

AR-NeRF: Unsupervised Learning of Depth and Defocus Effects from Natural Images with Aperture Rendering Neural Radiance Fields

Takuhiro Kaneko

NTT Communication Science Laboratories, NTT Corporation

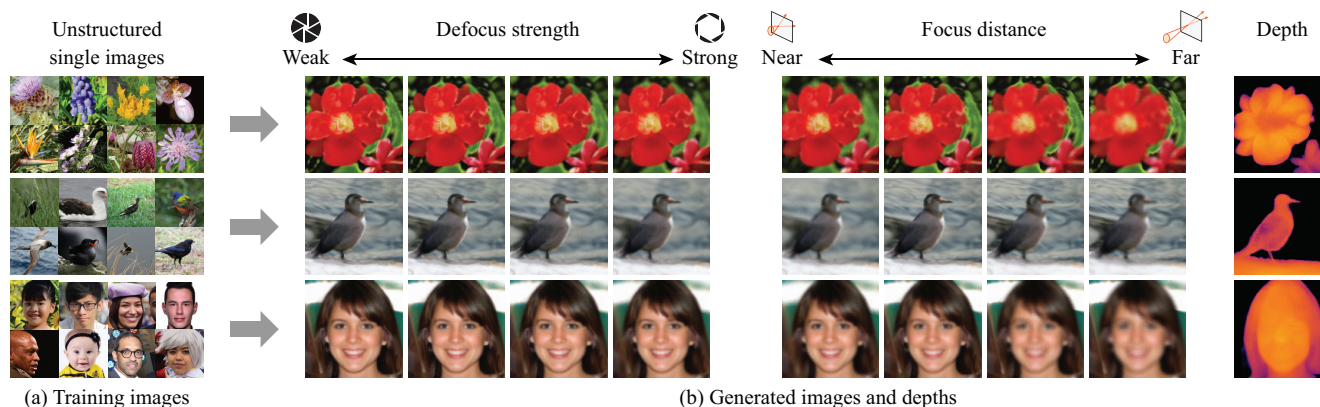


Figure 1. **Unsupervised learning of depth and defocus effects from unstructured (and view-limited) natural images.** (a) During training, we used *only* a collection of unstructured single natural images and did *not* use any supervision (e.g., ground-truth depth, pairs of multiview images, defocus supervision, or pretrained models). (b) After training, *AR-NeRF* can generate sets of images and depths. In particular, in the generation of an image, *AR-NeRF* can adjust the defocus strength and focus distance intuitively and continuously using photometric constraints. The project page is available at <https://www.kecl.ntt.co.jp/people/kaneko.takuhiro/projects/ar-nerf/>.

Abstract

Fully unsupervised 3D representation learning has gained attention owing to its advantages in data collection. A successful approach involves a viewpoint-aware approach that learns an image distribution based on generative models (e.g., generative adversarial networks (GANs)) while generating various view images based on 3D-aware models (e.g., neural radiance fields (NeRFs)). However, they require images with various views for training, and consequently, their application to datasets with few or limited viewpoints remains a challenge. As a complementary approach, an aperture rendering GAN (AR-GAN) that employs a defocus cue was proposed. However, an AR-GAN is a CNN-based model and represents a defocus independently from a viewpoint change despite its high correlation, which is one of the reasons for its performance. As an alternative to an AR-GAN, we propose an aperture rendering NeRF (AR-NeRF), which can utilize viewpoint and defocus cues in a unified manner by representing both factors in a common ray-tracing framework. Moreover, to learn defocus-aware and defocus-independent representations in a disentangled manner, we propose aperture randomized training, for which we learn to generate images while randomizing the aperture size and latent codes independently. During

our experiments, we applied *AR-NeRF* to various natural image datasets, including flower, bird, and face images, the results of which demonstrate the utility of *AR-NeRF* for unsupervised learning of the depth and defocus effects.

1. Introduction

Natural images are 2D projections of the 3D world. Solving the inverse problem, i.e., understanding the 3D world from natural images, is a principal challenge in computer vision and graphics and has been actively studied in various fields owing to its diverse applications, such as environmental understanding in robotics, content creation in advertisements, and photo editing in the arts.

After collecting pairs of 2D and 3D data or sets of multiview images, a successful approach is to learn the 3D predictor using direct or photometric-driven supervision. This approach demonstrates promising results in terms of fidelity. However, the collection of such data is often difficult or impractical. To reduce the collection costs, learning from single images (i.e., from a dataset that includes a single image per training instance) has been actively studied.

To obtain clues under such setting, several studies [19, 32, 75, 77, 93] have introduced object-specific shape models, including 3DMM [5] and SMPL [51], and searched

for solutions within the shape model constraints. Other studies have utilized auxiliary information such as 2D keypoints [33, 87] or 2D silhouettes [10, 22, 27, 44] to simplify the problem by aligning the object parts or separating the target objects from the background. These studies also demonstrate remarkable results; however, the construction of the shape model is not always easy and narrows the applicable objects, and auxiliary information incurs extra costs in terms of data collection.

To alleviate such restrictions, a *fully* unsupervised approach, which learns 3D representations from single images without *any* additional supervision (including auxiliary information and pre-trained models), has gained attention. Under this setting, the viewpoint is a principal clue, which typical methods utilize by learning an image distribution using a generative model (e.g., a generative adversarial network (GAN) [23]) while generating various viewpoint images based on viewpoint-aware 3D models, such as voxels [27, 58, 59], primitives [46], and neural radiance fields (NeRFs) [9, 24, 56, 60, 61, 76]. This allows learning a viewpoint-aware 3D representation; however, owing to the diverse viewpoints needed, the application to a dataset in which viewpoint cues are limited or unavailable without the use of a preprocessing (e.g., natural flower or bird images, as shown in Figure 1) remains a challenge.

As a complement to a viewpoint cue, an aperture rendering GAN (AR-GAN) [34] was proposed to exploit a defocus cue by equipping the aperture rendering [83] on top of the CNN GANs. This constraint allows the learning of both depth and depth-of-field (DoF) effects in an unsupervised manner. However, as a limitation, an AR-GAN employs the defocus cue independently from the viewpoint cue and cannot utilize both factors jointly despite these two factors being highly correlated with the ability to help each other.¹ Consequently, the quality of the depth prediction when using AR-GAN remains limited.

We thus aim to construct a *unified model* that can leverage defocus and viewpoint cues jointly by considering the application of unsupervised learning of the 3D representation (particularly *depth* and *defocus effects*) from natural unstructured (and view-limited) images (Figure 1). To achieve this, we propose a new extension of NeRF called *aperture rendering NeRF (AR-NeRF)*, which can represent defocus effects and viewpoint changes in a unified manner by representing both factors through a common ray-tracing framework. More precisely, in contrast to the standard NeRF, which represents each pixel using a single ray under the *pin-hole camera* assumption, AR-NeRF employs an *aperture camera* [79] that represents each pixel using a collection of rays that converge at the focus plane and whose scale is de-

¹More precisely, in [34], the combinations of an AR-GAN and viewpoint-aware GANs (particularly, HoloGAN [58] and RGBD-GAN [65]) are provided. These models can learn the defocus and viewpoint-aware representations *simultaneously* but *individually*; therefore, such models cannot utilize the learning of one representation for the learning of another.

termined according to the aperture size. Through such modeling, we can represent both viewpoint changes and defocus effects by simply changing the inputs and the integration of the implicit function (multilayer perceptron (MLP)), which converts the point position and view direction into the RGB color and volume density. Consequently, through training, we can optimize the MLP while reflecting both factors.

