

FWD: Real-time Novel View Synthesis with Forward Warping and Depth

Ang Cao, Chris Rockwell, Justin Johnson
University of Michigan, Ann Arbor
{ancao, cnris, justincj}@umich.edu

Abstract

Novel view synthesis (NVS) is a challenging task requiring systems to generate photorealistic images of scenes from new viewpoints, where both quality and speed are important for applications. Previous image-based rendering (IBR) methods are fast, but have poor quality when input views are sparse. Recent Neural Radiance Fields (NeRF) and generalizable variants give impressive results but are not real-time. In our paper, we propose a generalizable NVS method with sparse inputs, called FWD, which gives high-quality synthesis in real-time. With explicit depth and differentiable rendering, it achieves competitive results to the SOTA methods with 130-1000 \times speedup and better perceptual quality. If available, we can seamlessly integrate sensor depth during either training or inference to improve image quality while retaining real-time speed. With the growing prevalence of depths sensors, we hope that methods making use of depth will become increasingly useful.

1. Introduction

Given several posed images, *novel view synthesis (NVS)* aims to generate photorealistic images depicting the scene from unseen viewpoints. This long-standing task has applications in graphics, VR/AR, bringing life to still images. It requires a deep visual understanding of geometry and semantics, making it appealing to test visual understanding.

Early work on NVS focused on *image-based rendering (IBR)*, where models generate target views from a set of input images. Light field [39] or proxy geometry (like mesh surfaces) [12, 24, 61, 62] are typically constructed from posed inputs, and target views are synthesized by resampling or blending warped inputs. Requiring dense input images, these methods are limited by 3D reconstruction quality, and can perform poorly with *sparse* input images.

Recently, Neural Radiance Fields (NeRF) [48] have become the leading methods for NVS, using MLPs to represent the 5D radiance field of the scene *implicitly*. The color and density of each sampling point are queried from the network and aggregated by volumetric rendering to get the




Figure 1. **Real-time Novel View Synthesis.** We present a real-time and generalizable method to synthesize images from sparse inputs. NeRF variants model the scene via an MLP, which is queried millions of times during rendering and leads to low speeds. Our method utilizes explicit depths and point cloud renderers for fast rendering, inspired by SynSin [82]. The model is trained end-to-end with a novel fusion transformer to give high-quality results, where regressed depths and features are optimized for synthesis.

pixel color. With dense sampling points and *differentiable renderer*, explicit geometry isn’t needed, and densities optimized for synthesis quality are learned. Despite impressive results, they are not *generalizable*, requiring MLP fitting for each scene with dense inputs. Also, they are extremely *slow* because of tremendous MLP query times for a single image.

Generalizable NeRF variants like PixelNeRF [89], IBR-Net [78] and MVNeRF [9] emerge very recently, synthesizing novel views of unseen scenes without per-scene optimization by modeling an MLP conditioned on sparse inputs. However, they still query the MLP millions of times, leading to slow speeds. Albeit the progress of accelerating NeRF with per-scene optimization [88, 27, 18], fast and generalizable NeRF variants are still under-explored.

In this paper, we target a *generalizable NVS* method with *sparse* inputs, refraining dense view collections. Both *real-time speed* and *high-quality* synthesis are expected, allowing interactive applications. Classical IBR methods are fast but require dense input views for good results. Generalizable NeRF variants show excellent quality without per-scene optimization but require intense computations, leading to slow speeds. Our method, termed FWD, achieves this

target by Forward Warping features based on Depths.

Our key insight is that explicitly representing the *depth* of each input pixel allows us to apply *forward warping* to each input view using a differentiable point cloud renderer. This avoids the expensive volumetric sampling used in NeRF-like methods, enabling real-time speed while maintaining high image quality. This idea is deeply inspired by the success of SynSin [82], which employs a differentiable point cloud renderer for single image NVS. Our paper extends SynSin to multiple inputs settings and explores effective and efficient methods to fuse multi-view information.

