

3D Neural Field Generation using Triplane Diffusion

J. Ryan Shue^{*1} Eric Ryan Chan^{*2} Ryan Po^{*2} Zachary Ankner^{*3,4} Jiajun Wu² Gordon Wetzstein²

¹Milton Academy ²Stanford University ³Massachusetts Institute of Technology ⁴MosaicML

Abstract

Diffusion models have emerged as the state-of-the-art for image generation, among other tasks. Here, we present an efficient diffusion-based model for 3D-aware generation of neural fields. Our approach pre-processes training data, such as ShapeNet meshes, by converting them to continuous occupancy fields and factoring them into a set of axis-aligned triplane feature representations. Thus, our 3D training scenes are all represented by 2D feature planes, and we can directly train existing 2D diffusion models on these representations to generate 3D neural fields with high quality and diversity, outperforming alternative approaches to 3D-aware generation. Our approach requires essential modifications to existing triplane factorization pipelines to make the resulting features easy to learn for the diffusion model. We demonstrate state-of-the-art results on 3D generation on several object classes from ShapeNet.




Figure 1. Our method leverages existing 2D diffusion models for 3D shape generation using hybrid explicit–implicit neural representations. Top: triplane-based 3D shape diffusion process using our framework. Bottom: Interpolation between generated shapes.

1. Introduction

Diffusion models have seen rapid progress, setting state-of-the-art (SOTA) performance across a variety of image generation tasks. While most diffusion methods model 2D images, recent work [2, 14, 42, 86] has attempted to develop denoising methods for 3D shape generation. These 3D diffusion methods operate on discrete point clouds and, while successful, exhibit limited quality and resolution.

In contrast to 2D diffusion, which directly leverages the image as the target for the diffusion process, it is not directly obvious how to construct such 2D targets in the case of 3D diffusion. Interestingly, recent work on 3D-aware generative adversarial networks (GANs) (see Sec. 2 for an overview) has demonstrated impressive results for 3D shape generation using 2D generators. We build upon this idea of learning to generate triplane representations [6] that encode 3D scenes or

radiance fields as a set of axis-aligned 2D feature planes. The structure of a triplane is analogous to that of a 2D image and can be used as part of a 3D generative method that leverages conventional 2D generator architectures.

Inspired by recent efforts in designing efficient 3D GAN architectures, we introduce a neural field-based diffusion framework for 3D representation learning. Our approach follows a two-step process. In the first step, a training set of 3D scenes is factored into a set of per-scene triplane features and a single, shared feature decoder. In the second step, a 2D diffusion model is trained on these triplanes. The trained diffusion model can then be used at inference time to generate novel and diverse 3D scenes. By interpreting triplanes as multi-channel 2D images and thus decoupling generation from rendering, we can leverage current (and likely future) SOTA 2D diffusion model backbones nearly out of the box. Fig. 1 illustrates how a single object is generated with our framework (top), and how two generated objects—even with different topologies—can be interpolated (bottom).

Our core contributions are as follows:

^{*}Equal contribution.

Part of the work was done during an internship at Stanford.

Project page: <https://jryanshue.com/nfd>

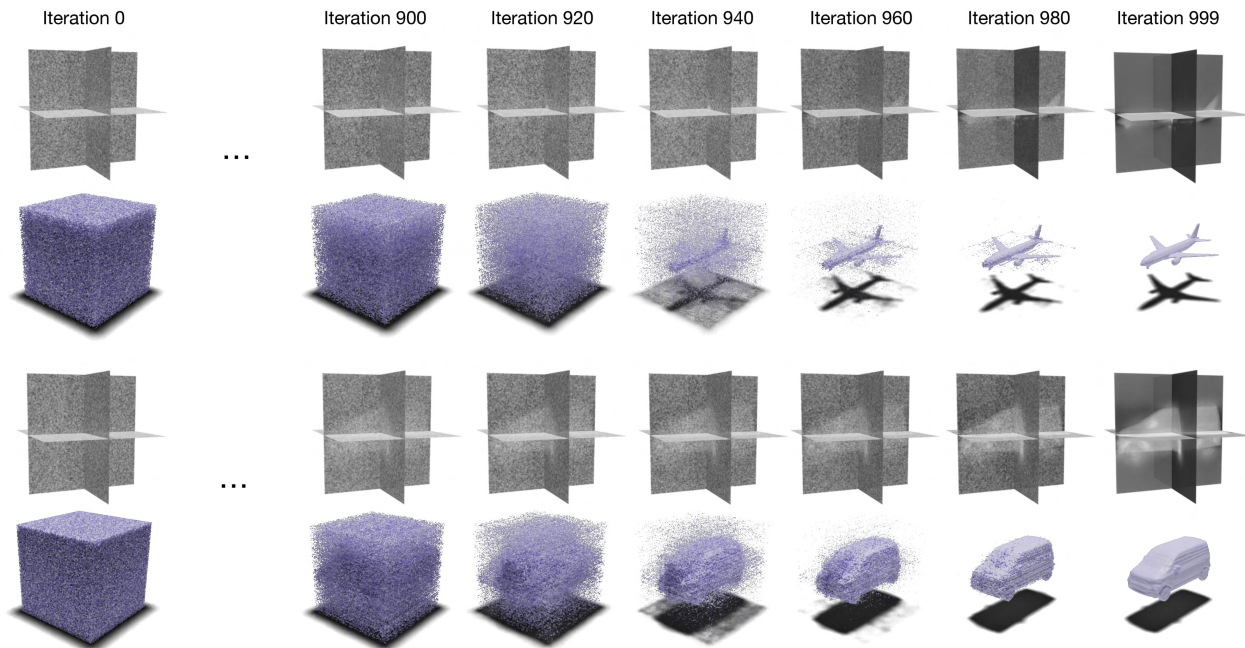


Figure 2. **Visualization of the denoising process.** Here, we show examples of triplanes as they are iteratively denoised at inference, as well as the shapes we obtain by “decoding” the noisy triplanes with our jointly-learned MLP. By interpreting triplane features simply as multi-channel feature images, we build our framework around 2D diffusion models.

