

SCADE: NeRFs from Space Carving with Ambiguity-Aware Depth Estimates

Mikaela Angelina Uy^{1,2} Ricardo Martin-Brualla² Leonidas Guibas^{1,2} Ke Li^{2,3}

¹Stanford University ²Google ³Simon Fraser University

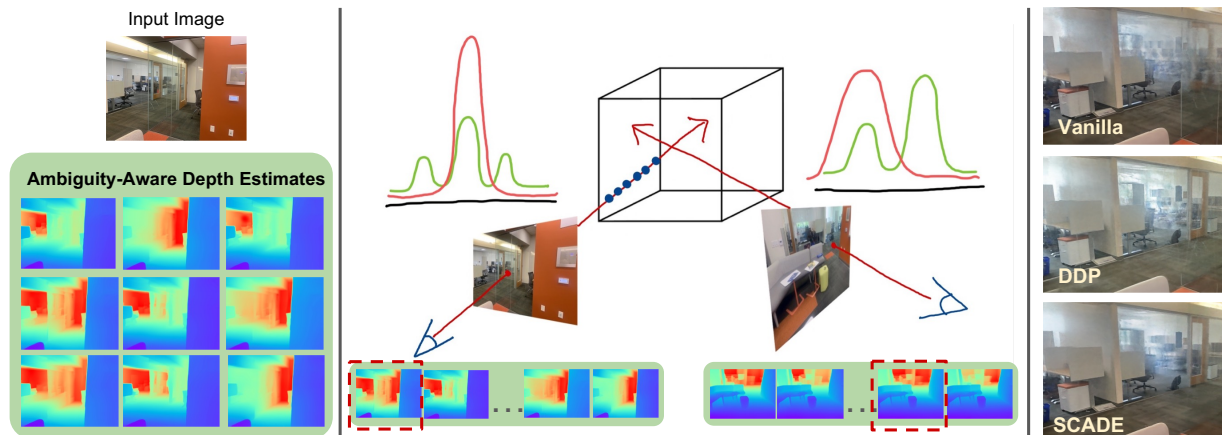


Figure 1. **SCADE Overview.** We present SCADE, a novel technique for NeRF reconstruction under *sparse, unconstrained views* for in-the-wild indoor scenes. We leverage on generalizable monocular depth priors and address to represent the inherent ambiguities of monocular depth by exploiting our **ambiguity-aware depth estimates** (left). Our approach accounts for *multimodality* of both distributions using our novel *space carving loss* that seeks to *disambiguate* and find the common mode to *fuse* the information between different views (middle). SCADE enables better photometric reconstruction especially in highly ambiguous scenes, *e.g.* non-opaque surfaces (right).

Abstract

Neural radiance fields (NeRFs) have enabled high fidelity 3D reconstruction from multiple 2D input views. However, a well-known drawback of NeRFs is the less-than-ideal performance under a small number of views, due to insufficient constraints enforced by volumetric rendering. To address this issue, we introduce SCADE, a novel technique that improves NeRF reconstruction quality on sparse, unconstrained input views for in-the-wild indoor scenes. To constrain NeRF reconstruction, we leverage geometric priors in the form of per-view depth estimates produced with state-of-the-art monocular depth estimation models, which can generalize across scenes. A key challenge is that monocular depth estimation is an ill-posed problem, with inherent ambiguities. To handle this issue, we propose a new method that learns to predict, for each view, a continuous, multimodal distribution of depth estimates using conditional Implicit Maximum Likelihood Estimation (cIMLE). In order to disambiguate exploiting multiple views, we introduce an original space carving loss that guides the NeRF representation to fuse multiple hypothesized depth maps from each view and distill from them a common geometry that is consistent with all views. Experiments show that our approach enables higher fidelity novel view synthesis from sparse views. Our project page can be found at scade-spacecarving-nerfs.github.io.

1. Introduction

Neural radiance fields (NeRF) [24] have enabled high fidelity novel view synthesis from dozens of input views. Such number of views are difficult to obtain in practice, however. When given only a small number of sparse views, vanilla NeRF tends to struggle with reconstructing shape accurately, due to inadequate constraints enforced by the volume rendering loss alone.

Shape priors can help remedy this problem. Various forms of shape priors have been proposed for NeRFs, such as object category-level priors [13], and handcrafted data-independent priors [25]. The former requires knowledge of category labels and is not applicable to scenes, and the latter is agnostic to the specifics of the scene and only encodes low-level regularity (e.g., local smoothness). A form of prior that is both scene-dependent and category-agnostic is per-view monocular depth estimates, which have been explored in prior work [6, 32]. Unfortunately, monocular depth estimates are often inaccurate, due to estimation errors and inherent ambiguities, such as albedo vs. shading (cf. check shadow illusion), concavity vs. convexity (cf. hollow face illusion), distance vs. scale (cf. miniature cinematography), etc. As a result, the incorrect or inconsistent priors imposed by such depth estimates may mislead the NeRF into reconstructing incorrect shape and produce artifacts in the generated views.

In this paper, we propose a method that embraces the uncertainty and ambiguities present in monocular depth esti-

mates, by modelling a probability distribution over depth estimates. The ambiguities are retained at the stage of monocular depth estimation, and are only resolved once information from multiple views are fused together. We do so with a principled loss defined on probability distributions over depth estimates for different views. This loss selects the subset of modes of probability distributions that are consistent across all views and matches them with the modes of the depth distribution along different rays as modelled by the NeRF. It turns out that this operation can be interpreted as a probabilistic analogue of classical depth-based space carving [17]. For this reason, we dub our method *Space Carving with Ambiguity-aware Depth Estimates*, or SCADE for short.

Compared to prior approaches of depth supervision [6,32] that only supervise the moments of NeRF’s depth distribution (rather than the modes), our key innovation is that we supervise the modes of NeRF’s depth distribution. The supervisory signal provided by the former is weaker than the latter, because the former only constrains the value of an integral aggregated along the ray, whereas the latter constrains the values at different individual points along the ray. Hence, the supervisory signal provided by the former is 2D (because it integrates over a ray), whereas the supervisory signal provided by the our method is 3D (because it is point-wise) and thus can be more fine-grained.

