

Efficient Geometry-aware 3D Generative Adversarial Networks

Eric R. Chan ^{*†1,2}, Connor Z. Lin^{*1}, Matthew A. Chan^{*1}, Koki Nagano^{*2}, Boxiao Pan¹, Shalini De Mello², Orazio Gallo², Leonidas Guibas¹, Jonathan Tremblay², Sameh Khamis², Tero Karras², and Gordon Wetzstein¹

¹Stanford University ²NVIDIA

Abstract

Unsupervised generation of high-quality multi-view-consistent images and 3D shapes using only collections of single-view 2D photographs has been a long-standing challenge. Existing 3D GANs are either compute-intensive or make approximations that are not 3D-consistent; the former limits quality and resolution of the generated images and the latter adversely affects multi-view consistency and shape quality. In this work, we improve the computational efficiency and image quality of 3D GANs without overly relying on these approximations. We introduce an expressive hybrid explicit-implicit network architecture that, together with other design choices, synthesizes not only high-resolution multi-view-consistent images in real time but also produces high-quality 3D geometry. By decoupling feature generation and neural rendering, our framework is able to leverage state-of-the-art 2D CNN generators, such as StyleGAN2, and inherit their efficiency and expressiveness. We demonstrate state-of-the-art 3D-aware synthesis with FFHQ and AFHQ Cats, among other experiments.

1. Introduction

Generative adversarial networks (GANs) have seen immense progress, with recent models capable of generating high-resolution, photorealistic images indistinguishable from real photographs [27–29]. Current state-of-the-art GANs, however, operate in 2D only and do not explicitly model the underlying 3D scenes.

Recent work on 3D-aware GANs has begun to tackle the problem of multi-view-consistent image synthesis and, to a lesser extent, extraction of 3D shapes without being supervised on geometry or multi-view image collections. How-



Figure 1. Our 3D GAN enables synthesis of scenes, producing high-quality, multi-view-consistent renderings and detailed geometry. Our approach trains from a collection of 2D images without target-specific shape priors, ground truth 3D scans, or multi-view supervision. Please see the accompanying video for more results.

ever, the image quality and resolution of existing 3D GANs have lagged far behind those of 2D GANs. Furthermore, their 3D reconstruction quality, so far, leaves much to be desired. One of the primary reasons for this gap is the computational inefficiency of previously employed 3D generators and neural rendering architectures.

In contrast to 2D GANs, 3D GANs rely on a combination of a 3D-structure-aware inductive bias in the generator network architecture and a neural rendering engine that aims at providing view-consistent results. The inductive bias can be modeled using explicit voxel grids [14, 21, 47, 48, 68, 74] or neural implicit representations [4, 47, 49, 58]. While successful in single-scene “overfitting” scenarios, neither of these representations is suitable for training a high-resolution 3D GAN because they are simply too memory inefficient or slow. Training a 3D GAN requires rendering tens of millions of images, but state-of-the-art neural volume rendering [45] at high-resolutions with these representations is computationally infeasible. CNN-based image up-

^{*}Equal contribution.

[†]Part of the work was done during an internship at NVIDIA.

Project page: <https://github.com/NVlabs/eg3d>

sampling networks have been proposed to remedy this [49], but such an approach sacrifices view consistency and impairs the quality of the learned 3D geometry.

We introduce a novel generator architecture for unsupervised 3D representation learning from a collection of single-view 2D photographs that seeks to improve the computational efficiency of rendering while remaining true to 3D-grounded neural rendering. We achieve this goal with a two-pronged approach. First, we improve the computational efficiency of 3D-grounded rendering with a hybrid explicit-implicit 3D representation that offers significant speed and memory benefits over fully implicit or explicit approaches without compromising on expressiveness. These advantages enable our method to skirt the computational constraints that have limited the rendering resolutions and quality of previous approaches [4, 58] and forced over-reliance on image-space convolutional upsampling [49]. Second, although we use some image-space approximations that stray from the 3D-grounded rendering, we introduce a dual-discrimination strategy that maintains consistency between the neural rendering and our final output to regularize their undesirable view-inconsistent tendencies. Moreover, we introduce pose-based conditioning to our generator, which decouples pose-correlated attributes (e.g., facial expressions) for a multi-view consistent output during inference while faithfully modeling the joint distributions of pose-correlated attributes inherent in the training data.

As an additional benefit, our framework decouples feature generation from neural rendering, enabling it to directly leverage state-of-the-art 2D CNN-based feature generators, such as StyleGAN2, to generalize over spaces of 3D scenes while also benefiting from 3D multi-view-consistent neural volume rendering. Our approach not only achieves state-of-the-art qualitative and quantitative results for view-consistent 3D-aware image synthesis, but also generates high-quality 3D shapes of the synthesized scenes due to its strong 3D-structure-aware inductive bias (see Fig. 1).

Our contributions are the following:

- We introduce a tri-plane-based 3D GAN framework, which is both efficient and expressive, to enable high-resolution geometry-aware image synthesis.
- We develop a 3D GAN training strategy that promotes multi-view consistency via dual discrimination and generator pose conditioning while faithfully modeling pose-correlated attribute distributions (e.g., expressions) present in real-world datasets.
- We demonstrate state-of-the-art results for unconditional 3D-aware image synthesis on the FFHQ and AFHQ Cats datasets along with high-quality 3D geometry learned entirely from 2D in-the-wild images.




Figure 2. Neural implicit representations use fully connected layers (FC) with positional encoding (PE) to represent a scene, which can be slow to query (a). Explicit voxel grids or hybrid variants using small implicit decoders are fast to query, but scale poorly with resolution (b). Our hybrid explicit-implicit tri-plane representation (c) is fast and scales efficiently with resolution, enabling greater detail for equal capacity.