- We introduce a generative framework for diffusion on 3D scenes that utilizes 2D diffusion model backbones and has a built-in 3D inductive bias.
- We show that our approach is capable of generating both high-fidelity and diverse 3D scenes that outperform state-of-the-art 3D GANs.

2. Related Work

Neural fields. Implicit neural representations, or neural fields, hold the SOTA for 3D scene representation [74, 80]. They either solely learn geometry [1, 4, 5, 10, 11, 15, 21, 23, 45, 47, 48, 57, 65, 73] or use posed images to jointly optimize geometry and appearance [6, 7, 18, 25, 29, 33, 36–39, 44, 49, 50, 54, 55, 59, 67, 72, 82–84]. Neural fields represent scenes as continuous functions, allowing them to scale well with scene complexity compared to their discrete counterparts [40, 66]. Initial methods used a single, large multilayer perceptron (MLP) to represent entire scenes [10, 47, 49, 57, 65], but reconstruction with this approach can be computationally inefficient because training such a representation requires thousands of forward passes through the large model per scene. Recent years have shown a trend towards locally conditioned representations, which either learn local functions [5, 9, 28, 62] or locally modulate a shared function with a hybrid explicit–implicit representation [4, 6, 12, 19–21, 37, 43, 45, 58]. These methods use small MLPs, which are efficient during inference and significantly

better at capturing local scene details. We adopt the expressive hybrid triplane representation introduced by Chan et al. [6]. Triplanes are efficient, scaling with the surface area rather than volume, and naturally integrate with expressive, fine-tuned 2D generator architectures. We modify the triplane representation for compatibility with our denoising framework.

Generative synthesis in 2D and 3D. Some of the most popular generative models include GANs [22, 31, 32], autoregressive models [16, 60, 77, 78], score matching models [68, 70, 71], and denoising diffusion probabilistic models (DDPMs) [13, 26, 52, 76]. DDPMs are arguably the SOTA approach for synthesizing high-quality and diverse 2D images [13]. Moreover, GANs can be difficult to train and suffer from issues like mode collapse [75] whereas diffusion models train stably and have been shown to better capture the full training distribution.

In 3D, however, GANs still outperform alternative generative approaches [6, 7, 17, 24, 34, 35, 37, 46, 51, 53, 56, 61, 64, 79, 85, 87]. Some of the most successful 3D GANs use an expressive 2D generator backbone (e.g., StyleGAN2 [32]) to synthesize triplane representations which are then decoded with a small, efficient MLP [6]. Because the decoder is small and must generalize across many local latents, these methods assign most of their expressiveness to the powerful backbone. In addition, these methods treat the triplane as a multi-channel image, allowing the generator backbone to be used almost out of the box.

Current 3D diffusion models [2, 14, 42, 81, 86] are still very limited. They either denoise a single latent or do not utilize neural fields at all, opting for a discrete point-cloud-based approach. For example, concurrently developed single-latent approaches [2, 14] generate a global latent for conditioning the neural field, relying on a 3D decoder to transform the scene representation from 1D to 3D without directly performing 3D diffusion. As a result, the diffusion model does not actually operate in 3D, losing this important inductive bias and generating blurry results. Point-cloud-based approaches [42, 86], on the other hand, give the diffusion model explicit 3D control over the shape, but limit its resolution and scalability due to the coarse discrete representation. While showing promise, both 1D-to-3D and point cloud diffusion approaches require specific architectures that cannot easily leverage recent advances in 2D diffusion models.

In our work, we propose to directly generate triplanes with out-of-the-box SOTA 2D diffusion models, granting the diffusion model near-complete control over the generated neural field. Key to our approach is our treatment of well-fit triplanes in a shared latent space as ground truth data for training our diffusion model. We show that the latent space of these triplanes is grounded spatially in local detail, giving the diffusion model a critical inductive bias for 3D generation. Our approach gives rise to an expressive 3D diffusion model.

3. Triplane Diffusion Framework

Here, we explain the architecture of our neural field diffusion (NFD) model for 3D shapes. In Section 3.1, we explain how we can represent the occupancy field of a single object using a triplane. In Section 3.2, we describe how we can extend this framework to represent an entire dataset of 3D objects. In Section 3.3, we describe the regularization techniques that we found necessary to achieve optimal results. Finally, Sections 3.4 and 3.5 illustrate training and sampling from our model. For an overview of the pipeline at inference, see Figure 3.

3.1. Representing a 3D Scene using a Triplane

Neural fields have been introduced as continuous and expressive 3D scene representations. In this context, a neural field $\text{NF} : \mathbb{R}^3 \rightarrow \mathbb{R}^M$ is a neural network-parameterized mapping function that takes as input a three-dimensional coordinate \mathbf{x} and that outputs an M -dimensional vector representing the neural field. Neural fields have been demonstrated for occupancy fields [47], signed distance functions [57], radiance fields [49], among many other types of signals [65]. For the remainder of this work, we focus on 3D scene representations using occupancy fields such that the output of the neural field is a binary value, indicating whether a coordinate is inside or outside an object and $M = 1$.