Like prior NVS methods, our approach can be trained with RGB data only, but it can be progressively enhanced if noisy sensor depth data is available during training or inference. Depth sensors are becoming more prevalent in consumer devices such as the iPhone 13 Pro and the LG G8 ThinQ, making RGB-D data more accessible than ever. For this reason, we believe that methods making use of RGB-D will become increasingly useful over time.

Our method estimates depths for each input view to build a point cloud of latent features, then synthesizes novel views via a point cloud renderer. To alleviate the inconsistencies between observations from various viewpoints, we introduce a view-dependent feature MLP into point clouds to model view-dependent effects. We also propose a novel Transformer-based fusion module to effectively combine features from multiple inputs. A refinement module is employed to inpaint missing regions and further improve synthesis quality. The whole model is trained end-to-end to minimize photometric and perceptual losses, learning depth and features optimized for synthesis quality.

Our design possesses several advantages compared with existing methods. First, it gives both high-quality and high-speed synthesis. Using explicit point clouds enables real-time rendering. In the meanwhile, differentiable renderer and end-to-end training empower high-quality synthesis results. Also, compared to NeRF-like methods, which cannot synthesize whole images during training because of intensive computations, our method could easily utilize perceptual loss and refinement module, which noticeably improves the visual quality of synthesis. Moreover, our model can seamlessly integrate sensor depths to further improve synthesis quality. Experimental results support these analyses.

We evaluate our method on the ShapeNet and DTU datasets, comparing it with representative NeRF-variants and IBR methods. It outperforms existing methods, considering speed and quality jointly: compared to IBR methods we improve both speed and quality; compared to recent NeRF-based methods we achieve competitive quality at real-time speeds (130-1000× speedup). A user study demonstrates that our method gives the most perceptually pleasing results among all methods. The code is available at https://github.com/Caoang327/fwd_code.

2. Related Work

Novel view synthesis is a long-standing problem in computer vision, allowing for the generation of novel views given several scene images. A variety of 3D representations (both implicit and explicit) have been used for NVS, including depth and multi-plane images [74, 94, 72, 57, 7, 67], voxels [69, 21], meshes [61, 23, 28, 62], point clouds [82, 40, 64] and neural scene representations [65, 41, 19, 34, 47, 55, 48]. In this work, we use point clouds as our 3D representations for computational and memory efficiency.

Image-based Rendering. IBR synthesizes novel views from a set of reference images by weighted blending [15, 39, 20, 24, 57, 61, 12, 62]. They generally estimate proxy geometry from dense captured images for synthesis. For instance, Riegler et al. [61] uses multi-view stereo [66, 87, 77, 77, 45, 29] to produce scene mesh surface and warps source view images to target views based on proxy geometry. Despite promising results in some cases, they are essentially limited by the quality of 3D reconstructions, where dense inputs (tens to hundreds) with large overlap and reasonable baselines are necessary for decent results. These methods estimate geometry as an intermediate task not directly optimized for image quality. In contrast we input sparse views and learn depth jointly to optimize for synthesis quality.

Neural Scene Representations. Recent work uses implicit scene representations for view synthesis [65, 41, 19, 34, 47, 55]. Given many views, neural radiance fields (NeRF) show impressive results [48, 92, 46, 56, 81], but require expensive per-scene optimization. Recent methods [78, 89, 75, 9, 31] generalize NeRF without per-scene optimization by learning a shared prior, with sparse inputs. However these methods require expensive ray sampling and therefore are very slow. In contrast, we achieve significant speedups using explicit representations. Some concurrent work accelerates NeRF by reformulating the computation [18], using pre-computation [88, 27], or adding view dependence to explicit 3D representations [41, 83, 2, 8, 49]; unlike ours, these all require dense input views and per-scene optimization.

Utilizing RGB-D in NVS. The growing availability of annotated depth maps [13, 5, 10, 1, 71, 68] facilitates depth utilization in NVS [54, 40, 26], which serves as extra supervision or input to networks. Our method utilizes explicit depths as 3D representations, allowing using sensor depths as additional inputs for better quality. Given the increasing popularity of depth sensors, integrating sensor depths is a promising direction for real-world applications.