An important technical challenge is modelling probability distributions over depth estimates. Classical approaches use simple distributions with closed-form probability densities such as Gaussian or Laplace distributions. Unfortunately these distributions are not very expressive, since they only have a single mode (known as “unimodal”) and have a fixed shape for the tails. Since each interpretation of an ambiguous image should be a distinct mode, these simple unimodal distributions cannot capture the complex ambiguities in depth estimates. Naïve extensions like a mixture of Gaussians are also not ideal because some images are more ambiguous than others, and so the number of modes needed may differ for the depth estimates of different images. Moreover, learning such a mixture requires backpropagating through E-M, which is nontrivial. Any attempt at modifying the probability density to make it capable of handling a variable number of modes can easily run into an intractable partition function [4,10,31], which makes learning difficult because maximum likelihood requires the ability to evaluate the probability density, which is a function of the partition function.

To get around this conundrum, we propose representing the probability distribution with a set of samples generated from a neural network. Such a distribution can be learned with a conditional GAN; however, because GANs suffer from mode collapse, they cannot model multiple modes [11] and are therefore unsuited to modelling ambiguity. Instead, we propose leveraging conditional Implicit Maximum Likelihood Estimation (cIMLE) [19,28] to learn the distribution, which is designed to avoid mode collapse.

We consider the challenging setting of leveraging out-of-

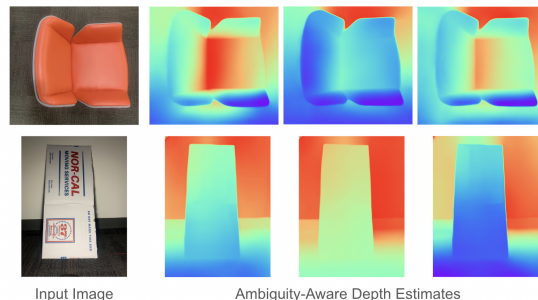


Figure 2. **Ambiguities from a single image.** We show samples from our ambiguity-aware depth estimates that is able to handle different types of ambiguities. (Top) An image of a chair taken from the top-view. Without context of the scene, it is unclear that it is an image of a chair. We show different samples from our ambiguity-aware depth estimates that are able to capture different degrees of convexity. (Bottom) An image of a cardboard under bad lighting conditions that capture the albedo vs shading ambiguity that is also represented by our different samples.

domain depth priors to train NeRFs on real-world indoor scenes with sparse views. Under this setting, the depth priors we use are trained on a different dataset (e.g., Taskonomy) from the scenes our NeRFs are trained on (e.g., ScanNet). This setting is more challenging than usual due to domain gap between the dataset the prior is trained on and the scenes NeRF is asked to reconstruct. We demonstrate that our method outperforms vanilla NeRF and NeRFs with supervision from depth-based priors in novel view synthesis.

In summary, our key contributions include:

- An method that allows the creation of NeRFs in unconstrained indoor settings from only a modest number of 2D views by introducing a sophisticated way to exploit ambiguity-aware monocular depth estimation.
- A novel way to sample distributions over image depth estimates based on conditional Implicit Maximum Likelihood Estimation that can represent depth ambiguities and capture a variable number of discrete depth modes.
- A new space-carving loss that can be used in the NeRF formulation to optimize for a mode-seeking 3D density that helps select consistent depth modes across the views and thus compensate for the under-constrained photometric information in the few view regime.

2. Related Work

Novel View Synthesis and 3D Reconstruction. Reconstructing the structure of a scene from a few views is a long standing problem in computer vision with a long literature. Neural Radiance Fields or NeRF [24] revolutionized the field by showing how to use weights of an MLP to represent a scene that is rendered using volume rendering [23]. A key innovation was the use of positional encoding [37,38] to increase the effective capacity of the MLPs that model a emitted radiance and density as a function of position and

viewing direction. Extensions of NeRF include works on unconstrained photo collections [22], dynamic scenes [20], deformable scenes [27], and reflective materials [1, 39].

A critical shortcoming of NeRF is its reliance on having many input views of a scene. Several approaches have been proposed, including adding patch likelihood losses [25], data-driven priors [36, 48], semantic consistency prior [12], image features [41], or surface [25, 39], occupancy [26], and depth [15, 32, 43] priors. DS-NeRF [6] uses the sparse point reconstructions recovered during Structure-from-Motion, to supervise the depth of sparse points in the recovered NeRF. In the same spirit, DDP [32] uses a depth completion network, that takes as input sparse point cloud projected to one an input view, and produces a depth estimate. Both works [15, 32] model depth with uncertainty but only supervise with *moments* of NeRF’s depth distribution. In contrast, our work is able to represent *multimodal* depth estimates, which can handle the inherent ambiguities of depth estimation, and is able to seek the modes for each depth estimate to make a consistent prediction of the scene structure.

Inspired by traditional multi-view stereo, MVS-NeRF [3] and RC-MVSNet [2] incorporate the use cost volumes [9] to NeRF. To construct the cost volumes, these works look for agreement between features to look for correspondences, which is difficult under the setting of large variations in appearance or viewpoint. In contrast, our approach introduces a novel space carving loss that does not rely on feature correspondences and instead directly defines the loss in 3D.

Another line of works focus on geometry reconstruction [40, 45, 49] by using the volumetric rendering scheme to learn a neural SDF representation. These works tackle a different problem on modeling geometry reconstruction, unlike NeRFs whose focus is modeling appearance for novel view synthesis, where our work falls under.

Depth Estimation from Single View. Depth estimation from a single image is a complex task due multiple ambiguities, including scene scale and shading / albedo ambiguities. Early efforts used MRFs to compute predictions on superpixels [33]. The advent of consumer depth cameras like Kinect [53] enabled the acquisition of larger scale indoor 3D datasets [35], leading to methods that used deep neural networks to predict depth from a single color image [7]. Nonetheless, scaling datasets for depth estimation remained a challenge. Deep3D [44] enabled the creation of stereo views by training on a large dataset of stereo movies, while [8] used the self-supervision of left-right consistency to learn depth from stereo views from driving cars. This approach was extended to videos where the egomotion is also estimated and self-supervision happens across time [54]. MegaDepth [21] uses the depth from SfM reconstructions of internet photos, together with semantic segmentation that conveys ordinal depth supervision cues, i.e. transient objects must be in front of the static scenes. Most recent, multi-task learning has shown promise in training depth estimators that work well in out-of-domain data [30].

In some cases, surface normal estimation has fewer pitfalls than direct depth estimation, as it suffers somewhat less from scale ambiguities. GeoNet [29] jointly estimates depth and surface normals, then uses the geometric relation between them to refine the estimates. LeRes [46] uses a second geometry reasoning module that refines a focal length estimate and the depth itself. Our depth estimation approach is derived from LeReS without the second stage, as the focal length is already estimated during SfM for NeRF scenes.

Most relevant is the work of Kendall and Gal [15], that learns depth estimation in a Bayesian setting, where their objective maximizes the likelihood under a predicted distribution. In our work, we use instead a multimodal depth estimation technique built on cIMLE [19], which makes our prior robust to ambiguities in depth prediction.