The triplane representation is a hybrid explicit-implicit network architecture for neural fields that is particularly effi-

cient to evaluate [6]. This representation uses three 2D feature planes $\mathbf{f}_{xy}, \mathbf{f}_{xz}, \mathbf{f}_{yz} \in \mathbb{R}^{N \times N \times C}$ with a spatial resolution of $N \times N$ and C feature channels each, and a multilayer perceptron (MLP) “decoder” tasked with interpreting features sampled from the planes. A 3D coordinate is queried by projecting it onto each of the axis-aligned planes (i.e., the $x-y$, $x-z$, and $y-z$ planes), querying and aggregating the respective features, and decoding the resulting feature using a lightweight MLP ϕ with parameters ϕ . Similar to Chan et al. [6], we found the sum to be an efficient feature aggregation function, resulting in the following formulation for the triplane architecture:

$$\text{NF}(\mathbf{x}) = \text{MLP}_\phi(\mathbf{f}_{xy}(\mathbf{x}) + \mathbf{f}_{yz}(\mathbf{x}) + \mathbf{f}_{xz}(\mathbf{x})). \quad (1)$$

The feature planes and MLP can be jointly optimized to represent the occupancy field of a shape.

3.2. Representing a Class of Objects with Triplanes

We aim to convert our dataset of shapes into a dataset of triplanes so that we can train a diffusion model on these learned feature planes. However, because the MLP and feature planes are typically jointly learned, we cannot simply train a triplane for each object of the dataset individually. If we did, the MLP’s corresponding to each object in our dataset would fail to generalize to triplanes generated by our diffusion model. Therefore, instead of training triplanes for each object in isolation, we jointly optimize the feature planes for many objects simultaneously, along with a decoder that is *shared* across all objects. This joint optimization results in a dataset of optimized feature planes and an MLP capable of interpreting any triplane from the dataset distribution. Thus, at inference, we can use this MLP to decode feature planes generated by our model.

In practice, during training, we are given a dataset of I objects, and we preprocess the coordinates and ground-truth occupancy values of J points per object. Typically, $J = 10M$, where $5M$ points are sampled uniformly throughout the volume and $5M$ points are sampled near the object surface. Our naive training objective is a simple $L2$ loss between predicted occupancy values $\text{NF}^{(i)}(\mathbf{x}_j^{(i)})$ and ground-truth occupancy values $O_j^{(i)}$ for each point, where $\mathbf{x}_j^{(i)}$ denotes the j^{th} point from the i^{th} scene:

$$\mathcal{L}_{\text{NAIVE}} = \sum_i \sum_j \left\| \text{NF}^{(i)}(\mathbf{x}_j^{(i)}) - O_j^{(i)} \right\|_2 \quad (2)$$

During training, we optimize Equation 2 for a shared MLP parameterized by ϕ , as well as the feature planes corresponding to every object in our dataset:

$$\left\{ \phi, \mathbf{f}_{xy}^{(i)}, \mathbf{f}_{xz}^{(i)}, \mathbf{f}_{yz}^{(i)} \right\} = \underset{\left\{ \phi, \mathbf{f}_{xy}^{(i)}, \mathbf{f}_{xz}^{(i)}, \mathbf{f}_{yz}^{(i)} \right\}}{\text{argmin}} \mathcal{L}_{\text{NAIVE}} \quad (3)$$




Figure 3. **Pipeline.** Sampling a 3D neural field from our model consists of two decoupled processes: 1) using a trained DDPM to iteratively denoise latent noise into feature maps and 2) using a locally conditioned Occupancy Network to decode the resulting triplane into the final neural field. This architecture allows the DDPM to generate samples with a 3D inductive bias while utilizing existing 2D DDPM backbones and a continuous output representation.

3.3. Regularizing Triplanes for Effective Generalization

Following the procedure outlined in the previous section, we can learn a dataset of triplane features and a shared triplane decoder; we can then train a diffusion model on these triplane features and sample novel shapes at inference. Unfortunately, the result of this naive training procedure is a generative model for triplanes that produces shapes with significant artifacts.

We find it necessary to regularize the triplane features during optimization to simplify the data manifold that the diffusion model must learn. Therefore, we include total variation (TV) regularization terms with weight λ_1 in the loss function to ensure that the feature planes of each training scene do not contain spurious high-frequency information. This strategy makes the distribution of triplane features more similar to the manifold of natural images (see supplement), which we found necessary to robustly train a diffusion model on them (see Sec. 4).

While the trained feature values are unbounded, our DDPM backbone requires training inputs with values in the range $[-1,1]$. We address this by normalizing the feature planes before training, but this process is sensitive to outliers. As a result, we include an L2 regularization term on the triplane features with weight λ_2 to discourage outlying values.

We also include an explicit density regularization (EDR) term. Due to our ground-truth occupancy data being concentrated on the surface of the shapes, there is often insufficient data to learn a smooth outside-of-shape volume. Our EDR term combats this issue by sampling a set of random points from the volume, offsetting the points by a random vector ω , feeding both sets through the MLP, and calculating the mean squared error. Notationally, this term can be represented as $\text{EDR}(\text{NF}(\mathbf{x}), \omega) = \|\text{NF}(\mathbf{x}) - \text{NF}(\mathbf{x} + \omega)\|_2^2$. We find this term necessary to remove floating artifacts in the volume (see Sec. 4)

Our training objective, with added regularization terms, is as follows:

$$\begin{aligned} \mathcal{L} = & \sum_i^N \sum_j^M \text{BCE} \left(\text{NF}^{(i)}(\mathbf{x}_j^{(i)}) - \text{O}_j^{(i)} \right) \\ & + \lambda_1 \left(\text{TV}(\mathbf{f}_{xy}^{(i)}) + \text{TV}(\mathbf{f}_{xz}^{(i)}) + \text{TV}(\mathbf{f}_{yz}^{(i)}) \right) \\ & + \lambda_2 \left(\|\mathbf{f}_{xy}^{(i)}\|_2 + \|\mathbf{f}_{yz}^{(i)}\|_2 + \|\mathbf{f}_{xz}^{(i)}\|_2 \right) \\ & + \text{EDR} \left(\text{NF}(\mathbf{x}_j^{(i)}), \omega \right) \end{aligned} \quad (4)$$

3.4. Training a Diffusion Model for Triplane Features

For unconditional generation, a diffusion model takes Gaussian noise as input and gradually denoises it in T steps. In our framework, the diffusion model operates on triplane features $\mathbf{f}_0 \dots \mathbf{f}_T \in \mathbb{R}^{N \times N \times 3C}$ that stack the feature channels of all three triplane axes into a single image. In this notation, $\mathbf{f}_T \sim \mathcal{N}(\mathbf{f}_T; 0, \mathbf{I})$ is the triplane feature image consisting of purely Gaussian noise, and $\mathbf{f}_0 \sim q(\mathbf{f}_0)$ is a random sample drawn from the data distribution. The data distribution in our framework includes the pre-factored triplanes of the training set, normalized by the mean and variance of the entire dataset such that each channel has a zero mean and a standard deviation of 0.5.