2. Related work

Neural scene representation and rendering. Emerging neural scene representations use differentiable 3D-aware representations [1, 3, 6, 8, 13, 17, 43, 44, 52, 65] that can be optimized using 2D multi-view images via neural rendering [15, 20, 24, 30, 34–37, 40, 45, 46, 50, 51, 54, 62, 63, 70–72]. Explicit representations, such as discrete voxel grids (Fig. 2b), are fast to evaluate but often incur heavy memory overheads, making them difficult to scale to high resolutions or complex scenes [38, 61]. Implicit representations, or coordinate networks (Fig. 2a), offer potential advantages in memory efficiency and scene complexity compared to discrete voxel grids by representing a scene as a continuous function (e.g., [43, 45, 52, 60, 66]). In practice, these implicit architectures use large fully connected networks that are slow to evaluate as each query requires a full pass through the network. Therefore, fully explicit and implicit representations provide complementary benefits.

Local implicit representations [3, 5, 23, 56] and hybrid explicit-implicit representations [11, 35, 39, 53] combine the benefits of both types of representations by offering computationally and memory-efficient architectures. Inspired by these ideas, we design a new hybrid explicit-implicit 3D-aware network that uses a memory-efficient tri-plane representation to explicitly store features on axis-aligned planes that are aggregated by a lightweight implicit feature decoder for efficient volume rendering (Fig. 2c). Our representation bears some resemblance to previous plane-based hybrid architectures [11, 53], but it is unique in its specific design. Our representation is key to enabling the high 3D GAN image quality that we demonstrate through efficient training comparable (in time scales) to modern 2D GANs [27].

Generative 3D-aware image synthesis. Generative adversarial networks [16] have recently achieved photorealistic

tic image quality for 2D image synthesis [25,28,29,55]. Extending these capabilities to 3D settings has started to gain momentum as well. Mesh-based approaches build on the most popular primitives used in computer graphics, but lack the expressiveness needed for high-fidelity image generation [33,64]. Voxel-based GANs directly extend the CNN generators used in 2D settings to 3D [14,21,47,48,68,74]. The high memory requirements of voxel grids and the computational burden of 3D convolutions, however, make high-resolution 3D GAN training difficult. Low-resolution 3D volume generation can be remedied with 2D CNN-based image upsampling layers [49], but without an inductive 3D bias the results often lack view consistency. Block-based sparse volume representations overcome some of these issues, but are applicable to mostly empty scenes [19,35] and difficult to generalize across scenes. As an alternative, fully implicit representation networks have been proposed for 3D scene generation [4,58], but these architectures are slow to query, which makes the GAN training inefficient, limiting the quality and resolution of generated images.

One of the primary insights of our work is that an efficient 3D GAN architecture with 3D-grounded inductive biases is crucial for successfully generating high-resolution view-consistent images and high-quality 3D shapes. Our framework achieves this in several ways. First, unlike most existing 3D GANs, we directly leverage a 2D CNN-based feature generator, i.e., StyleGAN2 [29], removing the need for inefficient 3D convolutions on explicit voxel grids. Second, our tri-plane representation allows us to leverage neural volume rendering as an inductive bias, but in a much more computationally efficient way than fully implicit 3D networks [4,45,58]. Similar to [49], we also employ 2D CNN-based upsampling after neural rendering, but our method introduces dual discrimination to avoid view inconsistencies introduced by the upsampling layers. Unlike existing StyleGAN2-based 2.5D GANs, which generate images and depth maps [59], our method works naturally for steep camera angles and in 360° viewing conditions.

The concurrently developed 3D-aware GANs StyleNeRF [18] and CIPS-3D [73] demonstrate impressive image quality. The central distinction between these and ours is that while StyleNeRF and CIPS-3D operate primarily in image-space, with less emphasis on the 3D representation, our method operates primarily in 3D. Our approach demonstrates greater view consistency, and is capable of generating high-quality 3D shapes. Furthermore, our experiments report superior FID image scores on FFHQ and AFHQ.

3. Tri-plane hybrid 3D representation

Training a high-resolution GAN requires a 3D representation that is both efficient and expressive. In this section, we introduce a new hybrid explicit-implicit tri-plane representation that offers both of these advantages. We introduce



Figure 3. A synthesized view of the multi-view *Family* scene, comparing a fully implicit Mip-NeRF representation (left), a dense voxel grid (center), and our tri-plane representation (right). Even though neither voxels nor tri-planes model view-dependent effects, they achieve high quality.

	MLP	Rel. Speed \uparrow	Rel. Mem. \downarrow
Mip-NeRF [2]	8×256	$1 \times$	$1 \times$
Voxels (hybrid)	4×128	$3.5 \times$	$0.33 \times$
Tri-plane (SSO)	4×128	$2.9 \times$	$0.32 \times$
Tri-plane (GAN)	1×64	$7.8 \times$	$0.06 \times$

Table 1. Relative speedups and memory consumption compared to Mip-NeRF. The proposed tri-plane representation is 3–8× faster than a fully implicit Mip-NeRF network and only requires a fraction of its memory. In this example, both voxel grid and tri-plane representation use an MLP-based decoder, as indicated. The number of voxels is chosen to match the total parameters of the tri-plane representation, thus the resolution is relatively low and the memory footprint lower than Mip-NeRF. In the SSO experiment (Fig. 3), we used a larger decoder for the tri-plane representation than for the GAN experiments discussed in Sec. 4 to optimize expressiveness over speed for this experiment.

the representation in this section for a single-scene overfitting (SSO) experiment, before discussing how it is integrated in our GAN framework in the next section.

In the tri-plane formulation, we align our explicit features along three axis-aligned orthogonal feature planes, each with a resolution of $N \times N \times C$ (Fig. 2c) with N being spatial resolution and C the number of channels. We query any 3D position $x \in \mathbb{R}^3$ by projecting it 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. An additional lightweight decoder network, implemented as a small MLP, interprets the aggregated 3D features F as color and density. These quantities are rendered into RGB images using (neural) volume rendering [41,45].