Depth has been used in neural scene representations for speedups [51, 73], sparser inputs [16] and dynamic scenes [84]. However, these works still require per-scene optimization. Utilizing RGB-D inputs to accelerate generalizable NeRF like [89, 78] is still an open problem.

Differentiable Rendering and Refinement. We use advances in differentiable rendering [42, 35, 11, 52, 43] to

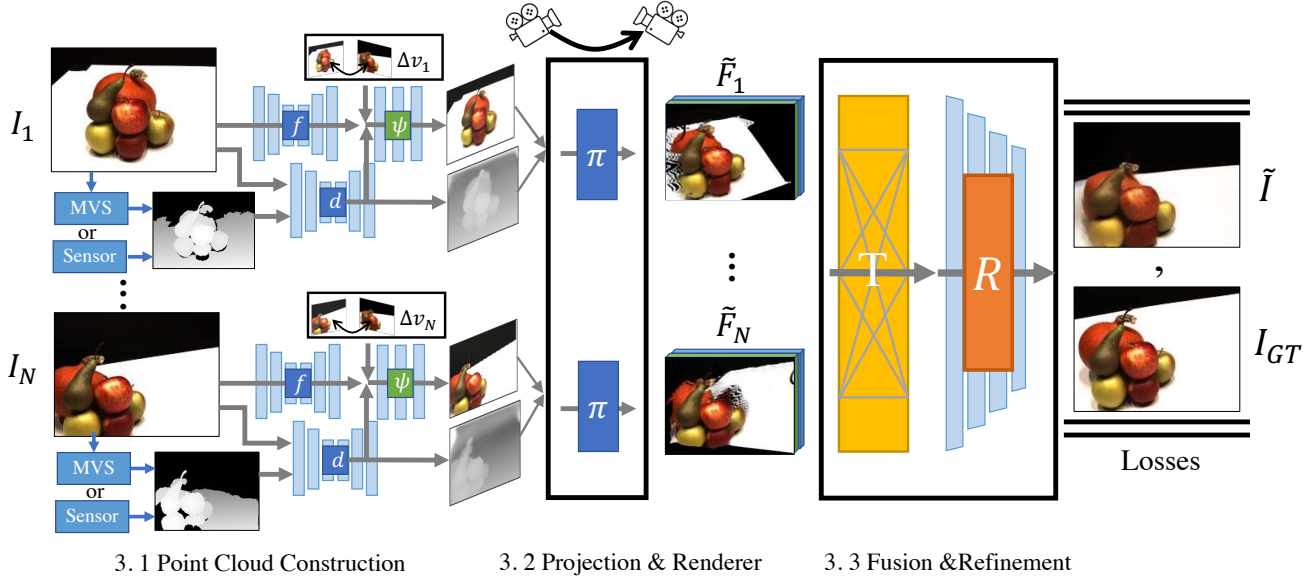


Figure 2. **System Overview.** Given a sparse set of images, we construct a point cloud \mathcal{P}_i for each image I_i using Feature Network f , View-Dependent Feature MLP ψ , and Depth Network d . Besides images, d takes MVS estimated depths or sensor depths as inputs and regresses refined depths. Per-pixel features F'_i are regressed by f and ψ based on images and relative view changes. A differentiable point cloud renderer π is employed to project and render point clouds to target views. We use Transformer T to fuse rendered results from arbitrary number inputs and apply refinement module R for final results. The model is trained with photometric loss and content loss.

learn 3D end-to-end. Learned geometries rely heavily on rendering and refinement [90, 86, 3, 79] to quickly synthesize realistic results. Refinement has improved dramatically owing to generative modeling [38, 36, 91, 95] and rendering frameworks [60, 32, 50, 30]. Instead of aggregating information across viewpoints before rendering [44], we render viewpoints separately and fuse using a Transformer [76, 17, 4], enabling attention across input views.