Moreover, to disentangle defocus-aware and defocus-independent representations in an unsupervised manner, we introduce *aperture randomized training*, in which we learn to generate images in a GAN framework while changing the aperture size and latent codes both randomly and independently. A similar technique is commonly used in viewpoint-aware representation learning [9, 24, 27, 46, 56, 58–61, 76], and this training is useful for disentangling the effect of the corresponding factor from latent codes.

We applied AR-NeRF to natural image datasets, including view-limited (Oxford Flowers [64] (flower) and CUB-200-2011 [90] (bird)) datasets and datasets with various views (FFHQ [39] (face)), and demonstrated that AR-NeRF is better than or comparable to the baseline models, including a state-of-art fully unsupervised depth-learning model (i.e., AR-GAN [34]) and generative NeRF (particularly pi-GAN [9]), in terms of the *depth* prediction accuracy. We also demonstrated that AR-NeRF can manipulate the *defocus effects* (i.e., defocus strength and focus distance) intuitively and continuously while retaining the image quality, whereas AR-GAN has difficulty doing so.

Overall, our contributions can be summarized as follows:

- To achieve an unsupervised learning of the depth and defocus effects, we propose a new extension of NeRF called *AR-NeRF*, which can employ viewpoint and defocus cues in a unified manner by representing both factors in a common ray-tracing framework.
- To disentangle defocus-aware and defocus-independent representations under unsupervised conditions, we introduce *aperture randomized training*, by which we learn to generate images while changing the aperture size and latent codes both randomly and independently.
- We empirically demonstrate the utility of AR-NeRF for the unsupervised learning of the *depth* and *defocus effects* using various natural image datasets, including view-limited (flower and bird) datasets and datasets of various views (face). We provide detailed analyses and extended results in the supplementary material.²

2. Related work

Implicit neural representations. Owing to their 3D-aware, continuous, and memory-efficient nature, implicit neural representations have gained attention in both

²The supplementary material is available at <https://www.kecl.ntt.co.jp/people/kaneko.takuhiro/projects/ar-nerf/>.

learning-based 3D [2, 11, 12, 20, 54, 55, 62, 66, 68, 74] and scene [8, 13, 31, 69] reconstructions. Typical representations are supervised using 3D data; however, to eliminate the need for 3D supervision, the incorporation of differentiable rendering has also been proposed [49, 50, 63, 82, 100]. The most relevant model is NeRF [56], which combines implicit neural representations with volume rendering for a novel view synthesis. Our AR-NeRF is based on NeRF and obtains a *viewpoint-aware* functionality by inheriting it. However, to obtain the *defocus-aware* functionality, AR-NeRF employs an *aperture camera* model instead of a *pinhole camera* model, which is typically used in NeRF. Moreover, the described studies aimed to learn a single network per scene using a *set of multiview* images, whereas we aimed to construct a generative model from the collection of *unstructured single* images. Owing to this difference, we do not aim to compare AR-NeRF with NeRF in this study; however, reimporting our idea (i.e., the usage of an aperture camera) to the original task remains for future research.

Generative adversarial networks. GANs [23] have shown remarkable results in 2D image modeling through a series of advancements (e.g., [7, 37–40]). A strong property of GANs is their ability to learn the data distribution through random sampling without directly defining the distribution. This property allows GANs to learn the data distribution through measurements [6, 35, 36, 43, 67] and architectural constraints [88, 94, 95, 99, 105]. Using the same logic, unsupervised 3D-aware GANs [9, 24, 27, 34, 46, 56, 58–61, 65, 76, 84] have succeeded in learning 3D-aware representations by incorporating 3D-2D projection modules and/or 3D-aware constraints. More specifically, most studies address *viewpoint-aware* representation learning using 3D representations based on voxels [27, 58, 59], primitives [46], and NeRFs [9, 24, 56, 60, 61, 76], and a few studies [34] have addressed the learning of *defocus-aware* representations. Herein, we introduce a unified model that can jointly leverage both *defocus* and *viewpoint* cues to strengthen the latter category model. We demonstrate the utility of their joint usage in the experiments described in Section 5.3.

Learning 3D representations from single images. As discussed in Section 1, to eliminate the cost of collecting 3D data or multiview images, the learning of a 3D representation from single images has garnered attention. Promising approaches involve the use of shape models [19, 32, 75, 77, 93] and the incorporation of auxiliary information such as 2D keypoints [33, 87] or 2D silhouettes [10, 22, 27, 44]. Although such approaches have yielded remarkable results, the requirement for shape models or auxiliary information remains a bottleneck. To eliminate this bottleneck, a *fully* unsupervised learning approach based on generative models has been actively studied. The learning targets differ according to the studies applied, and to date, the unsupervised learning of *viewpoints* [9, 24, 27, 46, 56, 58–61, 65, 73, 76, 84, 96], *albedo* [96], *texture* [84], *light* [96], *3D meshes* [73, 84], *depth* [34, 65, 96], and *defocus effects* [34] has been pro-

posed. Among these approaches, AR-NeRF shares the motivation with AR-GAN [34] with the aim of learning the *depth* and *defocus effects*. However, as the main difference, an AR-GAN represents an aperture renderer in a *discretized* CNN and is specific to the learning of the *depth* and *defocus effects*, whereas our AR-NeRF represents an aperture renderer with *continuous* radiance fields and can explain and utilize other ray-tracing-related phenomena (e.g., *viewpoints*) in a unified manner. We empirically demonstrate these merits in Section 5.2.

Learning of depth and defocus effects. There is a large body of studies conducted on depth learning. Representative approaches involve training the depth predictor using pairs or sets of data, such as image and depth pairs [15, 16, 41, 42, 48, 98], multiview image pairs [18, 21, 97], and consecutive frame sets [91, 101, 104]. Defocus synthesis has also garnered interest in computer vision and graphics, and both model-based [3, 26, 30, 78, 89] and learning-based [29, 70, 83, 92] defocus synthesizers have been proposed. Based on the high correlation between the depth and defocus strength, some studies [25, 83] have proposed learning the depth while reconstructing focused images from all-in-focus images under the assumption that pairs of focused and all-in-focus images are available for training. Although our study is motivated by the success of such studies, the main difference is that we address a challenging but practically important situation in which there are *no* training data available other than natural unstructured (and view-limited) images. The latest model addressing this problem is an AR-GAN [34]. As stated previously, we investigate the quantitative and qualitative differences in Section 5.2.

3. Preliminaries

3.1. GAN

We begin by describing the two previous work upon which our model is built. The first is a GAN [23], which learns the data distribution implicitly through a two-player min-max game using the following objective:

$$\mathcal{L}_{\text{GAN}} = \mathbb{E}_{\mathbf{I}^r \sim p^r(\mathbf{I})} [\log D(\mathbf{I}^r)] + \mathbb{E}_{\mathbf{z} \sim p^g(\mathbf{z})} [\log(1 - D(G(\mathbf{z})))] \quad (1)$$

where, given a latent code \mathbf{z} , a generator G generates an image \mathbf{I}^g that fools the discriminator D by minimizing this objective, whereas D distinguishes \mathbf{I}^g from a real image \mathbf{I}^r by maximizing this objective. Here, the superscripts r and g represent real and generated data, respectively. Through adversarial training, $p^g(\mathbf{I})$ reaches close to $p^r(\mathbf{I})$.

3.2. NeRF

NeRF [56] (in particular, we consider generative variants [9, 76] relevant to our study) represents a scene using an MLP that takes the 3D position $\mathbf{x} \in \mathbb{R}^3$ and view direction $\mathbf{d} \in \mathbb{S}^2$ as inputs and predicts the RGB color $\mathbf{c}(\mathbf{x}, \mathbf{d}) \in \mathbb{R}^3$

and volume density $\sigma(\mathbf{x}) \in \mathbb{R}^+$. More precisely, in [9, 76], positional encoding [56, 86] and sine nonlinearity [80] were used prior to or during the application of the MLP to encode positional information; however, we omit them for a general representation. Moreover, in a generative variant, the MLP also takes the latent code $\mathbf{z} \in \mathbb{R}^{L_z}$ as input to represent a variety of data. However, this is omitted for simplicity.