The *forward* or *diffusion processes* is a Markov chain that gradually adds Gaussian noise to the triplane features, according to a variance schedule $\beta_1, \beta_2, \dots, \beta_T$

$$q(\mathbf{f}_t | \mathbf{f}_{t-1}) = \mathcal{N} \left(\mathbf{f}_t; \sqrt{1 - \beta_t} \mathbf{f}_{t-1}, \beta_t \mathbf{I} \right). \quad (5)$$

This forward process can be directly sampled at step t using the closed-form solution $q(\mathbf{f}_t | \mathbf{f}_0) = \mathcal{N}(\mathbf{f}_t; \sqrt{\bar{\alpha}_t} \mathbf{f}_0, (1 - \bar{\alpha}_t) \mathbf{I})$, where $\bar{\alpha}_t = \prod_{s=1}^t \alpha_s$ with $\alpha_t = 1 - \beta_t$.

The goal of training a diffusion model is to learn the *reverse process*. For this purpose, a function approximator

ϵ_θ is needed that predicts the noise $\epsilon \sim \mathcal{N}(\mathbf{0}, \mathbf{I})$ from its noisy input. Typically, this function approximator is implemented as a variant of a convolutional neural network defined by its parameters θ . Following [26], we train our triplane diffusion model by optimizing the simplified variant of the variational bound on negative log-likelihood:

$$\mathcal{L}_{\text{DDPM}} = \mathbb{E}_{t, \mathbf{f}_0, \epsilon} \left[\left\| \epsilon - \epsilon_\theta \left(\sqrt{\bar{\alpha}_t} \mathbf{f}_0 + \sqrt{1 - \bar{\alpha}_t} \epsilon, t \right) \right\|^2 \right], \quad (6)$$

where t is sampled uniformly between 1 and T .

3.5. Sampling Novel 3D Shapes

The unconditional generation of shapes at inference is a two-stage process that involves sampling a triplane from the trained diffusion model and then querying the neural field.

Sampling a triplane from the diffusion model is identical to sampling an image from a diffusion model. Beginning with a random Gaussian noise $\mathbf{f}_T \sim \mathcal{N}(\mathbf{0}, \mathbf{I})$, we iteratively denoise the sample in T steps as

$$\mathbf{f}_{t-1} = \frac{1}{\sqrt{\alpha_t}} \left(\mathbf{f}_t - \frac{1 - \alpha_t}{\sqrt{1 - \bar{\alpha}_t}} \epsilon_\theta(\mathbf{f}_t, t) \right) + \sigma_t \epsilon, \quad (7)$$

where $\epsilon \sim \mathcal{N}(\mathbf{0}, \mathbf{I})$ for all but the very last step (*i.e.*, $t = 1$), at which $\epsilon = 0$ and $\sigma_t^2 = \beta_t$.

The result of the denoising process, \mathbf{f}_0 , is a sample from the normalized triplane feature image distribution. Denormalizing it using the dataset normalization statistics and splitting the generated features into the axis aligned planes $\mathbf{f}_{xy}, \mathbf{f}_{yz}, \mathbf{f}_{xz}$ yields a set of triplane features which, when combined with the pre-trained MLP, are used to query the neural field.

We use the marching cubes algorithm [41] to extract meshes from the resulting neural fields. Note that our framework is largely agnostic to the diffusion backbone used; we choose to use ADM [52], a 2D state-of-the-art diffusion model.

Source code and pre-trained models will be made available.

4. Experiments

Datasets. To compare NFD against existing 3D generative methods, we train our model on three object categories from the ShapeNet dataset individually. Consistent with previous work [85, 86], we choose the categories: *cars*, *chairs* and *airplanes*. Each mesh is normalized to lie within $[-1, 1]^3$ and then passed through watertighting. The generation of ground truth triplanes then works as follows: we precompute the occupancies of 10M points per object, where 5M points are distributed uniformly at random in the volume, and 5M points are sampled within a 0.01 distance from the mesh surface.

¹As PVD outputs 098 point clouds, we apply the ball pivoting algorithm (BPA) to PVD outputs before calculating FID. BPA was selected as it achieved a good balance between speed and quality.

Data	Method	FID ↓	Precision ↑	Recall ↑
Cars	PVD ¹	335.8	0.1	0.2
	SDF-StyleGAN	98.0	35.9	36.2
	NFD (Ours)	83.6	49.5	50.5
Chairs	PVD	305.8	0.2	1.7
	SDF-StyleGAN	36.5	90.9	87.4
	NFD (Ours)	26.4	92.4	94.8
Planes	PVD	244.4	2.7	3.8
	SDF-StyleGAN	65.8	64.5	72.8
	NFD (Ours)	32.4	70.5	81.1

Table 1. Render quality metrics on ShapeNet. We achieve state-of-the-art FID, which measures overall quality, as well as well as state-of-the-art precision and recall, which measure fidelity and diversity independently. Metrics calculated on shaded renderings of generated and ground-truth shapes.

We then train an MLP jointly with as many triplanes as we can fit in the GPU memory of a single A6000 GPU. In our case, we initially train on the first 500 objects in the dataset. After this initial joint optimization, we freeze the shared MLP and use it to optimize the triplanes of the remaining objects in the dataset. All triplanes beyond the first 500 are optimized individually with the same shared MLP; thus, the training of these triplanes can be effectively parallelized.

