Pan Pan, and Yi Yang. Multi-view consistent generative adversarial networks for 3d-aware image synthesis. In *2022 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 18429–18438, 2022.