The primary advantage of this hybrid representation is efficiency—by keeping the decoder small and shifting the bulk of the expressive power into the explicit features, we reduce the computational cost of neural rendering compared to fully implicit MLP architectures [2,45] without losing




Figure 4. Our 3D GAN framework comprises several parts: a pose-conditioned StyleGAN2-based feature generator and mapping network, a tri-plane 3D representation with a lightweight feature decoder, a neural volume renderer, a super-resolution module, and a pose-conditioned StyleGAN2 discriminator with dual discrimination. This architecture elegantly decouples feature generation and neural rendering, allowing the use of a powerful StyleGAN2 generator for 3D scene generalization. Moreover, the lightweight 3D tri-plane representation is both expressive and efficient in enabling high-quality 3D-aware view synthesis in real-time.

expressiveness. To validate that the tri-plane representation is compact yet sufficiently expressive, we evaluate it with a common novel-view synthesis setup. For this purpose, we directly optimize the features of the planes and the weights of the decoder to fit 360° views of a scene from the Tanks & Temples dataset [31] (Fig. 3). In this experiment, we use feature planes of resolution $N = 512$ and channels $C = 48$, paired with an MLP of four layers of 128 hidden units each and a Fourier feature encoding [66]. We compare the results against a dense feature volume of equal capacity. For reference, we include comparisons to a state-of-the-art fully implicit 3D representation [2]. Fig. 3 and Tab. 1 demonstrate that the tri-plane representation is capable of representing this complex scene, albeit without view-dependent effects, outperforming dense feature volume representations [38, 61] and fully implicit representations [45] in terms of PSNR and SSIM, while offering considerable advantages in computation and memory efficiency. For a side length of N features, tri-planes scale with $O(N^2)$ rather than $O(N^3)$ as dense voxels do, which means for equal capacity and memory, the tri-plane representation can use higher resolution features and capture greater detail. Finally, our tri-plane representation has one other key advantage over these alternatives: the feature planes can be generated with an off-the-shelf 2D CNN-based generator, enabling generalization across 3D representations using the GAN framework discussed next.

4. 3D GAN framework

Armed with an efficient and expressive 3D representation, we train a 3D GAN for geometry-aware image synthesis from 2D photographs, without any explicit 3D or multi-view supervision. We associate each training image with a set of camera intrinsics and extrinsics using off-the-shelf pose detectors [10, 32]; see the supplement for details.

Fig. 4 gives an overview of our network architecture. We use the tri-plane representation introduced in the last sec-

tion to efficiently render images through neural volume rendering, but make a number of modifications to adapt this representation to the 3D GAN setting. Unlike in the SSO experiment, where the features of the planes were directly optimized from the multiple input views, for the GAN setting we generate the tri-plane features, each containing 32 channels, with the help of a 2D convolutional StyleGAN2 backbone (Sec. 4.1). Instead of producing an RGB image, in the GAN setting our neural renderer aggregates features from each of the 32-channel tri-planes and predicts 32-channel feature images from a given camera pose. This is followed by a “super-resolution” module to upsample and refine these raw neurally rendered images (Sec. 4.2). The generated images are critiqued by a slightly modified StyleGAN2 discriminator (Sec. 4.3). The entire pipeline is trained end-to-end from random initialization, using the non-saturating GAN loss function [16] with R1 regularization [42], following the training scheme in StyleGAN2 [29]. To speed training, we use a two-stage training strategy in which we train with a reduced (64^2) neural rendering resolution followed by a short fine-tuning period at full (128^2) neural rendering resolution. Additional experiments found that regularization to encourage smoothness of the density field helped reduce artifacts in 3D shapes. The following sections discuss major components of our framework in detail. For additional descriptions, implementation details, and hyperparameters, please see the supplement.

4.1. CNN generator backbone and rendering

The features of the tri-plane representation, when used in our GAN setting, are generated by a StyleGAN2 CNN generator. The random latent code and camera parameters are first processed by a mapping network to yield an intermediate latent code which then modulates the convolution kernels of a separate synthesis network.

We change the output shape of the StyleGAN2 backbone such that, rather than producing a three-channel RGB image, we produce a $256 \times 256 \times 96$ feature image. This



Figure 5. Dual discrimination ensures that the raw neural rendering I_{RGB} and super-resolved output I_{RGB}^+ maintain consistency, enabling high-resolution and multi-view-consistent rendering.

feature image is split channel-wise and reshaped to form three 32-channel planes (see Fig. 4). We choose StyleGAN2 for predicting the tri-plane features because it is a well-understood and efficient architecture achieving state-of-the-art results for 2D image synthesis. Furthermore, our model inherits many of the desirable properties of StyleGAN: a well-behaved latent space that enables style-mixing and latent-space interpolation (see Sec. 5 and supplement).

We sample features from the tri-planes, aggregate by summation, and process the aggregated features with a lightweight decoder, as described in Sec. 3. Our decoder is a multi-layer perceptron with a single hidden layer of 64 units and softplus activation functions. The MLP does *not* use a positional encoding, coordinate inputs, or view-direction inputs. This hybrid representation can be queried for continuous coordinates and outputs a scalar density σ as well as a 32-channel feature, both of which are then processed by a neural volume renderer to project the 3D feature volume into a 2D feature image.

Volume rendering [41] is implemented using two-pass importance sampling as in [45]. Following [49], volume rendering in our GAN framework produces feature images, rather than RGB images, because feature images contain more information that can be effectively utilized for the image-space refinement described next. For the majority of the experiments reported in this manuscript, we render 32-channel feature images I_F at a resolution of 128^2 , with 96 total depth samples per ray.

4.2. Super resolution

Although the tri-plane representation is significantly more computationally efficient than previous approaches, it is still too slow to natively train or render at high resolutions while maintaining interactive framerates. We thus perform volume rendering at a moderate resolution (e.g., 128^2) and rely upon image-space convolutions to upsample the neural rendering to the final image size of 256^2 or 512^2 .