3. Background

3.1. Neural Radiance Fields (NeRF)

A neural radiance field [24], or NeRF for short, represents a field in 3D space, where each point represents an infinitesimal particle with a certain opacity that emits varying amounts of light along different viewing directions. The opacity at a point is represented as a volume density, and the amount of emitted light is represented as a colour. Mathematically, a NeRF is represented as two parameterized functions, a volume density $\sigma_\theta : \mathbb{R}^3 \rightarrow \mathbb{R}_{\geq 0}$ and a colour $c_\psi : \mathbb{R}^3 \times S^2 \rightarrow [0, 255]^3$. The former maps a 3D coordinate $\mathbf{x} \in \mathbb{R}^3$ to a volume density $\sigma \in \mathbb{R}_{\geq 0}$, and the latter maps a 3D coordinate $\mathbf{x} \in \mathbb{R}^3$ and a viewing direction $\mathbf{d} \in S^2$ to a colour $\mathbf{c} \in [0, 255]^3$.

To render a NeRF from a view, we shoot rays from each pixel on the camera sensor and integrate over the product of colour, volume density and visibility along the ray to arrive at each pixel value. Visibility at a point is represented with transmittance, which accumulates the exponentiated negative volume density multiplicatively up to the point. The higher transmittance at a point, the more visible the point is from the camera. Mathematically, if the camera ray is $\mathbf{r}(t) = \mathbf{o} + t\mathbf{d}$, where \mathbf{o} is the camera centre and \mathbf{d} is the ray direction, the pixel value $\hat{I}_{\theta, \psi}(\mathbf{o}, \mathbf{d})$ we would have is:

$$\hat{I}_{\theta, \psi}(\mathbf{o}, \mathbf{d}) := \int_{t_n}^{t_f} T_{\theta, \mathbf{o}, \mathbf{d}}(t) \sigma_\theta(\mathbf{o} + t\mathbf{d}) c_\psi(\mathbf{o} + t\mathbf{d}, \mathbf{d}) dt$$

where $T_{\theta, \mathbf{o}, \mathbf{d}}(t) = \exp\left(-\int_{t_n}^t \sigma_\theta(\mathbf{o} + s\mathbf{d}) ds\right)$.

(1)

In the expressions above, t_n and t_f denote the points the ray intersects with the near plane and far plane respectively.

3.2. Inverse Rendering with NeRF

Given a set of real-world images from known views, inverse rendering aims to find the scene whose rendered images match the real-world images. If we use $I(\mathbf{o}, \mathbf{d})$ to de-

note the pixel value for the ray $\mathbf{r}(t) = \mathbf{o} + t\mathbf{d}$, this problem can be cast as an optimization problem, namely:

$$\min_{\theta, \psi} \sum_{\mathbf{o}} \sum_{\mathbf{d}} \|\hat{I}_{\theta, \psi}(\mathbf{o}, \mathbf{d}) - I(\mathbf{o}, \mathbf{d})\|_2^2 \quad (2)$$

If $\hat{I}_{\theta, \psi}(\mathbf{o}, \mathbf{d})$ is differentiable w.r.t. the parameters θ, ψ , this problem can be tackled straightforwardly with gradient-based optimization. To enable this, the volume density function σ_{θ} and the colour function c_{ψ} are chosen to be neural networks. As a result, inverse rendering amounts to training the NeRF to reconstruct the scene in a way that is consistent with the images that are given. Once the NeRF is trained, novel view synthesis can be achieved by rendering the NeRF from new, unseen views.

Since this optimization problem is underconstrained, in general many views are needed to reconstruct the scene accurately. Yet, in practical applications, typically only a few views are available, hence priors are needed to provide sufficient constraints. In this paper, we consider priors in the form of monocular depth estimates from each view. The monocular depth estimates are trained on different datasets from the dataset NeRF is trained on to simulate real-world conditions where there is a domain gap between what the prior was trained on and what the NeRF is trained on.

3.3. Ray Termination Distance

In order to leverage depth priors, a natural way is to use them to constrain the ray termination distance. Because NeRF can represent non-opaque surfaces, there is no single ray termination distance. Instead, there is a distribution of ray termination distances. The cumulative distribution function (CDF) of this distribution represents the probability that the ray terminates before reaching a given point. It turns out that the probability density of this distribution $f_{\theta, \mathbf{o}, \mathbf{d}}(t)$ is given by [6]:

$$f_{\theta, \mathbf{o}, \mathbf{d}}(t) = T_{\theta, \mathbf{o}, \mathbf{d}}(t) \sigma_{\theta}(\mathbf{o} + t\mathbf{d}) \quad (3)$$

4. Method

4.1. Supervising Ray Termination Distance

If all surfaces were opaque and the ground truth depth were given, supervising NeRF with the ground truth depth would be straightforward. All that is needed is to make the NeRFs distribution of ray termination distances as close as possible to a delta distribution centred at the ground truth depth. However, when the ground truth depth is not known, we have to use monocular depth estimates as supervision. The challenge is that unlike ground truth depth, monocular depth estimates may not be stereo-consistent, due to estimation error and inherent ambiguity. As a result, it may not be possible to find a scene that is consistent with the depth estimates from every view. Therefore, in order to reconstruct the scene, we must model the uncertainty and ambiguity in the depth estimates, which can be mathematically characterized by distributions. Such ambiguities cannot be resolved from

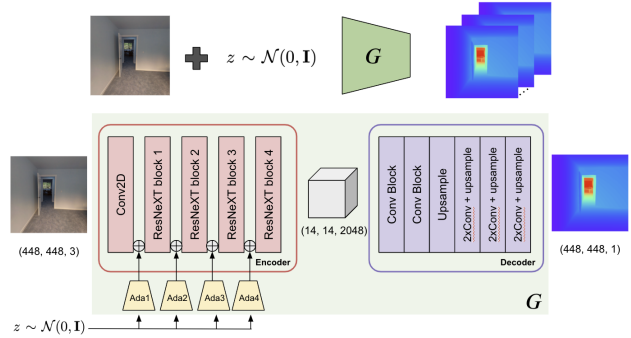


Figure 3. Network Architecture for our Ambiguity-aware Depth Estimates.

a single view and can often be resolved from multiple views. So we need a method that can fuse together uncertain depth estimates from different views.