Evaluation metrics. As in [85], we choose to evaluate our model using an adapted version of Fréchet inception distance (FID) that utilizes rendered shading images of our generated meshes. Shading-image FID [85] overcomes limitations of other mesh-based evaluation metrics such as the light-field-descriptor (LFD) [8] by taking human perception into consideration. Zheng et al. [85] provide a detailed discussion of the various evaluation metrics for 3D generative models. Following the method [85], shading images of each shape are rendered from 20 distinct views; FID is then compared across each view and averaged to obtain a final score:

$$\text{FID} = \frac{1}{20} \left[\sum_{i=1}^{20} \|\mu_g^i - \mu_r^i\|^2 + \text{Tr} \left(\Sigma_g^i + \Sigma_r^i - 2(\Sigma_r^i \Sigma_g^i)^{\frac{1}{2}} \right) \right], \quad (8)$$

where g and r represent the generated and training datasets, while μ^i, Σ^i represent the mean and covariance matrices for shading images rendered from the i^{th} view, respectively.

Along with FID, we also report precision and recall scores using the method proposed by Sajjadi et al. [63]. While FID correlates well with perceived image quality, the one-dimensional nature of the metric prevents it from identifying different failure modes. Sajjadi et al. [63] aim to disentangle FID into separate metrics known as precision and recall, where the former correlates to the quality of the generated images and the latter represents the diversity of




Figure 4. We compare 3D shapes generated by our model against generations of state-of-the-art baselines for ShapeNet *Cars*, *Chairs*, and *Planes*. Our model synthesizes shapes with noticeably sharper details than the previous state-of-the-art, while also capturing the broad diversity in each category.

the generative model.

Baselines. We compare our method against state-of-the-art point-based and neural-field-based 3D generative models, namely PVD [86] and SDF-StyleGAN [85]. For evaluation, we use the pre-trained models for both methods on the three ShapeNet categories listed above. Note that PVD is inherently a point-based generative method and therefore does not output a triangle mesh needed for shading image rendering. To circumvent this, we choose to convert generated point clouds to triangle meshes using the ball-pivoting algorithm [3].

Results. We provide qualitative results, comparing samples generated by our method to samples generated by baselines, in Figure 4. Our method generates a diverse and finely detailed collection of objects. Objects produced by our method contain sharp edges and features that we would expect to be difficult to accurately reconstruct—note that delicate features, such as the suspension of cars, the slats in chairs, and armaments of planes, are faithfully generated. Perhaps more importantly, samples generated by our model are diverse—our model successfully synthesizes many different types of cars, chairs,

and planes, including reproductions of several varieties that we would expect to be rare in the training dataset.

In comparison, while *PVD* also produces a wide variety of shapes, it is limited by its nature to generating only coarse object shapes. Furthermore, because *PVD* produces a fixed-size point cloud with only 2048 points, it cannot synthesize fine elements.

SDF-StyleGAN creates high-fidelity shapes, accurately reproducing many details, such as airplane engines and chair legs. However, our method is more capable of capturing very fine features. Note that while *SDF-StyleGAN* smooths over the division between tire and wheel well when generating cars, our method faithfully portrays this gap. Similarly, our method synthesizes the tails and engines of airplanes, and the legs and planks of chairs, with noticeably better definition. Our method also apparently generates a greater diversity of objects than *SDF-StyleGAN*. While *SDF-StyleGAN* capably generates varieties of each ShapeNet class, our method reproduces the same classes with greater variation. This is expected, as a noted advantage of diffusion models over GANs is better mode coverage.

We provide quantitative results in Table 1. The metrics




Figure 5. **Interpolation.** Our model learns a continuous latent space of triplanes. We can smoothly interpolate between two noise triplanes, resulting in semantically meaningful shape interpolation.

tell a similar story to the qualitative results. Quantitatively, NFD outperforms all baselines in FID, precision, and recall for each ShapeNet category. FID is a standard one-number metric for evaluating generative models, and our performance under this evaluation indicates the generally better quality of object renderings. Precision evaluates the renderings’ fidelity, and recall evaluates their diversity. Outperforming baselines in both precision and recall suggest that our model produces higher fidelity of shapes and a more diverse distribution of shapes. This is consistent with the qualitative results in Figure 4, where our method produced sharper and more complex objects while also covering more modes.

Semantically meaningful interpolation. Figure 5 shows latent space interpolation between pairs of generated neural fields. As shown in prior work [69], smooth interpolation in the latent space of diffusion models can be achieved by interpolation between noise tensors before they are iteratively denoised by the model. As in their method, we sample from our trained model using a deterministic DDIM, and we use spherical interpolation so that the intermediate latent noise retains the same distribution. Our method is capable of smooth latent space interpolation in the generated triplanes and their corresponding neural fields.

4.1. Ablation Studies

We validate the design of our framework by ablating components of our regularization strategies using the cars dataset.

Explicit density regularization. As discussed by Park et al. [57], the precision of the ground truth decoded meshes is limited by the finite number of point samples guiding the training of the decision boundaries. Because we rely on a limited number of pre-computed coordinate–occupancy pairs to train our triplanes, it is easy to overfit to this limited training set. Even when optimizing a single triplane in isolation (i.e., *without* learning a generative model), this overfitting manifests in “floater” artifacts in the optimized




Figure 6. **Ablation over density regularization.** Clear artifacts are visible in the resulting occupancy field without explicit density regularization. In this example, we optimize a single triplane on a single shape.

neural field. Figure 6 shows an example where we fit a single triplane with and without density regularization. Without density regularization, the learned occupancy field contains significant artifacts; with density regularization, the learned occupancy field captures a clean object.