NeRF employs a pinhole camera (Figure 2(a)) and predicts the color of each pixel $\mathbf{C}(\mathbf{r})$ and the corresponding depth $Z(\mathbf{r})$ by integrating over a single camera ray $\mathbf{r}(t) = \mathbf{o} + t\mathbf{d}$ (where \mathbf{o} and \mathbf{d} are the camera origin and direction, respectively) within a distance $t \in [t_n, t_f]$ using the volume rendering equation [53]:

$$\begin{aligned} \mathbf{C}(\mathbf{r}) &= \int_{t_n}^{t_f} T(t)\sigma(\mathbf{r}(t))\mathbf{c}(\mathbf{r}(t), \mathbf{d})dt, \\ Z(\mathbf{r}) &= \int_{t_n}^{t_f} T(t)\sigma(\mathbf{r}(t))tdt, \\ \text{where } T(t) &= \exp\left(-\int_{t_n}^t \sigma(\mathbf{r}(s))ds\right). \end{aligned} \quad (2)$$

In practice, the integral is intractable; thus, a discretized form with stratified and hierarchical sampling [56] is used.

4. Aperture rendering NeRF: AR-NeRF

4.1. Problem statement

We first clarify the problem statement. We address *fully* unsupervised learning of the depth and defocus effects, where *no* supervision or pretrained models are available and only a collection of unstructured single images are accessible during training. Owing to the lack of explicit supervision, it is difficult to learn a *conditional* model that can directly predict the depth and defocus effects from an input image. As an alternative, we aim to construct an *unconditional* generator $G(\mathbf{z})$ that can generate the image and depth as a set while varying the defocus effects. Similar to viewpoint-aware representation learning, which requires a dataset that includes various view images to acquire the viewpoint cue, our defocus-aware representation learning requires a dataset that includes variously defocused images to obtain a defocus cue. More formally, we impose the following assumption for the dataset:

Assumption 1 *The training images are captured using variously aperture-sized cameras, and the dataset includes diversely defocused images.*

Two factors affecting the defocus effects (as detailed in Section 4.2) are the aperture size and focus distance (distance between the ray origin and the plane where all objects are in focus). Hence, we can also impose the assumption of diversity of the focus distance instead of or in addition to Assumption 1. However, under a practical scenario, the focused target tends to be fixed when the scene is determined.

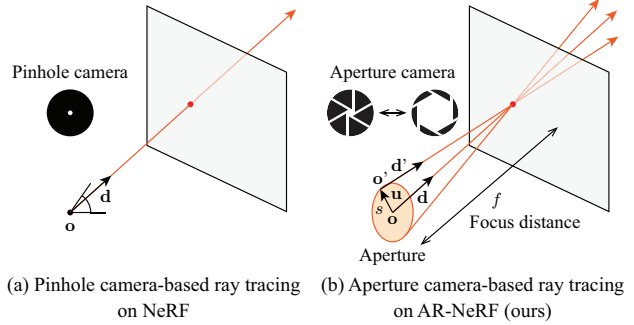


Figure 2. Comparison of ray tracing on NeRF and AR-NeRF.

Hence, in this study, only Assumption 1 is imposed.³

Note that we assume the existence of diversely defocused images but do *not* assume the existence of their *pairs/sets*. We observed that this assumption is satisfied in a typical natural image dataset, as shown in Figure 1.

4.2. Aperture rendering with NeRF

As described in Section 3.2, NeRF is a strongly 3D-aware model that can jointly represent an image and depth at the design level (Equation 2). To utilize this strong property in our problem (Section 4.1), we consider representing aperture rendering in a ray-tracing framework, which is the basis of NeRF. This is achieved by replacing *pinhole camera*-based ray tracing (Figure 2(a)), which is used in a standard NeRF, with *aperture camera*-based ray tracing [79] (Figure 2(b)).

For pinhole-camera-based ray tracing, we cast all rays from a single point \mathbf{o} . By contrast, with aperture camera-based ray tracing, we cast rays from an aperture of radius s . More formally, the origin of a ray from the aperture (\mathbf{o}') is written as

$$\mathbf{o}' = \mathbf{o} + \mathbf{u}, \quad (3)$$

where $|\mathbf{u}| \in [0, s]$, and the direction of \mathbf{u} is orthogonal to \mathbf{o} .

A bundle of rays emitted from the aperture converges to a point on the plane at a focal distance of f . Based on this definition, the direction of the ray from the aperture (\mathbf{d}') is calculated as follows:

$$\mathbf{d}' = (\mathbf{o} + f\mathbf{d} - \mathbf{o}')/f, \quad (4)$$

Based on Equations 3 and 4, we can calculate the ray from the origin \mathbf{o}' , that is, $\mathbf{r}'(t) = \mathbf{o}' + t\mathbf{d}'$, and render the corresponding color $\mathbf{C}(\mathbf{r}')$ and depth $Z(\mathbf{r}')$ using volume rendering (Equation 2). The final color and depth are calculated by integrating over $\mathbf{C}(\mathbf{r}')$ and $Z(\mathbf{r}')$ for all rays in $|\mathbf{u}| \in [0, s]$.⁴ However, similar to volume rendering, the integral is intractable in practice; therefore, a discretized form is used.

³A similar assumption (diversity of the DoF settings) has also been introduced in an AR-GAN [34]. However, it does not distinguish the effects of the two factors, and for a stricter assumption, we redefine it here.

⁴The depth does not need to be integrated, but we used this formulation to account for the ambiguity derived from the defocus blur. However, we empirically found that the effect on the depth accuracy is subtle.

More precisely, we generate a finite bundle of rays from the sampled $|\mathbf{u}| \in [0, s]$ and calculate the final output by taking the average of the corresponding $\mathbf{C}(\mathbf{r}')$ and $Z(\mathbf{r}')$.

4.3. Aperture randomized training

To learn defocus-aware and defocus-independent representations in a disentangled manner, we introduce *aperture randomized training*, in which we learn to generate images by varying the aperture size and latent codes both randomly and independently. More formally, we rewrite the GAN objective (Equation 1) as follows:⁵

$$\mathcal{L}_{\text{AR-NeRF}} = \mathbb{E}_{\mathbf{I}^r \sim p^r(\mathbf{I})} [\log D(\mathbf{I}^r)] + \mathbb{E}_{\mathbf{z} \sim p^g(\mathbf{z}), s \sim p^g(s)} [\log(1 - D(G(\mathbf{z}, s)))] \quad (5)$$

where the latent code \mathbf{z} and aperture size s are sampled randomly and independently. In practice, we represent $p^g(s)$ using a half-normal distribution and parameterize its standard deviation σ_s to determine the range of aperture sizes in a data-driven manner. As a side note, we represent the focus distance f , another variable in aperture rendering, using MLP, taking \mathbf{z} as the input under the assumption that f is determined according to the rendered target.

As discussed in Section 4.2, our aperture rendering has a strong 3D constraint based on ray tracing, and therefore when we train a model using Equation 5, \mathbf{z} must capture the representations that are independent and robust to the change in defocus driven by the fluctuation of s .

4.4. Advanced techniques for practice

To the best of our knowledge, unsupervised learning of the depth and defocus effects is a relatively new task (e.g., the first attempt was at CVPR 2021 [34]), and practical techniques (in particular, those specific to NeRF) have yet to be sufficiently developed. To advance this research direction, we discuss practical techniques considered for this task.