A natural way of representing uncertainty is through a probability distribution. When there are ambiguities, the distribution of depth estimates are typically *multimodal*, i.e., their probability densities have disjoint peaks, where each mode represents one interpretation of the scene structure. For example, albedo vs shading ambiguity as shown in Figure 2 can result in multiple plausible interpretations and make the distribution multimodal.

An added challenge arises when non-opaque surfaces are present. In this case, there is no single ground truth depth, and so the distribution of ground truth depths is multimodal. So even after training the NeRF, the distribution of ray termination distances induced by NeRF should be multimodal, as illustrated in the glass walls in Figure 1.

Prior methods compute moments (i.e., mean and/or variance) of the distribution of depth estimates or ray termination distance induced by NeRF. For example, DS-NeRF [6] predicts the mean depth estimate and uses reprojection error as a proxy for the variance. It then minimizes the KL divergence from the distribution of ray termination distances induced by NeRF to a Gaussian whose moments match those of the depth estimates. DDP [32] fits a Gaussian to distribution of ray termination distances and maximizes the likelihood of the depth estimate under the Gaussian.

Distilling complex distributions to moments has the drawback of ignoring their possible multimodality. Because both the distribution of depth estimates and ray termination distances are potentially multimodal, it is important to handle the multimodality.

4.2. SCADE

We now present our novel method *Space Carving with Ambiguity-aware Depth Estimates (SCADE)*. Our contribution addresses the existing problems and drawbacks mentioned in the previous section. First, SCADE accounts for the *multimodality* of both the distributions of monocular depth estimates and ray termination distances that arise due to inherent ambiguities (Sec 4.2.1) and non-opaque surfaces.

Second, SCADE is also able to resolve these ambiguities by fusing together information from multiple views through our space carving loss (Sec 4.2.3) that uses the reverse cross entropy loss. Because this a loss on the distribution rather than moments of the distribution (which was the paradigm in prior work [6, 32]), it achieves supervision in 3D rather than 2D. Third, our loss formulation is sample-based (Sec 4.2.2) and so is computationally efficient to optimize.

4.2.1 Our Ambiguity-aware Depth Estimates

We handle multimodality of monocular depth distributions by introducing our ambiguity-aware depth estimation module to account for ambiguities in depth estimation from a single view (Figure 1, 2). In contrast to existing monocular depth estimation networks [47] that predict single point estimates for depth, we model the inherent uncertainty in monocular depth estimation by representing the distribution as a *set of samples* generated by a neural network. We propose leveraging conditional Implicit Maximum Likelihood estimation (cIMLE) [19] to learn a multimodal distribution of depth estimates. We chose cIMLE rather than conditional GANs [14] in order to avoid mode collapse, which can lead to a unimodal distribution.

Figure 3 shows our network architecture. Concretely, we combine cIMLE with a state-of-the-art monocular depth estimation network, LeReS [47]. Our ambiguity-aware depth estimation module (G) takes an input image (I) and a latent code $z \sim \mathcal{N}(0, \mathbf{I})$ and outputs a conditional depth sample $G(I, z)$ for the input image I . To inject randomness into the network, we follow the technique used in [14] and incorporate AdaIn layers into the network backbone that predicts a scale and shift to the intermediate network features. Specifically, we added four AdaIn layers to the encoder of the depth estimation backbone.

4.2.2 Sample-based Losses on Distributions

We desire to achieve an agreement between the distributions from our ambiguity-aware prior and the ray termination distance from NeRF to constrain the NeRF optimization under the sparse view set-up using our learned prior. As distributions are continuous, taking their analytical integral is computationally intractable. Thus, in order to define a differentiable loss function for our optimization problem, we result to a sample-based loss on these two distributions.

Samples from depth prior. As previously described, we leverage cIMLE to sample from our ambiguity-aware depth prior. Sampling $z_1, \dots, z_M \sim \mathcal{N}(0, \mathbf{I})$, we get depth map samples $G(I, z_1), \dots, G(I, z_M)$ for input image I . Hence we denote corresponding samples to *estimate* and represent the distribution for ray \mathbf{r} as $y_1, y_2, \dots, y_M \sim G_{\mathbf{o}, \mathbf{d}}$.

Samples from NeRF. For the distribution given by the ray termination distances modelled by NeRF, we sample from its probability density function $f_{\theta, \mathbf{o}, \mathbf{d}}$ (Eq. 3) with inverse transform sampling. Concretely, we define the probability mass function for ray \mathbf{r} as $f_{\theta, \mathbf{o}, \mathbf{d}}(t_i)$ as given in Eq. 3 with



Figure 4. Our space carving loss supervises the ray termination in 3D as opposed to existing approaches that supervise on 2D expected depth. As shown, this clears out clouds of dust in space. Moreover supervising in 3D allows for wide baseline input views and still finds the common mode that correctly snaps the surfaces.

samples t_1, \dots, t_K , where $t_i \in [t_n, t_f] \forall i$. By inverse transform sampling, we first compute the cumulative distribution function (CDF) of the ray termination distribution at samples t_i , which we denote as $F_{\theta, \mathbf{o}, \mathbf{d}}(t_i) = \sum_{j < i} f_{\theta, \mathbf{o}, \mathbf{d}}(t_j)$. Thus our samples x_1, x_2, \dots, x_N of ray termination distances from NeRF are then given by:

$$x_i = F_{\theta, \mathbf{o}, \mathbf{d}}^{-1}(u_i), \quad (4)$$

where $u_i \sim \mathcal{U}(0, 1)$

4.2.3 Our Space Carving Loss

We now introduce our novel space carving loss utilizes the samples from both multimodal distributions to constrain the NeRF optimization. The goal is to find a subset of modes captured in the monocular depth distributions from each view that are globally consistent. In doing so, we can fuse together the information from the different views, since the inherent ambiguities are only resolved given information from multiple views. This would allow us to find a common shape that is consistent across all views and “snap” objects surfaces together.

We thus desire a loss that has **mode seeking** behavior, that is the modes of the distribution being trained should be a subset of the modes of the distribution that provides supervision. One such loss is the reverse cross entropy, which is the cross entropy from the NeRF ray termination distance distribution to the distribution of our ambiguity-aware depth estimates:

$$H(f_{\theta, \mathbf{o}, \mathbf{d}}, G_{\mathbf{o}, \mathbf{d}}) = -\mathbb{E}_{f_{\theta, \mathbf{o}, \mathbf{d}}}[\log G_{\mathbf{o}, \mathbf{d}}]. \quad (5)$$

As shown in the IMLE papers [19, 28], this is equivalent to penalizing the minimum of the L2 norm between the samples from the two distributions, please see [18] for the full proof. Hence, our space carving loss is given by

$$\mathcal{L}_{\text{space_carving}}(\mathbf{r}) = \sum_{i \in [N]} \min_{j \in [M]} \|x_i - y_j\|_2^2. \quad (6)$$

Because our novel space carving loss operates on distributions rather than moments of distributions, we can have

different supervision targets for different points along the ray. In contrast, if the loss were to only operate on moments, which are integrals along the ray, this only yields the same supervision target for all points along the ray. Therefore, our loss allows us to supervise in 3D, unlike prior methods [6, 32] (which use losses on moments) that supervise in 2D. 3D supervision allows us to clear out dust in space, since in order to move a sample from the ray termination distance distribution farther in depth, the loss would have to decrease the probability density of *all* the points in front of it. These advantages are highlighted in Figure 4.