Our super resolution module is composed of two blocks of StyleGAN2-modulated convolutional layers that upsample and refine the 32-channel feature image I_F into the final RGB image I_{RGB}^+ . We disable per-pixel noise inputs to reduce texture sticking [27] and reuse the mapping network of the backbone to modulate these layers.

4.3. Dual discrimination

As in standard 2D GAN training, the resulting renderings are critiqued by a 2D convolutional discriminator. We use a StyleGAN2 discriminator with two modifications.

First, we introduce *dual discrimination* as a method to avoid multi-view inconsistency issues observed in prior work [47, 49]. For this purpose, we interpret the first three feature channels of a neurally rendered feature image I_F as a low-resolution RGB image I_{RGB} . Intuitively, dual discrimination then ensures consistency between I_{RGB} and the super-resolved image I_{RGB}^+ . This is achieved by bilinearly upsampling I_{RGB} to the same resolution as I_{RGB}^+ and concatenating the results to form a six-channel image (see Fig. 4). The real images fed into the discriminator are also processed by concatenating each of them with an appropriately blurred copy of itself. We discriminate over these six-channel images instead of the three-channel images traditionally seen in GAN discriminators.

Dual discrimination not only encourages the final output to match the distribution of real images, but also offers additional effects: it encourages the neural rendering to match the distribution of downsampled real images; and it encourages the super-resolved images to be consistent with the neural rendering (see Fig. 5). The second point importantly allows us to leverage effective image-space super-resolution layers without introducing view-inconsistency artifacts.

Second, we make the discriminator aware of the camera poses from which the generated images are rendered. Specifically, following the conditional strategy from StyleGAN2-ADA [26], we pass the rendering camera intrinsics and extrinsics matrices (collectively \mathbf{P}) to the discriminator as a conditioning label. We find that this conditioning introduces additional information that guides the generator to learn correct 3D priors. We provide additional studies in the supplement showing the effect of this discriminator conditioning and the robustness of our framework to high levels of noise in the input camera poses.

4.4. Modeling pose-correlated attributes

Most real-world datasets like FFHQ include biases that correlate camera poses with other attributes (e.g., facial expressions), and naively handling them leads to view inconsistent results. For example, the camera angle with respect to a person’s face is correlated with smiling (see supplement). While faithfully modeling such attribute correlations inherent in the dataset is important for reproducing the best image quality, such unwanted attributes need to be decoupled during inference for multi-view consistent synthesis. Related work has been successful at being view consistent [4, 58, 59] or modeling pose-appearance correlations [47, 49], but cannot achieve both simultaneously.

We introduce *generator pose conditioning* as a means to model and decouple correlations between pose and other



Figure 6. Curated examples at 512^2 , synthesized by models trained with FFHQ [28] and AFHQv2 Cats [7]




Figure 7. Qualitative comparison between GIRAFFE, π -GAN, Lifting StyleGAN, ours, with FFHQ at 256^2 . Shapes are iso-surfaces extracted from the density field using marching cubes. We inspected the underlying 3D representations of GIRAFFE and found that its over-reliance on image-space approximations significantly harms the learning of the 3D geometry.

attributes observed in the training images. To this end, we provide the backbone mapping network not only a latent code vector z , but also the camera parameters \mathbf{P} as input, following the conditional generation strategy in [26]. By giving the backbone knowledge of the rendering camera position, we allow the target view to influence scene synthesis.

During training, pose conditioning allows the generator to model pose-dependent biases implicit to the dataset, allowing our model to faithfully reproduce the image distributions in the dataset. To prevent the scene from shifting with camera pose during inference, we condition the generator on a fixed camera pose when rendering from a moving camera trajectory. We noticed that always conditioning the generator with the rendering camera pose can lead to degenerate solutions where the GAN produces 2D billboards angled towards the camera (see supplement). To prevent this,

we randomly swap the conditioning pose in \mathbf{P} with another random pose with 50% probability during training.

5. Experiments and results

Datasets. We compare methods on the task of unconditional 3D-aware generation with FFHQ [28], a real-world human face dataset, and AFHQv2 Cats [7,27], a small, real-world cat face dataset. We augment both datasets with horizontal flips and use off-the-shelf pose estimators [10,32] to extract approximate camera extrinsics. For all methods on AFHQv2, we apply transfer learning [26] from corresponding FFHQ checkpoints; for our method on AFHQv2 512^2 , we additionally use adaptive data augmentation [26]. For more results, please see the accompanying video.

	FID↓	FFHQ			Cats
		ID↑	Depth↓	Pose↓	FID↓
GIRAFFE 256 ²	31.5	0.64	0.94	.089	16.1
π -GAN 128 ²	29.9	0.67	0.44	.021	16.0
Lift. SG 256 ²	29.8	0.58	0.40	.023	—
Ours 256 ²	4.8	0.76	0.31	.005	3.88
Ours 512 ²	4.7	0.77	0.39	.005	2.77[†]

Table 2. Quantitative evaluation using FID, identity consistency (ID), depth accuracy, and pose accuracy for FFHQ and AFHQ Cats. Labelled is the image resolution of training and evaluation. [†] Trained with adaptive data augmentation [26].

5.1. Comparisons

Baselines. We compare our methods against three state-of-the-art methods for 3D-aware image synthesis: π -GAN [4], GIRAFFE [49], and Lifting StyleGAN [59].

Qualitative results. Fig. 6 presents selected examples synthesized by our model with FFHQ and AFHQ at a resolution of 512², highlighting the image quality, view-consistency, and diversity of outputs produced by our method. Fig. 7 provides a qualitative comparison against baselines. While GIRAFFE synthesizes high-quality images, reliance on view-inconsistent convolutions produces poor-quality shapes and identity shift—note the hairline inconsistency between rendered views. π -GAN and Lifting StyleGAN generate adequate shapes and images but both struggle with photorealism and in capturing detailed shapes.