Triplane regularization. Regularization of the triplanes is essential for training a well-behaved diffusion model. Figure 7 compares generated samples produced by our entire framework, with and without regularization terms. If we train only with Equation 2, i.e., without regularization terms, we can optimize a dataset of triplane features and train a diffusion model to generate samples. However, while the surfaces of the optimized shapes will appear real, the triplane features themselves will have many high-frequency artifacts, and these convoluted feature images are a difficult manifold for even a powerful diffusion model to learn. Consequently, generated triplane features produced by a trained diffusion model decode into shapes with significant artifacts. We note that these artifacts are present *only* in generated samples; shapes directly factored from the ground-truth shapes are artifact-free, even without regularization.

Training with Equation 4 introduces TV, L2, and




Figure 7. **Ablation over regularized triplanes.** A generative model trained on unregularized triplanes produces samples with significant artifacts. Effective regularization of triplane features enables training of a generative model that produces shapes without artifacts. Top left: triplane features learned only with Equation 2 contain many high frequency artifacts. Bottom left: a diffusion model trained on these unregularized triplanes fails to produce convincing samples. Top right: triplane features learned with Equation 4 are noticeably smoother. Bottom right: A diffusion model trained on these regularized triplanes produces high-quality shapes.

Method	FID ↓	Precision ↑	Recall ↑
No regularization	285.8	1.6	0.6
Density + TV + L2 Reg.	83.6	49.5	50.5

Table 2. Quantitative results for the ablation on triplane regularization. Our model performs poorly without explicit regularization on the triplanes.

density regularizing factors. Triplanes learned with these regularization terms are noticeably smoother, with frequency distributions that more closely align with those found in natural images (see supplement). As we would expect, a diffusion model more readily learns the manifold of regularized triplane features. Samples produced by a diffusion model trained on these regularized shapes decode into convincing and artifact-free shapes.

5. Discussion

In summary, we introduce a 3D-aware diffusion model that uses a 2D diffusion backbone to generate triplane feature maps, which are assembled into 3D neural fields. Our approach improves the quality and diversity of generated objects over existing 3D-aware generative models by a large margin.

Limitations. Similarly to other generative methods, training a diffusion model is slow and computationally demanding. Diffusion models, including ours, are also slow to evaluate, whereas GANs, for example, can be evaluated in real-time once trained. Luckily, our method will benefit from improvements to 2D diffusion models in this research area.




Figure 8. Failure cases. We observe that our model at times generate axis-aligned artifacts and struggles to account for thin structures, likely caused by the use of a triplane representation.

Slow sampling at inference could be addressed by more efficient samplers [30] and potentially enable real-time synthesis. While a step forward in quality, some of the samples generated by our method suffer from artifacts, as depicted by Fig. 8. Strategies like guidance [13, 27], which trade off diversity for fidelity, may reduce the prevalence of these outliers.

Future Work. We have demonstrated an effective way to generate occupancy fields, but in principle, our approach can be extended to generating any type of neural field that can be represented by a triplane. In particular, triplanes have already been shown to be excellent representations for radiance fields, so it seems natural to extend our diffusion approach to generating NeRFs. While we demonstrate successful results for unconditional generation, conditioning our generative model on text, images, or other input would be an exciting avenue for future work.

Ethical Considerations. Generative models, including ours, could be extended to generate DeepFakes. These pose a societal threat, and we do not condone using our work to generate fake images or videos of any person intending to spread misinformation or tarnish their reputation.

Conclusion. 3D-aware object synthesis has many exciting applications in vision and graphics. With our work, which is among the first to connect powerful 2D diffusion models and 3D object synthesis, we take a significant step towards utilizing emerging diffusion models for this goal.

Acknowledgements

We thank Vincent Sitzmann for valuable discussions. This project was in part supported by Samsung, the Stanford Institute for Human-Centered AI (HAI), the Stanford Center for Integrated Facility Engineering (CIFE), NSF RI #2211258, Autodesk, and a PECASE from the ARO.