Our total loss to optimize and train our NeRF model is given by

$$\mathcal{L} = \mathcal{L}_{\text{photometric}} + \lambda \mathcal{L}_{\text{space_carving}}, \quad (7)$$

where $\mathcal{L}_{\text{photometric}}$ is the standard MSE loss on the predicted and ground truth image. See Fig. 1 for our overall pipeline.

5. Results

In this section, we present our experimental evaluation to demonstrate the advantages of **SCADE**.

5.1. Datasets and Evaluation Metrics

We evaluate our method on ScanNet [5] and an in-the-wild dataset we collected¹. We use the sparse-view ScanNet data used by DDP [32] in their evaluations, which comprises of three sample scenes each with 18 to 20 train images and 8 test images. To further test the robustness of our method, we further evaluated our method on three in-the-wild scenes collected using an iPhoneX. We captured sparse views in three different scenes – basement, kitchen and lounge, where each scene has 18 to 23 train images and 8 test images. Following similar data preprocessing steps as DDP [32], we ran SfM [34] on all images to obtain camera poses for NeRF training. For quantitative comparison, we follow the original NeRF [24] paper and report the PSNR, SSIM [42] and LPIPS [52] on the novel test views.

5.2. Implementation Details

We train our ambiguity-aware prior on the Taskonomy [50] dataset and initialize the weights with the pre-trained LeReS [47] model. We use $M = 20$ depth estimates in our experiments. See supplement for additional details.

5.3. Experiments on ScanNet

We evaluate on ScanNet following the evaluation from DDP [32]. We compare to the original NeRF [24] (**Vanilla NeRF**) and the recent state-of-the-art NeRF with depth prior-based supervision, Dense Depth Prior [32] (**DDP**). Table 1 shows SCADE quantitatively outperforming the baselines. Because we are interested in simulating real-world conditions with a domain gap between the prior and the NeRF, our setting requires the use of out-of-domain priors. Since DDP uses an in-domain prior, we retrain DDP’s depth completion

¹Also see supplement for results on the Tanks and Temples dataset [16].

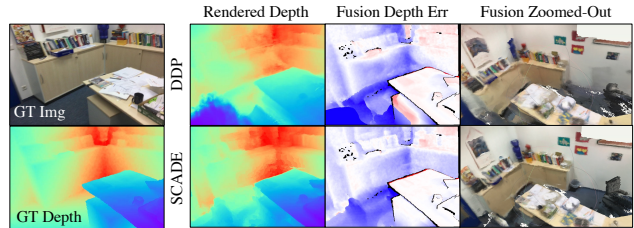


Figure 5. **Depth and Fusion Comparison.** Depth map and fusion comparison between SCADE and DDP.

	PSNR \uparrow	SSIM \uparrow	LPIPS \downarrow
Vanilla NeRF [24]	19.03	0.670	0.398
NerfingMVS [47]	16.29	0.626	0.502
IBRNet [41]	13.25	0.529	0.673
MVSNerf [3]	15.67	0.533	0.635
DS-NeRF [6]	20.85	0.713	0.344
DDP [32]	19.29	0.695	0.368
SCADE (Ours)	21.54	0.732	0.292

Table 1. **ScanNet Results.** Results for DS-NeRF and NerfingMVS follow what was reported in prior literature [32]. Because our setting requires out-of-domain priors, the results for DDP are with out-of-domain priors. The results of DDP with in-domain priors are (20.96, 0.737, 0.236) for PSNR, SSIM and LPIPS, respectively.

network using their official code and hyperparameters on Taskonomy [50], *i.e.* the same out-of-domain dataset that our prior is trained on.

Figure 6 shows our qualitative results. As shown, compared to the baselines, SCADE is able to avoid producing the clouds of dust that are present in the results of baselines (first, third and last column). Moreover, SCADE is also able to snap and recover objects in the scene such as the details on the blinds (second column), the back of the chair (fourth column) and the legs of the piano stool (fifth column). Moreover, Fig. 5 shows rendered depthmaps and fusion results using [51] on a Scannet scene. Notice that SCADE is able to recover better geometry compared to DDP – see corner of the calendar, cabinets and office chair in the right image.

5.4. Experiments on In-the-Wild Data

To further test the robustness of SCADE, we further evaluate on data collected in-the-wild with a standard phone camera (iPhoneX). The intrinsics for different views are different and are not known to the algorithm. As shown in Table 2, **SCADE** with the same out-of-domain prior (trained on Taskonomy) outperforms the baselines on in-the-wild captured data, which demonstrates its robustness. Figure 7 shows qualitative examples. Interestingly, we are able to recover the room behind the glass better the baselines, whose results are more blurry (first, third and fourth column). We are also able to recover objects behind the glass such as the monitor and chair (first column). As discussed in Sect. 4.1, non-opaque surfaces are the most challenging because both the ground truth distribution of ray termination distances induced by NeRF and the distribution of estimated depths