Representation of unbounded background with NeRF++. A typical generative NeRF [9, 76] renders an entire scene in a tightly bounded 3D space to efficiently model the foreground. However, this strategy is problematic when training images with an unbounded background (e.g., bird images in Figure 1). In particular, this problem is critical in terms of learning the defocus effects because its strength is determined according to depth. We cannot represent a strong defocus effect at the design level when using a tightly bounded 3D space. To address this problem, we implemented a synthesis network using NeRF++ [103], which is composed of a foreground NeRF in a unit sphere and background NeRF modeled using inverted sphere

⁵The relevant training scheme is DoF-mixture learning [34], in which the image generation is learned while generating various defocused images. The main difference from our approach is that their method manipulates the depth scale instead of the aperture size because they cannot scale the aperture directly owing to the discretized formulation. Consequently, their learned depth is relative, and therefore they must carefully tune the value when changing the focus distance by adding an offset.

parameterization. This implementation allows representing a strong defocus effect in a far background. For a fair description, we note that concurrent approaches [24, 60] have also incorporated NeRF++ for image generation to represent an unbounded background.

Learning depth and defocus effects with changes in viewpoint. Fully unsupervised learning of the depth and defocus effects is a challenging and ill-posed problem, although our aperture randomized training alleviates this difficulty. To obtain a hint from another source, we jointly learn viewpoint-aware and view-independent representations by randomly sampling the camera poses during training [9, 76]. To prevent sampled camera parameters beyond a real distribution, we restrict its range (using a standard deviation of 0.1 radian in practice). We found that this setting works reasonably well for both datasets, including limited and wide viewpoints (Section 5.3).

Aperture ray sampling scheme. In typical ray tracing used in computer graphics [79], a large number of rays (e.g., 100) are sampled per pixel in aperture rendering (Section 4.2) to improve the synthesis fidelity. However, this increases both the processing time and memory. To efficiently represent the aperture using limited rays, we used stratified sampling [56]. More concretely, we used five rays; the origin of one ray was placed at the center of the aperture, and the origins of the others were placed along the circumference of the aperture with equal intervals. We examine the effect of this approximation in the supplementary material.²

5. Experiments

5.1. Experimental settings

We conducted two experiments to verify the effectiveness of AR-NeRF from multiple perspectives: a comparative study (Section 5.2) in which we compared AR-NeRF to AR-GAN [34], which is a pioneering model with a similar objective, and an ablation study (Section 5.3) in which we investigated the importance of our ideas. In this section, we present the common settings and discuss the details of each in the next sections.

Dataset. Following the AR-GAN study [34], we evaluated AR-NeRF using three natural image datasets: two view-limited datasets, i.e., Oxford Flowers [64] (8,189 flower images with 102 categories) and CUB-200-2011 [90] (11,788 bird images with 200 categories), and a view-various dataset, that is, FFHQ [39] (70,000 face images). To effectively examine various cases, we resized the images to a pixel resolution of 64×64 . This strategy was also used in the AR-GAN study [34]. Therefore, we can compare AR-NeRF to AR-GAN under fair conditions. We provide detailed information about the datasets in the supplementary material.²

Evaluation metrics. We evaluated the effectiveness of AR-NeRF quantitatively using the same two metrics used in the AR-GAN study [34] for a direct comparison. The first is

the *kernel inception distance (KID)* [4], which measures the maximum mean discrepancy between real and generated images within the inception model [85]. We used the KID to evaluate the visual fidelity of the generated images. We calculated this score using 20,000 generated images and all real images. Based on our objective (i.e., training an unconditional model on unstructured natural images), preparing the ground truth depth is nontrivial. Following [34], as an alternative, we calculated the depth accuracy by (1) training the depth predictor using pairs of images and depths generated through GANs, (2) predicting the depths of real images using the trained depth predictor, and (3) comparing the predicted depths to those predicted by a highly generalizable monocular depth estimator [97] trained using stereo pairs.⁶ To measure the differences in depth, we used the *scale-invariant depth error (SIDE)* [15], which measures the difference between depths in a scale-invariant manner and is useful for comparing the depths predicted by different representation systems. For both metrics, the smaller the value, the better the performance.

Implementation. We implemented AR-NeRF based on pi-GAN [9],⁷ which is a state-of-the-art generative variant of NeRF. Because the original pi-GAN was not applied to the dataset used in our experiments, we carefully tuned the configurations and hyperparameters such that the baseline pi-GAN could generate images reasonably well. Next, we incorporated a background synthesis network into pi-GAN based on NeRF++ [103]⁸ (Section 4.4). Hereafter, we refer to this model as *pi-GAN++*. Subsequently, we incorporated aperture rendering (Section 4.2) and aperture randomized training (Section 4.3) into *pi-GAN++*. This is the model denoted by *AR-NeRF* below. We provide implementation details in the supplementary material.²

5.2. Comparative study

To determine the validity of AR-NeRF for unsupervised learning of the *depth* and *defocus effects*, we first investigated the comparative performance between AR-NeRF and AR-GAN [34], which is a state-of-art model for this problem. The main difference between AR-NeRF and AR-GAN is the architectural difference, where AR-NeRF is constructed based on the *continuous* radiance fields, whereas AR-GAN is constructed based on *discretized* CNNs. Another small but significant difference is that AR-NeRF represents the defocus distribution (i.e., aperture size distribution) using a half-normal distribution (Equation 5), whereas AR-GAN represents it (i.e., depth scale distribution in this case) using a binomial distribution (Equation 6 in [34]). To confirm the effects of this difference, we also evaluated a variant of AR-GAN (referred to as *AR-GAN++*), in which

⁶We used the official pretrained model: <https://github.com/KexianHust/Structure-Guided-Ranking-Loss>.

⁷We implemented it based on the official code <https://github.com/marcoamonteiro/pi-GAN>.

⁸We implemented this while referring to the official code: <https://github.com/Kai-46/nerfplusplus>.

	Oxford Flowers		CUB-200-2011		FFHQ	
	KID↓	SIDE↓	KID↓	SIDE↓	KID↓	SIDE↓
AR-GAN	11.23	4.46	14.30	3.58	5.75	4.21
AR-GAN++	10.18	4.42	13.91	3.61	5.43	4.88
RGBD-GAN	12.04	7.01	14.92	7.06	6.73	5.81
AR-NeRF (ours)	7.86	3.94	6.81	3.63	3.67	2.61

Table 1. **Comparison of KID↓ ($\times 10^3$) and SIDE↓ ($\times 10^2$) between baseline GANs and AR-NeRF (ours).**

the defocus distribution was expressed using a half-normal distribution, similar to AR-NeRF. Furthermore, as a reference, we report the scores of RGBD-GAN [65], which learns the depth information using viewpoint cues.⁹

Quantitative comparisons. We summarize quantitative comparison results in Table 1. AR-NeRF outperformed the baseline GANs in terms of the KID and SIDE, except for SIDE on CUB-200-2011, where AR-GAN/AR-GAN++ was comparable to AR-NeRF.¹⁰ These results validate the utility of AR-NeRF for unsupervised learning of *depth*. We believe that the strengths of AR-NeRF, i.e., the joint usage of the viewpoint and defocus cues and continuous representations based on implicit functions, contribute to this improvement. We present qualitative comparisons of the predicted depths in the supplementary material.²

Qualitative comparisons. We conducted qualitative comparisons to validate the effectiveness of unsupervised learning of the *defocus effects*. We present examples of generated images and depths in Figure 3. In AR-NeRF, we manipulated the defocus strength and focus distance by changing s and f (Figure 2(b)). As discussed above, the original AR-GAN discretely represents the defocus distribution. Therefore, differently from AR-NeRF, it is unsuitable for conducting continuous operations. Alternatively, we examine the performance of AR-GAN++, which represents a continuous defocus distribution. In AR-GAN++, we manipulated the defocus strength and focus distance by changing the scale and offset of depth, respectively.