Our method synthesizes not only images that are higher quality and more view-consistent but also higher-fidelity 3D geometry as seen in the detailed glasses and hair strands.

Quantitative evaluations. Table 2 provides quantitative metrics comparing the proposed approach against baselines. We measure image quality with Fréchet Inception Distance (FID) [22] between 50k generated images and all available real images. We evaluate shape quality by calculating MSE against pseudo-ground-truth depth-maps (Depth) and poses (Pose) estimated from synthesized images by [10]; a similar evaluation was introduced by [59]. We assess multi-view facial identity consistency (ID) by calculating the mean Arcface [9] cosine similarity score between pairs of views of the same synthesized face rendered from random camera poses. Additional evaluation details are provided in the supplement. Our model demonstrates significant improvements in FID across both datasets, bringing the 3D GAN to near the same level as StyleGAN2 512² (2.97 for FFHQ [29] and 2.99 for Cats [26]) while also maintaining state-of-the-art view consistency, geometry quality, and pose accuracy.

Runtime. Table 3 compares rendering speed at inference running on a single NVIDIA RTX 3090 GPU. Our end-to-end approach achieves real-time framerates at 512² final resolution with 128² neural rendering resolution and

Res.	GIRAFFE	π -GAN	Lift. SG	Ours	Ours + TC
256 ²	181	5	51	27	36
512 ²	161	1	—	26	35

Table 3. Runtime in frames per second at different rendering resolutions. We compare variants of our approach with and without tri-plane caching (TC). Run on a single RTX 3090 GPU.

	FID ↓	FACS Smile Std. ↓
Naive model	5.5	0.069
+ DD	6.5	0.054
+ DD, GPC (ours)	4.7	0.031

Table 4. Dual-discrimination (DD) improves multi-view expression consistency but hurts the model’s ability to capture pose-correlated attributes for image quality. Adding generator pose conditioning (GPC) allows the model to improve upon both aspects. Reported at 512², with FFHQ.

96 total depth samples per ray, suitable for applications such as real-time visualization. When rendering consecutive frames of a static scene, we need not regenerate the tri-plane features every frame; caching the generated features is a simple tweak that improves render speed. The proposed approach is significantly faster than fully implicit methods like π -GAN [4]. Although it is not as fast as Lifting StyleGAN [59] and GIRAFFE [49], we believe major improvements in image quality, geometry quality, and view-consistency outweigh the increased compute cost.

5.2. Ablation study

Without dual discrimination, generated images can include multi-view inconsistencies due to the unconstrained image-space super-resolution layers. We measure this effect quantitatively by extracting smile-related Facial Action Coding System (FACS) [12] coefficients from videos produced by models with and without dual discrimination, using a proprietary facial tracker. We measure the standard deviation of smile coefficients for the same scene across video frames. A view-consistent scene should exhibit little expression shift and thus produce little variation in smile coefficients. This is validated in Table 4 showing that introducing dual discrimination (second row) reduces the smile coefficient variation versus the naive model (first row), indicating improved expression consistency. However, dual discrimination also reduces image quality as seen by the slightly worse FID score, perhaps because the model is restricted from reproducing the pose-correlated attribute biases in the FFHQ dataset. By adding generator pose conditioning (third row), we allow the generator to faithfully model pose-correlated attributes while decoupling them at inference, leading to both the best FID score and view-consistent results.



Figure 8. Style-mixing [27–29] with FFHQ 512².

5.3. Applications

Style mixing. Since our 3D representation is designed with the StyleGAN2 backbone from the ground up, it inherits the well-studied properties of the StyleGAN2 latent space, allowing us to do semantic image manipulations. Fig. 8 shows our method’s results for style mixing [27–29].

Single-view 3D reconstruction. Fig. 9 shows the application of our learned latent space for single-view 3D reconstruction. We use pivotal tuning inversion (PTI) [57] to fit test images. The learned 3D prior over FFHQ enables surprisingly high-quality single-view geometry recovery. Further exploration of few-shot 3D reconstruction and novel-view-synthesis may prove a fruitful avenue for future work.

6. Discussion

Limitations and future work. Although our shapes show significant improvements over those generated by previous 3D-aware GANs, they may still contain artifacts and lack finer details, such as individual teeth. To further improve the quality of the learned shapes, we could instill a stronger geometry prior or regularize the density component of the radiance field following methods proposed by [51, 67, 69].

Our model requires knowledge of the camera pose distribution of the dataset. Although prior work has proposed learning the pose distribution on the fly [49], others have noticed such methods can diverge [18], so it would be fruitful to explore this direction further. Pose conditioning aids the generator in decoupling appearance from pose, but still does not fully disentangle the two. Furthermore, ambiguities that can be explained by geometry remain unresolved. For example, by creating concave eye sockets, the generator creates the illusion of eyes that “follow” the camera, an incorrect interpretation, though the renderings are view-consistent and reflect the underlying geometry.

We used StyleGAN 2, but other 2D backbones may find success in our framework. Alternative backbones, such as



Figure 9. We use PTI [57] to fit a target image and recover the underlying 3D shape. Target (left); reconstructed image (center); reconstructed shape (right). From a model trained on FFHQ 512².

as image-to-image translation or Transformer-based models, could enable new applications in conditional synthesis.

Ethical considerations. The single-view 3D reconstruction or style mixing applications could be misused for generating edited imagery of real people. Such misuse of image synthesis techniques poses a societal threat, and we do not condone using our work with the intent of spreading misinformation or tarnishing reputation. We also recognize a potential lack of diversity in our faces results, stemming from implicit biases of the datasets we process.

Conclusion. By combining an efficient explicit–implicit neural representation with an expressive pose-aware convolutional generator and a dual discriminator, our approach takes significant steps towards photorealistic 3D-aware image synthesis and high-quality unsupervised shape generation. This may enable rapid prototyping of 3D models, more controllable image synthesis, and novel techniques for shape reconstruction from temporal data.