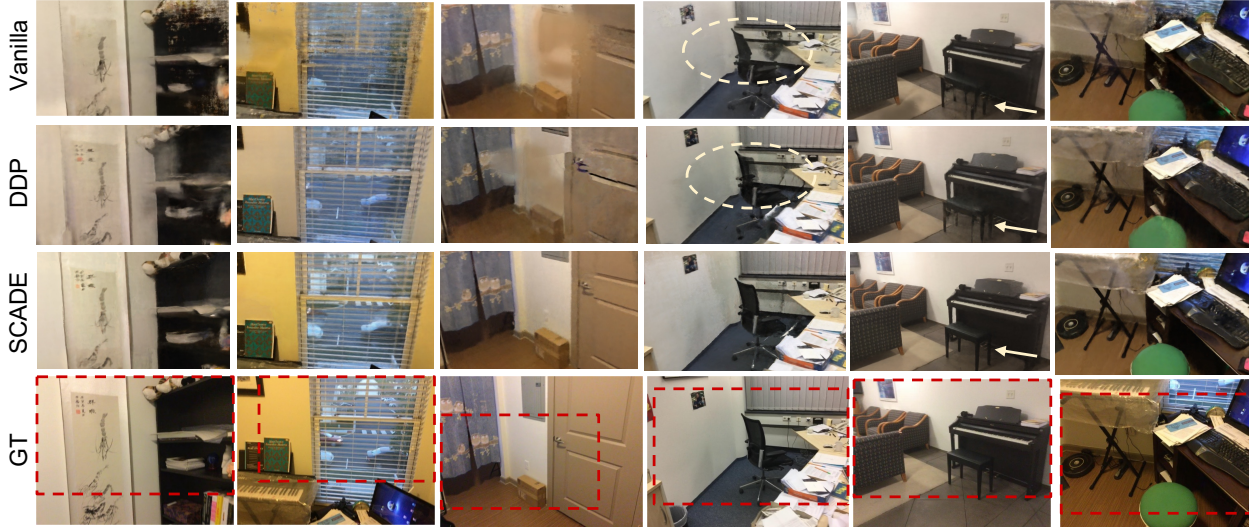


Figure 6. **Qualitative results on ScanNet.** As shown, SCADE is able to avoid producing the clouds of dust that are present in the results of baselines (1st, 3rd, last columns). Moreover, SCADE is also able to snap and recover objects in the scene such as the details on the blinds (2nd column), the back of the chair (4th column) and the legs of the piano stool (5th column).

	PSNR \uparrow	SSIM \uparrow	LPIPS \downarrow
Vanilla NeRF [24]	20.46	0.713	0.398
DDP [32]	21.28	0.727	0.366
SCADE (Ours)	22.82	0.743	0.347

Table 2. **In-the-wild Results.**

from the prior must be multimodal. This validates the ability of our model to capture the multimodal distributions.

Similar to the ScanNet results, we are able to recover crisp shapes of objects such as the thin table leg and the white chair (second and last column), and we are also able to clear up dust near the microwave and printer compared to the baselines (fifth and last column).

5.5. Ablation Study

Supervision on full distribution vs moments We first validate the importance of using our novel space carving loss that supervises the full distribution of ray termination distance rather than just its moments. The latter integrates along the ray, and so provides one target value for the entire ray, which effectively makes it 2D supervision. The former provides different target values for different samples along the same ray, which makes it 3D supervision. We adapt our method to use 2D supervision proposed in the recent MonoSDF [49], which computes the expected ray termination distance and aligns the output depth map with a monocular depth estimation prior using the MiDaS [30] loss. We use our learned prior as the monocular depth estimate. Table 3 validates the effectiveness of our space carving loss.

Multimodal prior vs unimodal prior. Our prior models does not constrain the depth distribution to a particular form and allows it to be multimodal. In contrast, DDP’s depth completion prior assumes a Gaussian, which is unimodal.

	PSNR \uparrow	SSIM \uparrow	LPIPS \downarrow
MonoSDF supervision	20.13	0.710	0.332
DDP prior - single sample	20.85	0.712	0.320
DDP prior - multiple samples	21.00	0.718	0.316
Our prior - single sample	21.22	0.714	0.318
SCADE (Ours)	21.54	0.732	0.292

Table 3. **Ablation Study.** Results on the Scannet dataset. MonoSDF supervision refers to supervising with MonoSDF loss on expected ray termination distance using our prior.

We ablate our prior against DDP’s by sampling from their Gaussian distribution. We train with our space carving loss using both a single sample, as well as M samples drawn from DDP’s prior. As shown in Table 3, while both perform better than the 2D supervision provided by MonoSDF, but they are subpar compared to using our multimodal prior.

Multiple vs single sample. Finally, we also ablate on using a single sample vs multiple samples for our space carving loss. For the single sample set-up, we supervise using the sample mean of the depth estimates, which is equivalent to the maximum likelihood estimate of the depth under Gaussian likelihood. Results in Table 3 show the importance of using multiple samples for our space carving loss.

Sparsity. We further show add’l results under varying number of views in Tab 4. Note that in general sparsity is not fully reflected by the absolute number of views, because all else being equal, a larger scene requires more views to attain complete coverage. Following prior work [32], we also report the average number of views that see the same point as another measure of sparsity. As shown, our setting is at the edge of the theoretical lower limit of 2 for general-purpose reconstruction, which shows view sparsity.




Figure 7. **Qualitative results on In-the-Wild data.** As shown, we are able to recover the room behind the glass better than the baselines, whose results are more blurry (1st, 3rd, 4th column). Also, SCADE is able to recover crisp shapes of objects such as the thin table leg and the white chair (2nd, last column), and we are also able to clear up dust near the microwave and printer compared to the baselines (5th, last column).

ave. #views visible	1.87	1.98	2.01	2.2
absolute #views	18	20	22	24
DDP / SCADE	19.35/ 21.66	22.4/ 23.67	23.10/ 24.00	23.56/ 24.84

Table 4. PSNR for DDP [32]/SCADE on scene781 of Scannet.




Figure 8. **Depth Mode Discussion.** a) Train images from Taskonomy [50] and their labels. Notice that non-opaque surfaces are labelled differently. b) Output of our prior on reflective surfaces.

5.6. Discussion on our Ambiguity-Aware Prior

We provide additional details on our ambiguity-aware prior on i) how and why it works, ii) its performance on reflective surfaces, which is a common failure case for depth estimate, and iii) provide an intuition on how what our depth modes look like. i) We achieve variable depth modes by exploiting inconsistently labeled training data. As shown in Fig 8-a, in Taskonomy [50], different training images with non-opaque surfaces label depth differently: shooting through the glass (left), on the glass (middle), or a mixture of both (right). Despite multiple possible depth labels, each image only has one ground truth label. Training with cIMLE allows our prior to model these multiple possible (ambiguous) outputs through sampling², even when given only one label per image. Interestingly when testing our prior on a test image with a mirror, we find that it is able to capture variable modes

²See supplement for more qualitative examples on our ambiguity-aware depth estimates.

on reflective surfaces, including the “correct” flat surface as shown in Fig 8-b. An intuition of what the depth modes look like is for example if a ray intersects $n - 1$ non-opaque surfaces, the n modes are intersections with the non-opaque surfaces and the terminating point on the opaque surface.