The results indicate that AR-NeRF can manipulate both the defocus strength and focus distance without generating significant artifacts. In particular, in the manipulation of the focus distance, AR-NeRF succeeds in refocusing on both the foreground and background, the appearances of which are the same as those in the all-in-focus images (in the leftmost column). By contrast, AR-GAN++ often generates

⁹In our preliminary experiments, we also examined the performances of AR-HoloGAN/AR-RGBD-GAN (combinations of AR-GAN and HoloGAN [58]/RGBD-GAN [65]) on FFHQ. We found that the SIDE scores for AR-HoloGAN and AR-RGBD-GAN were 4.79 and 4.40, respectively, and were worse than those for AR-GAN. This result indicates that the simultaneous but individual usage of the viewpoint and defocus cues in AR-HoloGAN/AR-RGBD-GAN does not improve the depth learning.

¹⁰SIDE has a limitation in that it can ignore certain types of degradation because it measures the difference based on l_2 , causing statistical averaging. This may be why SIDE is comparable on CUB-200-2011 despite the qualitative differences (Figure 3). To validate this hypothesis, we analyzed the gradient of the difference between the ground truth and predicted depths [14] and found that AR-NeRF outperformed AR-GAN/AR-GAN++ on this metric. We discuss the details in the supplementary material.²



Figure 3. **Comparison of generated images and depths between AR-GAN++ and AR-NeRF (ours).** To manipulate the defocus strength, we varied the strength within $[0, \sigma_s, 2\sigma_s, 3\sigma_s]$, where σ_s indicates the standard deviation of a half-normal distribution, which is used to represent the defocus distribution during training. To manipulate the focus distance, we used a range in which the foregrounds and backgrounds were focused.

unexpected artifacts, particularly when it attempts to refocus on the background (in the second-to-last column). As possible causes for this phenomenon, (1) AR-GAN++ discretely represents light fields in a 2D space; thus, the discretization error becomes critical when large manipulations are counted and (2) the predicted depths include artifacts (e.g., holes appearing in objects), causing errors when images are rendered based on depth. The properties of AR-NeRF, that is, (1) a continuous representation in a 3D space and (2) joint usage of defocus and viewpoint cues, are useful for addressing these defects. As another advantage of AR-NeRF, it can increase the resolution of the generated images by increasing the density of the input points owing to the nature of the implicit function [9]. We demonstrate this strength in Figure 1, where 128×128 images are generated using the same model as that used in Figure 3.

5.3. Ablation study

We conducted ablation studies to examine the utility of AR-NeRF as a generative variant of the NeRF. We compared AR-NeRF to five baselines: *pi-GAN* [9], where a background synthesis network and aperture rendering are ablated; *pi-GAN++*, where aperture rendering is ablated; *AR-NeRF-0*, where viewpoint changes (Section 4.4) are not applied during training; *AR-NeRF-F*, where the full viewpoint changes that are optimized to the face dataset (FFHQ) are used; and *pi-GAN++-F*, where aperture rendering is ablated from *AR-NeRF-F*. We tested the last two models on FFHQ only because viewpoint cues were limited on the

				Oxford Flowers		CUB-200-2011		FFHQ	
	(B)	(D)	(V)	KID↓	SIDE↓	KID↓	SIDE↓	KID↓	SIDE↓
pi-GAN			L	3.69	5.23	5.04	4.87	4.29	3.03
pi-GAN++	✓		L	8.30	4.83	9.84	3.88	4.43	2.69
AR-NeRF-0	✓	✓	0	6.81	4.03	8.67	3.74	3.83	3.61
AR-NeRF-F	✓	✓	F	–	–	–	–	4.59	2.75
pi-GAN++-F	✓		F	–	–	–	–	5.06	2.78
AR-NeRF	✓	✓	L	7.86	3.94	6.81	3.63	3.67	2.61

Table 2. **Comparison of KID↓ ($\times 10^3$) and SIDE↓ ($\times 10^2$) between AR-NeRF and ablated models.** Check marks (B) and (D) indicate the use of a background synthesis network and defocus cue, respectively. In column (V), L, F, and 0 indicate the use of local, full, and no viewpoint changes, respectively.

other datasets. For *pi-GAN++*, we set the number of rays to be the same as that in AR-NeRF by an ensemble of multiple rays with an aperture size of $s = 0$. We used this implementation to investigate the pure performance differences between the models with and without aperture rendering.

Results. We list the quantitative results in Table 2 and provide a qualitative comparison of the predicted depths in the supplementary material.² Our findings are as follows:

(1) *Effects of the background synthesis network (pi-GAN vs. pi-GAN++).* We found that pi-GAN outperforms pi-GAN++ in terms of the KID. We consider that the compact representation of the pi-GAN is advantageous for efficiently learning 2D image distributions. However, pi-GAN was outperformed by pi-GAN++ in terms of the SIDE. This result indicates that pi-GAN is unsuitable for our aims (i.e., unsupervised learning of depth information) despite its abil-

ity to generate high-fidelity images.

(2) *Effects of aperture rendering (pi-GAN++ vs. AR-NeRF)*. We found that AR-NeRF outperformed pi-GAN++ on both metrics, except for SIDE on FFHQ, where pi-GAN++ was comparable to AR-NeRF. The same tendency holds for the comparison between AR-NeRF-F and pi-GAN++-F. This is because FFHQ includes sufficient viewpoint variations to leverage the viewpoint cues. By contrast, Oxford Flowers and CUB-200-2011 do not contain them. In this case, the defocus cues used in AR-NeRF contributed to an improvement.

(3) *Comparison between viewpoint and defocus cues (pi-GAN++ vs. AR-NeRF-0)*. With these models, the defocus and viewpoint manipulations are ablated. Therefore, we can analyze each effect by comparing them. We found that, in FFHQ, pi-GAN++ outperformed AR-NeRF-0 in terms of SIDE, whereas in the other datasets, AR-NeRF-0 outperformed pi-GAN++. This can be explained by differences in the availability of the viewpoint cues, as discussed in (2).

(4) *Comparison between local and full viewpoint changes (AR-NeRF vs. AR-NeRF-F)*. We found that AR-NeRF outperformed AR-NeRF-F on both metrics. This result indicates that we do not need to carefully tune the camera parameters for unsupervised depth learning. The same tendency was observed in the comparison between pi-GAN++ and pi-GAN++-F. Note that AR-NeRF-F has an advantage in the viewpoint manipulation capability because it can learn full view variations, whereas AR-NeRF can only learn local view variations.

Detailed analyses. For further analyses, we examined (1) the importance of learning defocus effects from images, (2) the effect of the aperture ray sampling scheme, (3) simultaneous control of the viewpoint and defocus, (4) generation of higher-resolution images, (5) application to defocus renders, (6) the Fréchet inception distance (FID) [28], and (7) the gradient of the difference in depth [14]. See the supplementary material² for further details.

6. Discussion

6.1. Limitations and future work

AR-NeRF has two limitations, stemming from radiance field representations and fully unsupervised learning.

Limitations caused by radiance field representations. In radiance field representations, the computational complexity increases not only with the image size but also with the depth along each ray. Consequently, the calculation cost is higher than that of a CNN GAN (e.g., AR-GAN [34]); therefore, application to high-resolution images is difficult. AR-NeRF requires multiple rays per pixel to represent aperture rendering. Therefore, it incurs a larger calculation cost than the standard NeRF, which represents each pixel using a single ray. Through our experiments, we found that AR-NeRF outperforms the baseline NeRF, which has a similar calculation cost (in particular, pi-GAN++).