Acknowledgements

We thank David Luebke, Jan Kautz, Jaewoo Seo, Jonathan Granskog, Simon Yuen, Alex Evans, Stan Birchfield, Alexander Bergman, and Joy Hsu for feedback on drafts, Alex Chan, Giap Nguyen, and Trevor Chan for help with diagrams, and Colette Kress and Bryan Catanzaro for allowing use of their photographs. This project was in part supported by Stanford HAI and a Samsung GRO. Koki Nagano and Eric Chan were partially supported by DARPA’s Semantic Forensics (SemaFor) contract (HR0011-20-3-0005). The views and conclusions contained in this document are those of the authors and should not be interpreted as representing the official policies, either expressed or implied, of the U.S. Government. Distribution Statement “A” (Approved for Public Release, Distribution Unlimited).

References

- [1] Matan Atzmon and Yaron Lipman. SAL: Sign agnostic learning of shapes from raw data. In IEEE Conference on Computer Vision and Pattern Recognition (CVPR), 2020. 2
- [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. In ICCV, 2021. 3, 4
- [3] Rohan Chabra, Jan Eric Lenssen, Eddy Ilg, Tanner Schmidt, Julian Straub, Steven Lovegrove, and Richard Newcombe. Deep local shapes: Learning local SDF priors for detailed 3D reconstruction. In European Conference on Computer Vision (ECCV), 2020. 2
- [4] Eric R Chan, Marco Monteiro, Petr Kellnhofer, Jiajun Wu, and Gordon Wetzstein. pi-GAN: Periodic implicit generative adversarial networks for 3D-aware image synthesis. In IEEE Conference on Computer Vision and Pattern Recognition (CVPR), 2021. 1, 2, 3, 5, 7
- [5] Yinbo Chen, Sifei Liu, and Xiaolong Wang. Learning continuous image representation with local implicit image function. In IEEE Conference on Computer Vision and Pattern Recognition (CVPR), 2021. 2
- [6] Zhiqin Chen and Hao Zhang. Learning implicit fields for generative shape modeling. In IEEE Conference on Computer Vision and Pattern Recognition (CVPR), 2019. 2
- [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. 6
- [8] Thomas Davies, Derek Nowrouzezahrai, and Alec Jacobson. Overfit neural networks as a compact shape representation. arXiv preprint arXiv:2009.09808, 2020. 2
- [9] Jiankang Deng, Jia Guo, Xue Niannan, and Stefanos Zafeiriou. Arcface: Additive angular margin loss for deep face recognition. In CVPR, 2019. 7
- [10] 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. 4, 6, 7
- [11] Terrance DeVries, Miguel Angel Bautista, Nitish Srivastava, Graham W. Taylor, and Joshua M. Susskind. Unconstrained scene generation with locally conditioned radiance fields. arXiv preprint arXiv:2104.00670, 2021. 2
- [12] Paul Ekman and Wallace V. Friesen. Facial Action Coding System: A Technique for the Measurement of Facial Movement. Consulting Psychologists Press, 1978. 7
- [13] SM Ali Eslami, Danilo Jimenez Rezende, Frederic Besse, Fabio Viola, Ari S Morcos, Marta Garnelo, Avraham Ruderman, Andrei A Rusu, Ivo Danihelka, Karol Gregor, et al. Neural scene representation and rendering. Science, 2018. 2
- [14] Matheus Gadelha, Subhransu Maji, and Rui Wang. 3D shape induction from 2D views of multiple objects. In International Conference on 3D Vision, 2017. 1, 3
- [15] Stephan J Garbin, Marek Kowalski, Matthew Johnson, Jamie Shotton, and Julien Valentin. FastNeRF: High-fidelity neural rendering at 200fps. arXiv preprint arXiv:2103.10380, 2021. 2
- [16] 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 (NeurIPS), 2014. 2, 4
- [17] Amos Gropp, Lior Yariv, Niv Haim, Matan Atzmon, and Yaron Lipman. Implicit geometric regularization for learning shapes. In International Conference on Machine Learning (ICML), 2020. 2
- [18] Jiatao Gu, Lingjie Liu, Peng Wang, and Christian Theobalt. StyleNeRF: A style-based 3D-aware generator for high-resolution image synthesis. arXiv preprint arXiv:2110.08985, 2021. 3, 8
- [19] Zekun Hao, Arun Mallya, Serge Belongie, and Ming-Yu Liu. GANcraft: Unsupervised 3D neural rendering of minecraft worlds. In IEEE International Conference on Computer Vision (ICCV), 2021. 3
- [20] Peter Hedman, Pratul P. Srinivasan, Ben Mildenhall, Jonathan T. Barron, and Paul Debevec. Baking neural radiance fields for real-time view synthesis. In IEEE International Conference on Computer Vision (ICCV), 2021. 2
- [21] Philipp Henzler, Niloy J Mitra, and Tobias Ritschel. Escaping Plato’s cave: 3D shape from adversarial rendering. In IEEE International Conference on Computer Vision (ICCV), 2019. 1, 3
- [22] Martin Heusel, Hubert Ramsauer, Thomas Unterthiner, Bernhard Nessler, Günter Klambauer, and Sepp Hochreiter. GANs trained by a two time-scale update rule converge to a nash equilibrium. In Advances in Neural Information Processing Systems (NeurIPS), 2017. 7
- [23] Chiyu Jiang, Avneesh Sud, Ameesh Makadia, Jingwei Huang, Matthias Nießner, and Thomas Funkhouser. Local implicit grid representations for 3D scenes. In IEEE Conference on Computer Vision and Pattern Recognition (CVPR), 2020. 2
- [24] Yue Jiang, Dantong Ji, Zhizhong Han, and Matthias Zwicker. SDFDiff: Differentiable rendering of signed distance fields for 3D shape optimization. In IEEE Conference on Computer Vision and Pattern Recognition (CVPR), 2020. 2
- [25] Tero Karras, Timo Aila, Samuli Laine, and Jaakko Lehtinen. Progressive growing of GANs for improved quality, stability, and variation. In International Conference on Learning Representations (ICLR), 2018. 3
- [26] 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 (NeurIPS), 2020. 5, 6, 7
- [27] Tero Karras, Miika Aittala, Samuli Laine, Erik Härkönen, Janne Hellsten, Jaakko Lehtinen, and Timo Aila. Alias-free generative adversarial networks. In Advances in Neural Information Processing Systems (NeurIPS), 2021. 1, 2, 5, 6, 8
- [28] Tero Karras, Samuli Laine, and Timo Aila. A style-based generator architecture for generative adversarial net-