6. Conclusion

In this paper, we present a new approach towards NeRF reconstruction that can work with a modest number of in-the-wild views of an indoor scene. We address the under-constrained nature of this problem by regularizing the NeRF optimization with additional depth estimates for each view. Our key technical contribution is to model multimodality in the depth estimates, which can capture inherent ambiguities in monocular depth estimation as well as the possible presence of non-opaque surfaces. We resolve ambiguities using a novel space carving loss that fuses the multimodal depth estimates from different views and seeks the modes that are consistent across views so as to arrive at a globally consistent 3D reconstruction. The improved recovery of shape and appearance enables higher fidelity novel view synthesis from sparse views.

Limitations and Future Work. The performance of our method is constrained by the quality of the monocular depth priors. While we found that our prior generalizes well across domains, if the domain gap is too great, the performance of our method will degrade. A future direction would be to detect when this happens and dynamically adjust the strength of depth supervision in response.

Acknowledgements. We sincerely thank the Teleportation and CCI team in Google for all the insightful discussions during the summer. We also thank Mirko Visontai for internship logistics, Guandao Yang for experiment set-up assist, and Weicheng Kuo for looking over the paper.

References

- [1] Sai Bi, Zexiang Xu, Pratul Srinivasan, Ben Mildenhall, Kalyan Sunkavalli, Miloš Hašan, Yannick Hold-Geoffroy, David Kriegman, and Ravi Ramamoorthi. Neural reflectance fields for appearance acquisition. *arXiv preprint arXiv:2008.03824*, 2020. **3**
- [2] Di Chang, Aljaž Božič, Tong Zhang, Qingsong Yan, Yingcong Chen, Sabine Süsstrunk, and Matthias Nießner. Rcmvsnet: Unsupervised multi-view stereo with neural rendering. *ECCV 2022*, 2022. **3**
- [3] Anpei Chen, Zexiang Xu, Fuqiang Zhao, Xiaoshuai Zhang, Fanbo Xiang, Jingyi Yu, and Hao Su. Mvsnerf: Fast generalizable radiance field reconstruction from multi-view stereo. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 14124–14133, 2021. **3, 6**
- [4] Kyung Hyun Cho, Tapani Raiko, and Alexander Ilin. Gaussian-bernoulli deep boltzmann machine. In *The 2013 International Joint Conference on Neural Networks (IJCNN)*, pages 1–7. IEEE, 2013. **2**
- [5] Angela Dai, Angel X Chang, Manolis Savva, Maciej Halber, Thomas Funkhouser, and Matthias Nießner. Scannet: Richly-annotated 3d reconstructions of indoor scenes. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 5828–5839, 2017. **6**
- [6] Kangle Deng, Andrew Liu, Jun-Yan Zhu, and Deva Ramanan. Depth-supervised nerf: Fewer views and faster training for free. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 12882–12891, 2022. **1, 2, 3, 4, 5, 6**
- [7] David Eigen, Christian Puhrsch, and Rob Fergus. Depth map prediction from a single image using a multi-scale deep network. *Advances in neural information processing systems*, 27, 2014. **3**
- [8] Clément Godard, Oisín Mac Aodha, and Gabriel J Brostow. Unsupervised monocular depth estimation with left-right consistency. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 270–279, 2017. **3**
- [9] Xiaodong Gu, Zhiwen Fan, Siyu Zhu, Zuozhuo Dai, Feitong Tan, and Ping Tan. Cascade cost volume for high-resolution multi-view stereo and stereo matching. 2019. **3**
- [10] Geoffrey E Hinton. A practical guide to training restricted boltzmann machines. In *Neural networks: Tricks of the trade*, pages 599–619. Springer, 2012. **2**
- [11] Phillip Isola, Jun-Yan Zhu, Tinghui Zhou, and Alexei A Efros. Image-to-image translation with conditional adversarial networks. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 1125–1134, 2017. **2**
- [12] Ajay Jain, Matthew Tancik, and Pieter Abbeel. Putting nerf on a diet: Semantically consistent few-shot view synthesis. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 5885–5894, October 2021. **3**
- [13] Wonbong Jang and Lourdes Agapito. Codenerf: Disentangled neural radiance fields for object categories. In *ICCV*, 2021. **1**
- [14] Tero Karras, Samuli Laine, and Timo Aila. A style-based generator architecture for generative adversarial networks. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 4401–4410, 2019. **5**
- [15] Alex Kendall and Yarin Gal. What uncertainties do we need in bayesian deep learning for computer vision? *Advances in neural information processing systems*, 30, 2017. **3**
- [16] Arno Knapitsch, Jaesik Park, Qian-Yi Zhou, and Vladlen Koltun. Tanks and temples: Benchmarking large-scale scene reconstruction. *ACM Transactions on Graphics*, 36(4), 2017. **6**
- [17] K.N. Kutulakos and S.M. Seitz. A theory of shape by space carving. In *ICCV*, 1999. **2**
- [18] Ke Li and Jitendra Malik. Implicit maximum likelihood estimation. *arXiv preprint arXiv:1809.09087*, 2018. **5**
- [19] Ke Li*, Shichong Peng*, Tianhao Zhang*, and Jitendra Malik. Multimodal image synthesis with conditional implicit maximum likelihood estimation. *International Journal of Computer Vision*, May 2020. **2, 3, 5**
- [20] Zhengqi Li, Simon Niklaus, Noah Snavely, and Oliver Wang. Neural scene flow fields for space-time view synthesis of dynamic scenes. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 6498–6508, 2021. **3**
- [21] Zhengqi Li and Noah Snavely. Megadepth: Learning single-view depth prediction from internet photos. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 2041–2050, 2018. **3**
- [22] 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 *CVPR*, 2021. **3**
- [23] Nelson Max. Optical models for direct volume rendering. *IEEE Transactions on Visualization and Computer Graphics*, 1(2):99–108, 1995. **2**
- [24] 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 *Computer Vision—ECCV 2020: 16th European Conference, Glasgow, UK, August 23–28, 2020, Proceedings, Part I*, pages 405–421, 2020. **1, 2, 3, 6, 7**
- [25] Michael Niemeyer, Jonathan T. Barron, Ben Mildenhall, Mehdi S. M. Sajjadi, Andreas Geiger, and Noha Radwan. Regnerf: Regularizing neural radiance fields for view synthesis from sparse inputs. In *CVPR*, 2022. **1, 3**
- [26] Michael Oechsle, Songyou Peng, and Andreas Geiger. Unisurf: Unifying neural implicit surfaces and radiance fields for multi-view reconstruction. In *International Conference on Computer Vision (ICCV)*, 2021. **3**
- [27] Keunhong Park, Utkarsh Sinha, Jonathan T. Barron, Sofien Bouaziz, Dan B Goldman, Steven M. Seitz, and Ricardo Martin-Brualla. Nerfies: Deformable neural radiance fields. *ICCV*, 2021. **3**
- [28] Shichong Peng, Seyed Alireza Moazenipourasil, and Ke Li. Chimle: Conditional hierarchical imle. In *NeurIPS*, 2022. **2, 5**
- [29] Xiaojuan Qi, Renjie Liao, Zhengzhe Liu, Raquel Urtasun, and Jiaya Jia. Geonet: Geometric neural network for joint depth and surface normal estimation. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 283–291, 2018. **3**
- [30] René Ranftl, Katrin Lasinger, David Hafner, Konrad Schindler, and Vladlen Koltun. Towards robust monocular