3. Method

Given a sparse set of input images $\{I_i\}_{i=1}^N$ and corresponding camera poses $\{R_i, T_i\}$, our goal is to synthesize a novel view with camera pose $\{R_t, T_t\}$ fast and effectively. The depths $\{D_i^{sen}\}$ of I_i captured from sensors are optionally available, which are generally incomplete and noisy.

The insight of our method is that using explicit depths and forward warping enables real-time rendering speed and tremendous accelerations. Meanwhile, to alleviate quality degradations caused by inaccurate depth estimations, a differentiable renderer and well-designed fusion & refinement modules are employed, encouraging the model to learn geometry and features optimized for synthesis quality.

As illustrated in Figure 2, with estimated depths, input view I_i is converted to a 3D point cloud \mathcal{P}_i containing geometries and view-dependent semantics of the view. A differentiable neural point cloud renderer π is used to project point clouds to target viewpoints. Rather than directly aggregating point clouds across views before rendering, we propose a Transformer-based module T fusing rendered re-

sults at target view. Finally, a refinement module R is employed to generate final outputs. The whole model is trained end-to-end with photometric and perceptual loss.

3.1. Point Cloud Construction

We use point clouds to represent scenes due to their efficiency, compact memory usage, and scalability to complex scenes. For input view I_i , point cloud \mathcal{P}_i is constructed by estimating depth D_i and feature vectors F'_i for each pixel in the input image, then projecting the feature vectors into 3D space using known camera intrinsics. The depth D_i is estimated by a *depth network* d ; features F'_i are computed by a *spatial feature encoder* f and *view-dependent MLP* ψ .

Spatial Feature Encoder f . Scene semantics of input view I_i are mapped to per-pixel feature vectors F'_i by spatial feature encoder f . Each feature vector in F'_i is 61-dimensions and is concatenated with RGB channels, which is 64 dimensions in total. f is built on BigGAN architecture [3].

Depth Network d . Estimating depth from a single image has scaling/shifting ambiguity, losing valuable multi-view cues and leading to inconsistent estimations across views. Applying multi-view stereo algorithms (MVS) [66, 87, 77, 85] solely on sparse inputs is challenging because of limited overlap and huge baselines between input views, leading to inaccurate and low-confidence estimations. Therefore, we employ a hybrid design cascading a U-Net after the MVS module. The U-Net takes image I_i and estimated depths from the MVS module as inputs, refining depths with multi-view stereo cues and image cues. PatchmatchNet [77] is

utilized as the MVS module, which is fast and lightweight. **Depth Estimation with sensor depths.** As stated, U-Net receives an initial depth estimation from the MVS module and outputs a refined depth used to build the point cloud. If sensor depth D_i^{sen} is available, it is directly input to the U-Net as the initial depth estimations. In this setting, U-Net servers as completion and refinement module taking D_i^{sen} and I_i as inputs, since D_i^{sen} is usually noisy and incomplete. During training, loss \mathcal{L}_s is employed to encourage the U-Net output to match the sensor depth.

$$\mathcal{L}_s = \|M_i \odot D_i - M_i \odot D_i^{sen}\| \quad (1)$$

where M_i is a binary mask indicating valid sensor depths. **View-Dependent Feature MLP ψ .** The appearance of the point could vary across views because of lighting and view direction, causing inconsistency between multiple views. Therefore, we propose to insert view direction changes into scene semantics to model this view-dependent effects. An MLP ψ is designed to compute view-dependent features F'_i by taking F_i and relative view changes Δv from input to target view as inputs. For each point in the cloud, Δv is calculated based on normalized view directions v_i and v_t , from the point to camera centers of input view i and target view t . The relative view direction change is calculated as:

$$\Delta v = [(v_i - v_t) / \|v_i - v_t\|, v_i \cdot v_t], v_i, v_t \in \mathbb{R}^3. \quad (2)$$

and the view-dependent feature F'_i is:

$$F'_i = \psi(F_i, \delta(\Delta v)) \quad (3)$$

where δ is a two-layer MLP mapping Δv to a 32-dimensions vector and ψ is also a two-layer MLP.