- works. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2019. 1, 3, 6, 8
- [29] Tero Karras, Samuli Laine, Miika Aittala, Janne Hellsten, Jaakko Lehtinen, and Timo Aila. Analyzing and improving the image quality of StyleGAN. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2020. 1, 3, 4, 7, 8
- [30] Petr Kellnhofer, Lars Jebe, Andrew Jones, Ryan Spicer, Kari Pulli, and Gordon Wetzstein. Neural lumigraph rendering. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2021. 2
- [31] Arno Knapitsch, Jaesik Park, Qian-Yi Zhou, and Vladlen Koltun. Tanks and temples: Benchmarking large-scale scene reconstruction. *ACM Transactions on Graphics*, 36(4), 2017. 4
- [32] Taehee Brad Lee. Cat hipsterizer, 2018. https://github.com/kairess/cat_hipsterizer. 4, 6
- [33] Yiyi Liao, Katja Schwarz, Lars Mescheder, and Andreas Geiger. Towards unsupervised learning of generative models for 3D controllable image synthesis. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2020. 3
- [34] David B Lindell, Julien NP Martel, and Gordon Wetzstein. AutoInt: Automatic integration for fast neural volume rendering. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2021. 2
- [35] Lingjie Liu, Jiatao Gu, Kyaw Zaw Lin, Tat-Seng Chua, and Christian Theobalt. Neural sparse voxel fields. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2020. 2, 3
- [36] Shichen Liu, Shunsuke Saito, Weikai Chen, and Hao Li. Learning to infer implicit surfaces without 3D supervision. *arXiv preprint arXiv:1911.00767*, 2019. 2
- [37] Shaohui Liu, Yinda Zhang, Songyou Peng, Boxin Shi, Marc Pollefeys, and Zhaopeng Cui. DIST: Rendering deep implicit signed distance function with differentiable sphere tracing. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2020. 2
- [38] Stephen Lombardi, Tomas Simon, Jason Saragih, Gabriel Schwartz, Andreas Lehrmann, and Yaser Sheikh. Neural volumes: Learning dynamic renderable volumes from images. *ACM Transactions on Graphics (SIGGRAPH)*, 2019. 2, 4
- [39] Julien N.P. Martel, David B. Lindell, Connor Z. Lin, Eric R. Chan, Marco Monteiro, and Gordon Wetzstein. ACORN: Adaptive coordinate networks for neural representation. *ACM Transactions on Graphics (SIGGRAPH)*, 2021. 2
- [40] Ricardo Martin-Brualla, Noha Radwan, Mehdi S. M. Sajjadi, Jonathan T. Barron, Alexey Dosovitskiy, and Daniel Duckworth. NeRF in the wild: Neural radiance fields for unconstrained photo collections. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2021. 2
- [41] N. Max. Optical models for direct volume rendering. *IEEE Transactions on Visualization and Computer Graphics (TVCG)*, 1995. 3, 5
- [42] Lars Mescheder, Andreas Geiger, and Sebastian Nowozin. Which training methods for GANs do actually converge? In *International Conference on Machine Learning (ICML)*, 2018. 4
- [43] Lars Mescheder, Michael Oechsle, Michael Niemeyer, Sebastian Nowozin, and Andreas Geiger. Occupancy networks: Learning 3D reconstruction in function space. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2019. 2
- [44] Mateusz Michalkiewicz, Jhony K Pontes, Dominic Jack, Mahsa Baktashmotlagh, and Anders Eriksson. Implicit surface representations as layers in neural networks. In *IEEE International Conference on Computer Vision (ICCV)*, 2019. 2
- [45] Ben Mildenhall, Pratul P Srinivasan, Matthew Tancik, Jonathan T Barron, Ravi Ramamoorthi, and Ren Ng. NeRF: Representing scenes as neural radiance fields for view synthesis. In *European Conference on Computer Vision (ECCV)*, 2020. 1, 2, 3, 4, 5
- [46] 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. 2
- [47] Thu Nguyen-Phuoc, Chuan Li, Lucas Theis, Christian Richardt, and Yong-Liang Yang. HoloGAN: Unsupervised learning of 3D representations from natural images. In *IEEE International Conference on Computer Vision (ICCV)*, 2019. 1, 3, 5
- [48] 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 (NeurIPS)*, 2020. 1, 3
- [49] Michael Niemeyer and Andreas Geiger. GIRAFFE: Representing scenes as compositional generative neural feature fields. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2021. 1, 2, 3, 5, 7, 8
- [50] Michael Niemeyer, Lars Mescheder, Michael Oechsle, and Andreas Geiger. Differentiable volumetric rendering: Learning implicit 3d representations without 3d supervision. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2020. 2
- [51] Michael Oechsle, Songyou Peng, and Andreas Geiger. UNISURF: Unifying neural implicit surfaces and radiance fields for multi-view reconstruction. In *IEEE International Conference on Computer Vision (ICCV)*, 2021. 2, 8
- [52] Jeong Joon Park, Peter Florence, Julian Straub, Richard Newcombe, and Steven Lovegrove. DeepSDF: Learning continuous signed distance functions for shape representation. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2019. 2
- [53] Songyou Peng, Michael Niemeyer, Lars Mescheder, Marc Pollefeys, and Andreas Geiger. Convolutional occupancy networks. In *European Conference on Computer Vision (ECCV)*, 2020. 2
- [54] Albert Pumarola, Enric Corona, Gerard Pons-Moll, and Francesc Moreno-Noguer. D-NeRF: Neural radiance fields for dynamic scenes. *arXiv preprint arXiv:2011.13961*, 2020. 2