Figure 4. Failure case.

This demonstrates the validity of our research direction. However, improving the calculation cost remains an important topic for future research. Recent concurrent studies [17, 24, 47, 71, 81, 102] have addressed reducing the calculation cost of NeRF, and the incorporation of these methods is also a promising research area.

Limitations caused by fully unsupervised learning. Fully unsupervised learning of the depth and defocus effects is highly challenging, and some limitations remain. During our experiments, we found that our model is better than or comparable to the models trained under the same conditions. However, its performance is lower than that of supervised models. In particular, application to complex images will be difficult because AR-NeRF is a generative approach that assumes that it can learn the image generation reasonably well. Furthermore, the use of an unbounded background based on NeRF++ [103] allows strong defocus effects that occur in the far plane to be represented. However, it is still difficult to distinguish the defocus blur from the flat texture when the defocus blur is extremely strong (e.g., Figure 4). Addressing these problems is a possible direction for future research.

6.2. Potential negative social impact

The method presented in this paper enables the creation of realistic images. This poses a potential risk to the creation of misleading content (e.g., deepfake). In particular, our model can increase the credibility of fake content in terms of 3D consistency and may potentially deceive systems that rely on 3D structures, such as face recognition systems. Therefore, we believe that it is essential for the community to develop technology to distinguish fake images from real images and carefully monitor advancements in the corresponding research fields [1, 45, 52, 57, 72].

7. Conclusion

To advance the research on the fully unsupervised learning of depth and defocus effects, we introduced AR-NeRF, which extends NeRF by incorporating aperture rendering. AR-NeRF is noteworthy because it can employ defocus and viewpoint cues in a unified manner by representing both factors through a common ray-tracing framework. We empirically demonstrated the effectiveness of AR-NeRF for unsupervised learning of the depth and defocus effects. Although we focused on a generative variant of NeRF in this study, our idea, that is, the incorporation of aperture rendering in NeRF, is general, and we expect that its usage will broaden the applications of NeRF under practical scenarios.

References

- [1] Shruti Agarwal, Hany Farid, Yuming Gu, Mingming He, Koki Nagano, and Hao Li. Protecting world leaders against deep fakes. In *CVPR Workshops*, 2019. 8
- [2] Matan Atzmon and Yaron Lipman. SAL: Sign agnostic learning of shapes from raw data. In *CVPR*, 2020. 3
- [3] Jonathan T. Barron, Andrew Adams, YiChang Shih, and Carlos Hernández. Fast bilateral-space stereo for synthetic defocus. In *CVPR*, 2015. 3
- [4] Mikotaj Bińkowski, Dougal J. Sutherland, Michael Arbel, and Arthur Gretton. Demystifying MMD GANs. In *ICLR*, 2018. 6
- [5] Volker Blanz and Thomas Vetter. A morphable model for the synthesis of 3D faces. In *SIGGRAPH*, 1999. 1
- [6] Ashish Bora, Eric Price, and Alexandros G. Dimakis. AmbientGAN: Generative models from lossy measurements. In *ICLR*, 2018. 3
- [7] Andrew Brock, Jeff Donahue, and Karen Simonyan. Large scale GAN training for high fidelity natural image synthesis. In *ICLR*, 2019. 3
- [8] Rohan Chabra, Jan E. Lenssen, Eddy Ilg, Tanner Schmidt, Julian Straub, Steven Lovegrove, and Richard Newcombe. Deep local shapes: Learning local SDF priors for detailed 3D reconstruction. In *ECCV*, 2020. 3
- [9] 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 *CVPR*, 2021. 2, 3, 4, 5, 6, 7
- [10] Wenzheng Chen, Huan Ling, Jun Gao, Edward Smith, Jaakko Lehtinen, Alec Jacobson, and Sanja Fidler. Learning to predict 3D objects with an interpolation-based differentiable renderer. In *NeurIPS*, 2019. 2, 3
- [11] Zhiqin Chen, Kangxue Yin, Matthew Fisher, Siddhartha Chaudhuri, and Hao Zhang. BAE-NET: Branched autoencoder for shape co-segmentation. In *ICCV*, 2019. 3
- [12] Zhiqin Chen and Hao Zhang. Learning implicit fields for generative shape modeling. In *CVPR*, 2019. 3
- [13] Julian Chibane, Aymen Mir, and Gerard Pons-Moll. Neural unsigned distance fields for implicit function learning. In *NeurIPS*, 2020. 3
- [14] David Eigen and Rob Fergus. Predicting depth, surface normals and semantic labels with a common multi-scale convolutional architecture. In *ICCV*, 2015. 6, 8
- [15] David Eigen, Christian Puhrsch, and Rob Fergus. Depth map prediction from a single image using a multi-scale deep network. In *NIPS*, 2014. 3, 6
- [16] Huan Fu, Mingming Gong, Chaohui Wang, Kayhan Batmanghelich, and Dacheng Tao. Deep ordinal regression network for monocular depth estimation. In *CVPR*, 2018. 3
- [17] Stephan J. Garbin, Marek Kowalski, Matthew Johnson, Jamie Shotton, and Julien Valentin. FastNeRF: High-fidelity neural rendering at 200FPS. In *ICCV*, 2021. 8
- [18] Ravi Garg, Vijay Kumar B G, Gustavo Carneiro, and Ian Reid. Unsupervised CNN for single view depth estimation: Geometry to the rescue. In *ECCV*, 2016. 3
- [19] Baris Gecer, Stylianos Ploumpis, Irene Kotsia, and Stefanos Zafeiriou. GANFIT: Generative adversarial network fitting for high fidelity 3D face reconstruction. In *CVPR*, 2019. 1, 3
- [20] Kyle Genova, Forrester Cole, Daniel Vlasic, Aaron Sarna, William T. Freeman, and Thomas Funkhouser. Learning shape templates with structured implicit functions. In *ICCV*, 2019. 3
- [21] Clément Godard, Oisín Mac Aodha, and Gabriel J. Brostow. Unsupervised monocular depth estimation with left-right consistency. In *CVPR*, 2017. 3
- [22] Shubham Goel, Angjoo Kanazawa, and Jitendra Malik. Shape and viewpoint without keypoints. In *ECCV*, 2020. 2, 3
- [23] Ian J. Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron Courville, and Yoshua Bengio. Generative adversarial nets. In *NIPS*, 2014. 2, 3
- [24] Jiatao Gu, Lingjie Liu, Peng Wang, and Christian Theobalt. StyleNeRF: A style-based 3D-aware generator for high-resolution image synthesis. In *ICLR*, 2022. 2, 3, 5, 8
- [25] Shir Gur and Lior Wolf. Single image depth estimation trained via depth from defocus cues. In *CVPR*, 2019. 3
- [26] Samuel W. Hasinoff and Kiriakos N. Kutulakos. A layer-based restoration framework for variable-aperture photography. In *ICCV*, 2007. 3
- [27] Philipp Henzler, Niloy Mitra, and Tobias Ritschel. Escaping Plato’s cave using adversarial training: 3D shape from unstructured 2D image collections. In *ICCV*, 2019. 2, 3
- [28] 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 *NIPS*, 2017. 8
- [29] Andrey Ignatov, Jagruti Patel, and Radu Timofte. Rendering natural camera bokeh effect with deep learning. In *CVPR Workshops*, 2020. 3
- [30] David E. Jacobs, Jongmin Baek, and Marc Levoy. Focal stack compositing for depth of field control. *Stanford Computer Graphics Laboratory Technical Report*, 1(1):2012, 2012. 3
- [31] Chiyu Max Jiang, Avneesh Sud, Ameet Makadia, Jingwei Huang, Matthias Nießner, and Thomas Funkhouser. Local implicit grid representations for 3D scenes. In *CVPR*, 2020. 3
- [32] Angjoo Kanazawa, Michael J. Black, David W. Jacobs, and Jitendra Malik. End-to-end recovery of human shape and pose. In *CVPR*, 2018. 1, 3
- [33] Angjoo Kanazawa, Shubham Tulsiani, Alexei A. Efros, and Jitendra Malik. Learning category-specific mesh reconstruction from image collections. In *ECCV*, 2018. 2, 3
- [34] Takuhiro Kaneko. Unsupervised learning of depth and depth-of-field effect from natural images with aperture rendering generative adversarial networks. In *CVPR*, 2021. 2, 3, 4, 5, 6, 8
- [35] Takuhiro Kaneko and Tatsuya Harada. Noise robust generative adversarial networks. In *CVPR*, 2020. 3
- [36] Takuhiro Kaneko and Tatsuya Harada. Blur, noise, and compression robust generative adversarial networks. In *CVPR*, 2021. 3
- [37] Tero Karras, Timo Aila, Samuli Laine, and Jaakko Lehtinen. Progressive growing of GANs for improved quality, stability, and variation. In *ICLR*, 2017. 3
- [38] Tero Karras, Miika Aittala, Samuli Laine, Erik Härkönen, Janne Hellsten, Jaakko Lehtinen, and Timo Aila. Alias-free generative adversarial networks. In *NeurIPS*, 2021. 3