- depth estimation: Mixing datasets for zero-shot cross-dataset transfer. *IEEE transactions on pattern analysis and machine intelligence*, 2020. 3, 7
- [31] Marc Aurelio Ranzato, Y-Lan Boureau, Sumit Chopra, and Yann LeCun. A unified energy-based framework for unsupervised learning. In *Artificial Intelligence and Statistics*, pages 371–379. PMLR, 2007. 2
- [32] Barbara Roessle, Jonathan T Barron, Ben Mildenhall, Pratul P Srinivasan, and Matthias Nießner. Dense depth priors for neural radiance fields from sparse input views. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 12892–12901, 2022. 1, 2, 3, 4, 5, 6, 7, 8
- [33] Ashutosh Saxena, Min Sun, and Andrew Y Ng. Make3d: Learning 3d scene structure from a single still image. *IEEE transactions on pattern analysis and machine intelligence*, 31(5):824–840, 2008. 3
- [34] Johannes Lutz Schönberger and Jan-Michael Frahm. Structure-from-motion revisited. In *Conference on Computer Vision and Pattern Recognition (CVPR)*, 2016. 6
- [35] Nathan Silberman, Derek Hoiem, Pushmeet Kohli, and Rob Fergus. Indoor segmentation and support inference from rgb-d images. In *European conference on computer vision*, pages 746–760. Springer, 2012. 3
- [36] Matthew Tancik, Ben Mildenhall, Terrance Wang, Divi Schmidt, Pratul P. Srinivasan, Jonathan T. Barron, and Ren Ng. Learned initializations for optimizing coordinate-based neural representations. In *CVPR*, 2021. 3
- [37] Matthew Tancik, Pratul Srinivasan, Ben Mildenhall, Sara Fridovich-Keil, Nithin Raghavan, Utkarsh Singhal, Ravi Ramamoorthi, Jonathan Barron, and Ren Ng. Fourier features let networks learn high frequency functions in low dimensional domains. *Advances in Neural Information Processing Systems*, 33:7537–7547, 2020. 2
- [38] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. Attention is all you need. *Advances in neural information processing systems*, 30, 2017. 2
- [39] Dor Verbin, Peter Hedman, Ben Mildenhall, Todd Zickler, Jonathan T. Barron, and Pratul P. Srinivasan. Ref-NeRF: Structured view-dependent appearance for neural radiance fields. *CVPR*, 2022. 3
- [40] Peng Wang, Lingjie Liu, Yuan Liu, Christian Theobalt, Taku Komura, and Wenping Wang. Neus: Learning neural implicit surfaces by volume rendering for multi-view reconstruction. *NeurIPS*, 2021. 3
- [41] Qianqian Wang, Zhicheng Wang, Kyle Genova, Pratul Srinivasan, Howard Zhou, Jonathan T. Barron, Ricardo Martin-Brualla, Noah Snavely, and Thomas Funkhouser. Ibrnet: Learning multi-view image-based rendering. In *CVPR*, 2021. 3, 6
- [42] Zhou Wang, Alan C Bovik, Hamid R Sheikh, and Eero P Simoncelli. Image quality assessment: from error visibility to structural similarity. *IEEE transactions on image processing*, 13(4):600–612, 2004. 6
- [43] Yi Wei, Shaohui Liu, Yongming Rao, Wang Zhao, Jiwen Lu, and Jie Zhou. Nerfingmvs: Guided optimization of neural radiance fields for indoor multi-view stereo. In *ICCV*, 2021. 3
- [44] Junyuan Xie, Ross Girshick, and Ali Farhadi. Deep3d: Fully automatic 2d-to-3d video conversion with deep convolutional neural networks. In *European conference on computer vision*, pages 842–857. Springer, 2016. 3
- [45] Lior Yariv, Jiatao Gu, Yoni Kasten, and Yaron Lipman. Volume rendering of neural implicit surfaces. In *NeurIPS*, 2021. 3
- [46] Wei Yin, Jianming Zhang, Oliver Wang, Simon Niklaus, Simon Chen, Yifan Liu, and Chunhua Shen. Towards accurate reconstruction of 3d scene shape from a single monocular image. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2022. 3
- [47] Wei Yin, Jianming Zhang, Oliver Wang, Simon Niklaus, Long Mai, Simon Chen, and Chunhua Shen. Learning to recover 3d scene shape from a single image. In *Proc. IEEE Conf. Comp. Vis. Patt. Recogn. (CVPR)*, 2021. 5, 6
- [48] Alex Yu, Vickie Ye, Matthew Tancik, and Angjoo Kanazawa. pixelNeRF: Neural radiance fields from one or few images. In *CVPR*, 2021. 3
- [49] Zehao Yu, Songyou Peng, Michael Niemeyer, Torsten Sattler, and Andreas Geiger. Monosdf: Exploring monocular geometric cues for neural implicit surface reconstruction. *arXiv preprint arXiv:2206.00665*, 2022. 3, 7
- [50] Amir R Zamir, Alexander Sax, William Shen, Leonidas J Guibas, Jitendra Malik, and Silvio Savarese. Taskonomy: Disentangling task transfer learning. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 3712–3722, 2018. 6, 8
- [51] Andy Zeng, Shuran Song, Matthias Nießner, Matthew Fisher, Jianxiong Xiao, and Thomas Funkhouser. 3dmatch: Learning local geometric descriptors from rgb-d reconstructions. In *CVPR*, 2017. 6
- [52] Richard Zhang, Phillip Isola, Alexei A Efros, Eli Shechtman, and Oliver Wang. The unreasonable effectiveness of deep features as a perceptual metric. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 586–595, 2018. 6
- [53] Zhengyou Zhang. Microsoft kinect sensor and its effect. *IEEE multimedia*, 19(2):4–10, 2012. 3
- [54] Tinghui Zhou, Matthew Brown, Noah Snavely, and David G Lowe. Unsupervised learning of depth and ego-motion from video. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 1851–1858, 2017. 3