- [55] Alec Radford, Luke Metz, and Soumith Chintala. Unsupervised representation learning with deep convolutional generative adversarial networks. In International Conference on Learning Representations (ICLR), 2016. 3
- [56] Christian Reiser, Songyou Peng, Yiyi Liao, and Andreas Geiger. KiloNeRF: Speeding up neural radiance fields with thousands of tiny MLPs. In IEEE International Conference on Computer Vision (ICCV), 2021. 2
- [57] Daniel Roich, Ron Mokady, Amit H Bermano, and Daniel Cohen-Or. Pivotal tuning for latent-based editing of real images. arXiv preprint arXiv:2106.05744, 2021. 8
- [58] 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. 1, 2, 3, 5
- [59] Yichun Shi, Divyansh Aggarwal, and Anil K Jain. Lifting 2D stylegan for 3D-aware face generation. In IEEE Conference on Computer Vision and Pattern Recognition (CVPR), 2021. 3, 5, 7
- [60] Vincent Sitzmann, Julien N.P. Martel, Alexander W. Bergman, David B. Lindell, and Gordon Wetzstein. Implicit neural representations with periodic activation functions. In Advances in Neural Information Processing Systems (NeurIPS), 2020. 2
- [61] Vincent Sitzmann, Justus Thies, Felix Heide, Matthias Nießner, Gordon Wetzstein, and Michael Zollhöfer. DeepVoxels: Learning persistent 3D feature embeddings. In IEEE Conference on Computer Vision and Pattern Recognition (CVPR), 2019. 2, 4
- [62] Vincent Sitzmann, Michael Zollhöfer, and Gordon Wetzstein. Scene representation networks: Continuous 3D-structure-aware neural scene representations. In Advances in Neural Information Processing Systems (NeurIPS), 2019. 2
- [63] Pratul P. Srinivasan, Boyang Deng, Xiuming Zhang, Matthew Tancik, Ben Mildenhall, and Jonathan T. Barron. NeRV: Neural reflectance and visibility fields for relighting and view synthesis. In IEEE Conference on Computer Vision and Pattern Recognition (CVPR), 2021. 2
- [64] Attila Szabó, Givi Meishvili, and Paolo Favaro. Unsupervised generative 3D shape learning from natural images. arXiv preprint arXiv:1910.00287, 2019. 3
- [65] Towaki Takikawa, Joey Litalien, Kangxue Yin, Karsten Kreis, Charles Loop, Derek Nowrouzezahrai, Alec Jacobson, Morgan McGuire, and Sanja Fidler. Neural geometric level of detail: Real-time rendering with implicit 3D shapes. In IEEE Conference on Computer Vision and Pattern Recognition (CVPR), 2021. 2
- [66] Matthew Tancik, Pratul P. Srinivasan, Ben Mildenhall, Sara Fridovich-Keil, Nithin Raghavan, Utkarsh Singhal, Ravi Ramamoorthi, Jonathan T. Barron, and Ren Ng. Fourier features let networks learn high frequency functions in low dimensional domains. In Advances in Neural Information Processing Systems (NeurIPS), 2020. 2, 4
- [67] Peng Wang, Lingjie Liu, Yuan Liu, Christian Theobalt, Taku Komura, and Wenping Wang. NeuS: Learning neural implicit surfaces by volume rendering for multi-view reconstruction. Advances in Neural Information Processing Systems (NeurIPS), 2021. 8
- [68] Jiajun Wu, Chengkai Zhang, Tianfan Xue, William T. Freeman, and Joshua B. Tenenbaum. Learning a probabilistic latent space of object shapes via 3D generative-adversarial modeling. In Advances in Neural Information Processing Systems (NeurIPS), 2016. 1, 3
- [69] Lior Yariv, Jiatao Gu, Yoni Kasten, and Yaron Lipman. Volume rendering of neural implicit surfaces. arXiv preprint arXiv:2106.12052, 2021. 8
- [70] Lior Yariv, Yoni Kasten, Dror Moran, Meirav Galun, Matan Atzmon, Ronen Basri, and Yaron Lipman. Multiview neural surface reconstruction by disentangling geometry and appearance. In Advances in Neural Information Processing Systems (NeurIPS), 2020. 2
- [71] Alex Yu, Ruilong Li, Matthew Tancik, Hao Li, Ren Ng, and Angjoo Kanazawa. PlenOctrees for real-time rendering of neural radiance fields. In IEEE International Conference on Computer Vision (ICCV), 2021. 2
- [72] Kai Zhang, Gernot Riegler, Noah Snavely, and Vladlen Koltun. Nerf++: Analyzing and improving neural radiance fields. arXiv preprint arXiv:2010.07492, 2020. 2
- [73] Peng Zhou, Lingxi Xie, Bingbing Ni, and Qi Tian. CIPS-3D: A 3D-Aware Generator of GANs Based on Conditionally-Independent Pixel Synthesis. arXiv preprint arXiv:2110.09788, 2021. 3
- [74] Jun-Yan Zhu, Zhoutong Zhang, Chengkai Zhang, Jiajun Wu, Antonio Torralba, Joshua B. Tenenbaum, and William T. Freeman. Visual object networks: Image generation with disentangled 3D representations. In Advances in Neural Information Processing Systems (NeurIPS), 2018. 1, 3