- [39] Tero Karras, Samuli Laine, and Timo Aila. A style-based generator architecture for generative adversarial networks. In *CVPR*, 2019. 2, 3, 5
- [40] Tero Karras, Samuli Laine, Miika Aittala, Janne Hellsten, Jaakko Lehtinen, and Timo Aila. Analyzing and improving the image quality of StyleGAN. In *CVPR*, 2020. 3
- [41] Yevhen Kuznietsov, Jorg Stuckler, and Bastian Leibe. Semi-supervised deep learning for monocular depth map prediction. In *CVPR*, 2017. 3
- [42] Iro Laina, Christian Rupprecht, Vasileios Belagiannis, Federico Tombari, and Nassir Navab. Deeper depth prediction with fully convolutional residual networks. In *3DV*, 2016. 3
- [43] Steven Cheng-Xian Li, Bo Jiang, and Benjamin Marlin. MisGAN: Learning from incomplete data with generative adversarial networks. In *ICLR*, 2019. 3
- [44] Xueting Li, Sifei Liu, Kihwan Kim, Shalini De Mello, Varun Jampani, Ming-Hsuan Yang, and Jan Kautz. Self-supervised single-view 3D reconstruction via semantic consistency. In *ECCV*, 2020. 2, 3
- [45] Yuezun Li, Xin Yang, Pu Sun, Honggang Qi, and Siwei Lyu. Celeb-DF: A large-scale challenging dataset for Deep-Fake forensics. In *CVPR*, 2020. 8
- [46] Yiyi Liao, Katja Schwarz, Lars Mescheder, and Andreas Geiger. Towards unsupervised learning of generative models for 3D controllable image synthesis. In *CVPR*, 2020. 2, 3
- [47] David B. Lindell, Julien N. P. Martel, and Gordon Wetstein. AutoInt: Automatic integration for fast neural volume rendering. In *CVPR*, 2021. 8
- [48] Fayao Liu, Chunhua Shen, Guosheng Lin, and Ian Reid. Learning depth from single monocular images using deep convolutional neural fields. *IEEE Trans. Pattern Anal. Mach. Intell.*, 38(10):2024–2039, 2015. 3
- [49] Shichen Liu, Shunsuke Saito, Weikai Chen, and Hao Li. Learning to infer implicit surfaces without 3D supervision. In *NeurIPS*, 2019. 3
- [50] 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 *CVPR*, 2020. 3
- [51] Matthew Loper, Naureen Mahmood, Javier Romero, Gerard Pons-Moll, and Michael J. Black. SMPL: A skinned multi-person linear model. *ACM Trans. Graph.*, 34(6):1–16, 2015. 1
- [52] Francesco Marra, Diego Gagnaniello, Davide Cozzolino, and Luisa Verdoliva. Detection of GAN-generated fake images over social networks. In *MIPR*, 2018. 8
- [53] Nelson Max. Optical models for direct volume rendering. *IEEE Trans Vis Comput Graph*, 1(2):99–108, 1995. 4
- [54] Lars Mescheder, Michael Oechsle, Michael Niemeyer, Sebastian Nowozin, and Andreas Geiger. Occupancy networks: Learning 3D reconstruction in function space. In *CVPR*, 2019. 3
- [55] Mateusz Michalkiewicz, Jhony K. Pontes, Dominic Jack, Mahsa Baktashmotlagh, and Anders Eriksson. Implicit surface representations as layers in neural networks. In *ICCV*, 2019. 3
- [56] Ben Mildenhall, Pratul P. Srinivasan, Matthew Tancik, Jonathan T. Barron, Ravi Ramamoorthi, and Ren Ng. NeRF: Representing scenes as neural radiance fields for view synthesis. In *ECCV*, 2020. 2, 3, 4, 5
- [57] Lakshmanan Nataraj, Tajuddin Manhar Mohammed, B. S. Manjunath, Shivkumar Chandrasekaran, Arjuna Flenner, Jawadul H. Bappy, and Amit K. Roy-Chowdhury. Detecting GAN generated fake images using co-occurrence matrices. In *EI*, 2019. 8
- [58] Thu Nguyen-Phuoc, Chuan Li, Lucas Theis, Christian Richardt, and Yong-Liang Yang. HoloGAN: Unsupervised learning of 3D representations from natural images. In *ICCV*, 2019. 2, 3, 6
- [59] Thu Nguyen-Phuoc, Christian Richardt, Long Mai, Yong-Liang Yang, and Niloy Mitra. BlockGAN: Learning 3D object-aware scene representations from unlabelled images. In *NeurIPS*, 2020. 2, 3
- [60] Michael Niemeyer and Andreas Geiger. CAMPARI: Camera-aware decomposed generative neural radiance fields. In *3DV*, 2021. 2, 3, 5
- [61] Michael Niemeyer and Andreas Geiger. GIRAFFE: Representing scenes as compositional generative neural feature fields. In *CVPR*, 2021. 2, 3
- [62] Michael Niemeyer, Lars Mescheder, Michael Oechsle, and Andreas Geiger. Occupancy flow: 4D reconstruction by learning particle dynamics. In *ICCV*, 2019. 3
- [63] Michael Niemeyer, Lars Mescheder, Michael Oechsle, and Andreas Geiger. Differentiable volumetric rendering: Learning implicit 3D representations without 3D supervision. In *CVPR*, 2020. 3
- [64] Maria-Elena Nilsback and Andrew Zisserman. Automated flower classification over a large number of classes. In *ICVGIP*, 2008. 2, 5
- [65] Atsuhiko Noguchi and Tatsuya Harada. RGBD-GAN: Unsupervised 3D representation learning from natural image datasets via RGBD image synthesis. In *ICLR*, 2020. 2, 3, 6
- [66] Michael Oechsle, Lars Mescheder, Michael Niemeyer, Thilo Strauss, and Andreas Geiger. Texture fields: Learning texture representations in function space. In *ICCV*, 2019. 3
- [67] Arthur Pajot, Emmanuel de Bezenac, and Patrick Gallinari. Unsupervised adversarial image reconstruction. In *ICLR*, 2018. 3
- [68] Jeong Joon Park, Peter Florence, Julian Straub, Richard Newcombe, and Steven Lovegrove. DeepSDF: Learning continuous signed distance functions for shape representation. In *CVPR*, 2019. 3
- [69] Songyou Peng, Michael Niemeyer, Lars Mescheder, Marc Pollefeys, and Andreas Geiger. Convolutional occupancy networks. In *ECCV*, 2020. 3
- [70] Ming Qian, Congyu Qiao, Jiamin Lin, Zhenyu Guo, Chenghua Li, Cong Leng, and Jian Cheng. BGGAN: Bokeh-glass generative adversarial network for rendering realistic bokeh. In *ECCV Workshops*, 2020. 3
- [71] Christian Reiser, Songyou Peng, Yiyi Liao, and Andreas Geiger. KiloNeRF: Speeding up neural radiance fields with thousands of tiny MLPs. In *ICCV*, 2021. 8
- [72] Andreas Rossler, Davide Cozzolino, Luisa Verdoliva, Christian Riess, Justus Thies, and Matthias Nießner. FaceForensics++: Learning to detect manipulated facial images. In *ICCV*, 2019. 8