3.2. Point Cloud Renderer

To observe the constructed point cloud \mathcal{P}_i at target views, we employ a neural point cloud renderer π . \mathcal{P}_i is first transformed to target view coordinates based on camera poses and then rendered by π . The rendered feature maps \tilde{F}_i share the same dimension as feature F'_i at each pixel. With explicit geometry transformation, our rendered results are geometrically consistent and correct across views.

We use the *differentiable* renderer design of [82], which splats 3D points to the image plane and gets pixel values by blending point features. The blending weights are computed based on z -buffer depths and distances between pixel and point centers. It is implemented using Pytorch3D [60].

This fully differentiable renderer allows our model to be trained end-to-end, where photometric and perceptual loss gradients can be propagated to points' position and features. In this way, the model learns to estimate depths and features optimized for synthesis quality, leading to superior quality. We show the effectiveness of it in experiments.




Figure 3. **Fusion Transformer.** We use a lightweight transformer T to fuse the features from N input views on each pixel. We use a learnable token to query the fusion results.

3.3. Fusion and Refinement

Unlike SynSin [82] using a single image for NVS, fusing multi-view inputs is required in our method. A naïve fusion transforms each point cloud to target view and aggregates them into a large one for rendering. Despite high efficiency, it is vulnerable to inaccurate depths since points with wrong depths may occlude points from other views, leading to degraded results. Methods like PointNet [58] may be feasible to apply on the aggregated point cloud for refinement, but they are not efficient with significant point numbers.

Instead, we render each point cloud individually at target viewpoints and fuse rendered results by a fusion Transformer T . A refinement module R is used to inpaint missing regions, decode feature maps and improve synthesis quality. **Fusion Transformer T .** Given an arbitrary number of rendered feature maps $\{\tilde{F}_i\}$, fusion should be effective, fast, and permutation invariant. Inspired by the success of Transformers, we propose a pixel-wise Transformer T for fusion, detailed in Figure 3. At each pixel, T inputs rendered feature vectors and queries fused results using a learnable “token”. Applied on features, T utilizes semantics for fusion.

Rendered results may lose geometry cues for fusion when rendered from 3D to 2D. For instance, depths may reveal occlusion relationships across views, and relative view changes from input to target views relate to each input’s importance for fusion. Therefore, we also explored to use geometry features as position encoding while not helpful.

Refinement Module R . Built with 8 ResNet [22] blocks, R decodes fused feature maps \tilde{F} to RGB images \tilde{I} at target views. It inpaints regions invisible for inputs in a semantically and geometrically meaningful manner. Also, it corrects local errors caused by inaccurate depths and improves perceptual quality based on semantics contained by feature maps, leading to coherent and high-quality synthesis.

3.4. Training and Implementation Details

Our model is trained end-to-end with photometric \mathcal{L}_{l_2} and perceptual \mathcal{L}_c losses between generated and ground-truth target images. The whole loss function is:

$$\mathcal{L} = \lambda_{l_2} \mathcal{L}_{l_2} + \lambda_c \mathcal{L}_c \quad (4)$$

Table 1. **Model variants settings.** We predefine three model variants with different settings. FWD utilizes a pre-trained MVS module, in which way it gets access to depths during training.

| Name | Test Depth | Train Depth | Depth Network | MVS Module | Losses |
|-------|------------|-------------|---------------|-------------|---|
| FWD-U | | | MVS + U-Net | Random ini. | $\mathcal{L}_{l_2} + \mathcal{L}_c$ |
| FWD | | ✓ | MVS + U-Net | Pre-trained | $\mathcal{L}_{l_2} + \mathcal{L}_c$ |
| FWD-D | ✓ | ✓ | RGB-D + U-Net | - | $\mathcal{L}_{l_2} + \mathcal{L}_c + \mathcal{L}_s$ |

where $\lambda_{l_2} = 5.0, \lambda_c = 1.0$. The model is trained end-to-end on 4 2080Ti GPUs for 3 days, using Adam [37] with learning rate 10^{-4} and $\beta_1=0.9, \beta_2=0.999$. When sensors depths are available as inputs, \mathcal{L}_s is used with $\lambda_s = 5.0$.