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] Miguel Ángel Bautista, Pengsheng Guo, Samira Abnar, Walter Talbott, Alexander Toshev, Zhuoyuan Chen, Laurent Dinh, Shuangfei Zhai, Hanlin Goh, Daniel Ulbricht, Afshin Dehghan, and Josh M. Susskind. GAUDI: A neural architect for immersive 3d scene generation. *CoRR*, abs/2207.13751, 2022. 1, 3
- [3] Fausto Bernardini, Joshua Mittleman, Holly E. Rushmeier, Cláudio T. Silva, and Gabriel Taubin. The ball-pivoting algorithm for surface reconstruction. *IEEE Transactions on Visualization and Computer Graphics*, 5:349–359, 1999. 6
- [4] Alexandre Boulch and Renaud Marlet. POCO: point convolution for surface reconstruction. *CoRR*, abs/2201.01831, 2022. 2
- [5] 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
- [6] Eric R. Chan, Connor Z. Lin, Matthew A. Chan, Koki Nagano, Boxiao Pan, Shalini De Mello, Orazio Gallo, Leonidas J. Guibas, Jonathan Tremblay, Sameh Khamis, Tero Karras, and Gordon Wetzstein. Efficient geometry-aware 3d generative adversarial networks. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 16123–16133, June 2022. 1, 2, 3
- [7] 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. 2
- [8] Ding-Yun Chen, Xiao-Pei Tian, Edward Yu-Te Shen, and Ming Ouhyoung. On visual similarity based 3d model retrieval. *Computer Graphics Forum*, 22, 2003. 5
- [9] 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
- [10] Zhiqin Chen and Hao Zhang. Learning implicit fields for generative shape modeling. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2019. 2
- [11] Thomas Davies, Derek Nowrouzezahrai, and Alec Jacobson. Overfit neural networks as a compact shape representation. *arXiv preprint arXiv:2009.09808*, 2020. 2
- [12] 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
- [13] Prafulla Dhariwal and Alexander Nichol. Diffusion models beat gans on image synthesis. In M. Ranzato, A. Beygelzimer, Y. Dauphin, P.S. Liang, and J. Wortman Vaughan, editors, *Advances in Neural Information Processing Systems*, volume 34, pages 8780–8794. Curran Associates, Inc., 2021. 2, 8
- [14] Emilien Dupont, Hyunjik Kim, S. M. Ali Eslami, Danilo J. Rezende, and Dan Rosenbaum. From data to functa: Your data point is a function and you should treat it like one. *CoRR*, abs/2201.12204, 2022. 1, 3
- [15] SM Ali Eslami, Danilo Jimenez Rezende, Frederic Besse, Fabio Viola, Ari S Morcos, Marta Garnelo, Avraham Ruderman, Andrei A Rusu, Ivo Danihelka, Karol Gregor, et al. Neural scene representation and rendering. *Science*, 2018. 2
- [16] Patrick Esser, Robin Rombach, and Björn Ommer. Taming transformers for high-resolution image synthesis, 2020. 2
- [17] Matheus Gadelha, Subhansu Maji, and Rui Wang. 3d shape induction from 2d views of multiple objects. In *2017 International Conference on 3D Vision (3DV)*, pages 402–411, 2017. 2
- [18] 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
- [19] Kyle Genova, Forrester Cole, Avneesh Sud, Aaron Sarna, and Thomas Funkhouser. Local deep implicit functions for 3D shape. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2020. 2
- [20] Kyle Genova, Forrester Cole, Daniel Vlasic, Aaron Sarna, William T Freeman, and Thomas Funkhouser. Learning shape templates with structured implicit functions. In *IEEE International Conference on Computer Vision (ICCV)*, 2019. 2
- [21] Simon Giebenhain and Bastian Goldlücke. Air-nets: An attention-based framework for locally conditioned implicit representations. In *3DV*, pages 1054–1064. IEEE, 2021. 2
- [22] 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
- [23] 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
- [24] 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. 2
- [25] 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
- [26] Jonathan Ho, Ajay Jain, and Pieter Abbeel. Denoising diffusion probabilistic models. In *Advances in Neural Information Processing Systems*, volume 33, pages 6840–6851, 2020. 2, 5
- [27] Jonathan Ho and Tim Salimans. Classifier-free diffusion guidance. *arXiv preprint arXiv:2207.12598*, 2022. 8
- [28] 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
- [29] 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
- [30] Tero Karras, Miika Aittala, Timo Aila, and Samuli Laine. Elucidating the design space of diffusion-based generative models. *ArXiv*, abs/2206.00364, 2022. 8