- [73] Mihir Sahasrabudhe, Zhixin Shu, Edward Bartrum, Riza Alp Guler, Dimitris Samaras, and Iasonas Kokkinos. Lifting AutoEncoders: Unsupervised learning of a fully-disentangled 3D morphable model using deep non-rigid structure from motion. In *ICCV Workshops*, 2019. 3
- [74] Shunsuke Saito, Zeng Huang, Ryota Natsume, Shigeo Morishima, Angjoo Kanazawa, and Hao Li. PIFu: Pixel-aligned implicit function for high-resolution clothed human digitization. In *ICCV*, 2019. 3
- [75] Soubhik Sanyal, Timo Bolkart, Haiwen Feng, and Michael J. Black. Learning to regress 3D face shape and expression from an image without 3D supervision. In *CVPR*, 2019. 1, 3
- [76] Katja Schwarz, Yiyi Liao, Michael Niemeyer, and Andreas Geiger. GRAF: Generative radiance fields for 3D-aware image synthesis. In *NeurIPS*, 2020. 2, 3, 4, 5
- [77] Jiayang Shang, Tianwei Shen, Shiwei Li, Lei Zhou, Mingmin Zhen, Tian Fang, and Long Quan. Self-supervised monocular 3D face reconstruction by occlusion-aware multi-view geometry consistency. In *ECCV*, 2020. 1, 3
- [78] Xiaoyong Shen, Xin Tao, Hongyun Gao, Chao Zhou, and Jiaya Jia. Deep automatic portrait matting. In *ECCV*, 2016. 3
- [79] Peter Shirley. Ray tracing in one weekend. <https://raytracing.github.io/books/RayTracingInOneWeekend.html>, December 2020. 2, 4, 5
- [80] Vincent Sitzmann, Julien Martel, Alexander Bergman, David Lindell, and Gordon Wetzstein. Implicit neural representations with periodic activation functions. In *NeurIPS*, 2020. 4
- [81] Vincent Sitzmann, Semon Rezchikov, William T. Freeman, Joshua B. Tenenbaum, and Fredo Durand. Light field networks: Neural scene representations with single-evaluation rendering. In *NeurIPS*, 2021. 8
- [82] Vincent Sitzmann, Michael Zollhöfer, and Gordon Wetzstein. Scene representation networks: Continuous 3D-structure-aware neural scene representations. In *NeurIPS*, 2019. 3
- [83] Pratul P. Srinivasan, Rahul Garg, Neal Wadhwa, Ren Ng, and Jonathan T. Barron. Aperture supervision for monocular depth estimation. In *CVPR*, 2018. 2, 3
- [84] Attila Szabó, Givi Meishvili, and Paolo Favaro. Unsupervised generative 3D shape learning from natural images. *arXiv preprint arXiv:1910.00287*, 2019. 3
- [85] Christian Szegedy, Vincent Vanhoucke, Sergey Ioffe, Jon Shlens, and Zbigniew Wojna. Rethinking the Inception architecture for computer vision. In *CVPR*, 2016. 6
- [86] 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 *NeurIPS*, 2020. 4
- [87] Luan Tran and Xiaoming Liu. Nonlinear 3D face morphable model. In *CVPR*, 2018. 2, 3
- [88] Carl Vondrick, Hamed Pirsiavash, and Antonio Torralba. Generating videos with scene dynamics. In *NIPS*, 2016. 3
- [89] Neal Wadhwa, Rahul Garg, David E. Jacobs, Bryan E. Feldman, Nori Kanazawa, Robert Carroll, Yair Movshovitz-Attias, Jonathan T. Barron, Yael Pritch, and Marc Levoy. Synthetic depth-of-field with a single-camera mobile phone. *ACM Trans. Graph.*, 37(4):1–13, 2018. 3
- [90] Catherine Wah, Steve Branson, Peter Welinder, Pietro Perona, and Serge Belongie. The Caltech-UCSD Birds-200-2011 Dataset. Technical Report CNS-TR-2011-001, California Institute of Technology, 2011. 2, 5
- [91] Chaoyang Wang, José Miguel Buenaposada, Rui Zhu, and Simon Lucey. Learning depth from monocular videos using direct methods. In *CVPR*, 2018. 3
- [92] Lijun Wang, Xiaohui Shen, Jianming Zhang, Oliver Wang, Zhe Lin, Chih-Yao Hsieh, Sarah Kong, and Huchuan Lu. DeepLens: Shallow depth of field from a single image. *ACM Trans. Graph.*, 37(6):1–11, 2018. 3
- [93] Mengjiao Wang, Zhixin Shu, Shiyang Cheng, Yannis Panagakis, Dimitris Samaras, and Stefanos Zafeiriou. An adversarial neuro-tensorial approach for learning disentangled representations. *Int. J. Comput. Vis.*, 127(6-7):743–762, 2019. 1, 3
- [94] Xiaolong Wang and Abhinav Gupta. Generative image modeling using style and structure adversarial networks. In *ECCV*, 2016. 3
- [95] 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 *NIPS*, 2016. 3
- [96] Shangzhe Wu, Christian Rupprecht, and Andrea Vedaldi. Unsupervised learning of probably symmetric deformable 3D objects from images in the wild. In *CVPR*, 2020. 3
- [97] Ke Xian, Jianming Zhang, Oliver Wang, Long Mai, Zhe Lin, and Zhiguo Cao. Structure-guided ranking loss for single image depth prediction. In *CVPR*, 2020. 3, 6
- [98] Dan Xu, Elisa Ricci, Wanli Ouyang, Xiaogang Wang, and Nicu Sebe. Multi-scale continuous CRFs as sequential deep networks for monocular depth estimation. In *CVPR*, 2017. 3
- [99] Jianwei Yang, Anitha Kannan, Dhruv Batra, and Devi Parikh. LR-GAN: Layered recursive generative adversarial networks for image generation. In *ICLR*, 2017. 3
- [100] 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 *NeurIPS*, 2020. 3
- [101] Zhichao Yin and Jianping Shi. GeoNet: Unsupervised learning of dense depth, optical flow and camera pose. In *CVPR*, 2018. 3
- [102] Alex Yu, Ruilong Li, Matthew Tancik, Hao Li, Ren Ng, and Angjoo Kanazawa. PlenOctrees for real-time rendering of neural radiance fields. In *ICCV*, 2021. 8
- [103] Kai Zhang, Gernot Riegler, Noah Snavely, and Vladlen Koltun. NeRF++: Analyzing and improving neural radiance fields. *arXiv preprint arXiv:2010.07492*, 2020. 5, 6, 8
- [104] Tinghui Zhou, Matthew Brown, Noah Snavely, and David G. Lowe. Unsupervised learning of depth and ego-motion from video. In *CVPR*, 2017. 3
- [105] 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 representation. In *NeurIPS*, 2018. 3