- [31] Tero Karras, Samuli Laine, and Timo Aila. A style-based generator architecture for generative adversarial networks. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2019. 2
- [32] 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. 2
- [33] 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
- [34] Adam R Kosiorek, Heiko Strathmann, Daniel Zoran, Pol Moreno, Rosalia Schneider, Sona Mokra, and Danilo Jimenez Rezende. Nerf-vae: A geometry aware 3d scene generative model. In Marina Meila and Tong Zhang, editors, *Proceedings of the 38th International Conference on Machine Learning*, volume 139 of *Proceedings of Machine Learning Research*, pages 5742–5752. PMLR, 18–24 Jul 2021. 2
- [35] Yiyi Liao, Katja Schwarz, Lars Mescheder, and Andreas Geiger. Towards unsupervised learning of generative models for 3d controllable image synthesis. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2020. 2
- [36] 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
- [37] 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
- [38] Shichen Liu, Shunsuke Saito, Weikai Chen, and Hao Li. Learning to infer implicit surfaces without 3D supervision. *arXiv preprint arXiv:1911.00767*, 2019. 2
- [39] 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
- [40] 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
- [41] William E Lorensen and Harvey E Cline. Marching cubes: A high resolution 3D surface construction algorithm. *ACM Transactions on Graphics (ToG)*, 1987. 5
- [42] Shitong Luo and Wei Hu. Diffusion probabilistic models for 3d point cloud generation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 2837–2845, June 2021. 1, 3
- [43] 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
- [44] 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
- [45] Ishit Mehta, Michaël Gharbi, Connelly Barnes, Eli Shechtman, Ravi Ramamoorthi, and Manmohan Chandraker. Modulated periodic activations for generalizable local functional representations. In *ICCV*, pages 14194–14203. IEEE, 2021. 2
- [46] Quan Meng, Anpei Chen, Haimin Luo, Minye Wu, Hao Su, Lan Xu, Xuming He, and Jingyi Yu. Gnerf: Gan-based neural radiance field without posed camera. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 6351–6361, October 2021. 2
- [47] 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, 3
- [48] 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
- [49] 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. 2, 3
- [50] 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
- [51] Thu Nguyen-Phuoc, Chuan Li, Lucas Theis, Christian Richardt, and Yong-Liang Yang. Hologan: Unsupervised learning of 3d representations from natural images. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, October 2019. 2
- [52] Alexander Quinn Nichol and Prafulla Dhariwal. Improved denoising diffusion probabilistic models. In Marina Meila and Tong Zhang, editors, *Proceedings of the 38th International Conference on Machine Learning*, volume 139 of *Proceedings of Machine Learning Research*, pages 8162–8171. PMLR, 18–24 Jul 2021. 2, 5
- [53] Michael Niemeyer and Andreas Geiger. Giraffe: Representing scenes as compositional generative neural feature fields. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 11453–11464, June 2021. 2
- [54] 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
- [55] 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
- [56] 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. [2](#)
- [57] 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](#), [3](#), [7](#)
- [58] Songyou Peng, Michael Niemeyer, Lars Mescheder, Marc Pollefeys, and Andreas Geiger. Convolutional occupancy networks. In *European Conference on Computer Vision (ECCV)*, 2020. [2](#)
- [59] 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](#)
- [60] Ali Razavi, Aaron van den Oord, and Oriol Vinyals. Generating diverse high-fidelity images with vq-vae-2. In H. Wallach, H. Larochelle, A. Beygelzimer, F. d'Alché-Buc, E. Fox, and R. Garnett, editors, *Advances in Neural Information Processing Systems*, volume 32. Curran Associates, Inc., 2019. [2](#)
- [61] Daniel Rebain, Mark Matthews, Kwang Moo Yi, Dmitry Lagun, and Andrea Tagliasacchi. Lolnerf: Learn from one look. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 1558–1567, June 2022. [2](#)
- [62] 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](#)
- [63] Mehdi S. M. Sajjadi, Olivier Bachem, Mario Lucic, Olivier Bousquet, and Sylvain Gelly. Assessing generative models via precision and recall. In *NeurIPS*, 2018. [5](#)
- [64] 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. [2](#)
- [65] 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](#), [3](#)
- [66] 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](#)
- [67] 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](#)
- [68] Jascha Sohl-Dickstein, Eric Weiss, Niru Maheswaranathan, and Surya Ganguli. Deep unsupervised learning using nonequilibrium thermodynamics. In Francis Bach and David Blei, editors, *Proceedings of the 32nd International Conference on Machine Learning*, volume 37 of *Proceedings of Machine Learning Research*, pages 2256–2265, Lille, France, 07–09 Jul 2015. PMLR. [2](#)
- [69] Jiaming Song, Chenlin Meng, and Stefano Ermon. Denoising diffusion implicit models. In *International Conference on Learning Representations*, 2021. [7](#)
- [70] Yang Song and Stefano Ermon. Generative modeling by estimating gradients of the data distribution. In H. Wallach, H. Larochelle, A. Beygelzimer, F. d'Alché-Buc, E. Fox, and R. Garnett, editors, *Advances in Neural Information Processing Systems*, volume 32. Curran Associates, Inc., 2019. [2](#)
- [71] Yang Song and Stefano Ermon. Improved techniques for training score-based generative models. In H. Larochelle, M. Ranzato, R. Hadsell, M.F. Balcan, and H. Lin, editors, *Advances in Neural Information Processing Systems*, volume 33, pages 12438–12448. Curran Associates, Inc., 2020. [2](#)
- [72] 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](#)
- [73] 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](#)
- [74] Ayush Tewari, Justus Thies, Ben Mildenhall, Pratul Srinivasan, Edgar Tretschk, W Yifan, Christoph Lassner, Vincent Sitzmann, Ricardo Martin-Brualla, Stephen Lombardi, et al. Advances in neural rendering. In *Computer Graphics Forum*, volume 41, pages 703–735. Wiley Online Library, 2022. [2](#)
- [75] Hoang Thanh-Tung and Truyen Tran. Catastrophic forgetting and mode collapse in gans. In *IJCNN*, pages 1–10. IEEE, 2020. [2](#)
- [76] Arash Vahdat, Karsten Kreis, and Jan Kautz. Score-based generative modeling in latent space. In M. Ranzato, A. Beygelzimer, Y. Dauphin, P.S. Liang, and J. Wortman Vaughan, editors, *Advances in Neural Information Processing Systems*, volume 34, pages 11287–11302. Curran Associates, Inc., 2021. [2](#)
- [77] Aäron van den Oord, Nal Kalchbrenner, and Koray Kavukcuoglu. Pixel recurrent neural networks. In Maria Florina Balcan and Kilian Q. Weinberger, editors, *Proceedings of The 33rd International Conference on Machine Learning*, volume 48 of *Proceedings of Machine Learning Research*, pages 1747–1756, New York, New York, USA, 20–22 Jun 2016. PMLR. [2](#)
- [78] Aaron van den Oord, Oriol Vinyals, and koray kavukcuoglu. Neural discrete representation learning. In I. Guyon, U. Von Luxburg, S. Bengio, H. Wallach, R. Fergus, S. Vishwanathan, and R. Garnett, editors, *Advances in Neural Information Processing Systems*, volume 30. Curran Associates, Inc., 2017. [2](#)
- [79] Jiajun Wu, Chengkai Zhang, Tianfan Xue, Bill Freeman, and Josh Tenenbaum. Learning a probabilistic latent space of object shapes via 3d generative-adversarial modeling. In D. Lee, M. Sugiyama, U. Luxburg, I. Guyon, and R. Garnett, editors, *Advances in Neural Information Processing Systems*, volume 29. Curran Associates, Inc., 2016. [2](#)
- [80] Yiheng Xie, Towaki Takikawa, Shunsuke Saito, Or Litany, Shiqin Yan, Numair Khan, Federico Tombari, James Tompkin,

- Vincent Sitzmann, and Srinath Sridhar. Neural fields in visual computing and beyond. *Comput. Graph. Forum*, 41(2):641–676, 2022. [2](#)
- [81] Guangming Yao, Hongzhi Wu, Yi Yuan, and Kun Zhou. Dd-nerf: Double-diffusion neural radiance field as a generalizable implicit body representation. *arXiv preprint arXiv:2112.12390*, 2021. [3](#)
- [82] 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](#)
- [83] 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](#)
- [84] Kai Zhang, Gernot Riegler, Noah Snavely, and Vladlen Koltun. Nerf++: Analyzing and improving neural radiance fields. *arXiv preprint arXiv:2010.07492*, 2020. [2](#)
- [85] Xin-Yang Zheng, Yang Liu, Peng-Shuai Wang, and Xin Tong. Sdf-stylegan: Implicit sdf-based stylegan for 3d shape generation. *CoRR*, abs/2206.12055, 2022. [2](#), [5](#), [6](#)
- [86] Linqi Zhou, Yilun Du, and Jiajun Wu. 3d shape generation and completion through point-voxel diffusion. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 5826–5835, October 2021. [1](#), [3](#), [5](#), [6](#)
- [87] 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. [2](#)