4. Experiments

The goal of our paper is *real-time* and *generalizable* novel view synthesis with *sparse* inputs, which can optionally use sensor depths. To this end, our experiments aim to identify the speed and quality at which our method can synthesize novel images and explore the advantage of explicit depths. We evaluate our methods on ShapeNet [6] and DTU [33] datasets, comparing results with the SOTA methods and alternative approaches. Experiments take place with held-out test scenes and no per-scene optimization. We conduct ablations to validate the effectiveness of designs.

Metrics. We conduct A/B test to measure the visual quality, in which workers select the image most similar to the ground truth from competing methods. Automatic image quality metrics including PSNR, SSIM [80] and LPIPS [93] are also reported, and we find LPIPS best reflects the image quality as perceived by humans. Frames per second (FPS) during rendering is measured on the same platform (single 2080Ti GPU with 4 CPU cores). All evaluations are conducted using the same protocol (same inputs and outputs).

Model Variants. Three models are evaluated with various accessibility to depths for training and test, as defined in Table 1. FWD utilizes a pretrained PatchmatchNet [77] as the MVS module for depth estimations, which is also updated during end-to-end training with photometric and perceptual loss. FWD-U learns depth estimations in an Unsupervised manner, sharing the same model and settings as FWD while PatchmatchNet is randomly initialized without any pretraining. FWD-D takes sensor depths as additional inputs during both training and inference. It doesn't use any MVS module since sensor depths provide abundant geometry cues. For pretraining PatchmatchNet, we train it following typical MVS settings and using the same data splitting as NVS.

4.1. ShapeNet Benchmarks

We first evaluate our model for category-agnostic synthesis task on ShapeNet [6]. Following the setting of [89], we train and evaluate a single model on 13 ShapeNet [6] categories. Each instance contains 24 fixed views of 64×64 resolution. During training, one random view is selected




Figure 4. **Qualitative results of category-agnostic NVS on ShapeNet.** We test the capacity of our model by training it across 13 categories of ShapeNet with single view input, and compare with PixelNeRF [89]. No gt depths are available during inference and training. Our results have better visual quality and details.

Table 2. **Category-agnostic NVS on ShapeNet.** Quantitative results for category-agnostic view-synthesis are presented.

| model | 1-view | | | | 2-view | | | |
|-----------|--------------|--------------|--------------|------------|--------------|--------------|--------------|------------|
| | PSNR | SSIM | LPIPS | FPS | PSNR | SSIM | LPIPS | FPS |
| DVR [53] | 22.70 | 0.860 | 0.130 | 1.5 | - | - | - | - |
| SRN [70] | 23.28 | 0.849 | 0.139 | 24 | - | - | - | - |
| PixelNeRF | 26.80 | 0.910 | 0.108 | 1.2 | 28.88 | 0.936 | 0.076 | 1.1 |
| FWD-U | 26.66 | 0.911 | 0.055 | 364 | 28.43 | 0.931 | 0.043 | 336 |

as input and the rests are served as target views. For testing, we synthesize all the other views from a fixed informative view. The model is finetuned with two random input views for 2-view experiments. We find that U-Net is sufficient for good results on this dataset without the MVS module.

Qualitative comparisons to PixelNeRF are shown in Figure 4, where FWD-U gets noticeably superior results. Our synthesized results are more realistic and closely matching to target views, while PixelNeRF's results tend to be blurry. We observe the same trend in the DTU benchmark and evaluate the visual quality quantitatively there.

We show quantitative results in Table 2, adding SRN [70] and DVR [53] as other baselines. Our method outperforms others significantly for LPIPS, indicating a much better perceptual quality, as corroborated by qualitative results. PixelNeRF has a slightly better PSNR while its results are blurry. Most importantly, FWD-U runs at a speed of over 300 FPS, which is $300\times$ faster than PixelNeRF.

4.2. DTU MVS Benchmarks

We also evaluate our model on DTU MVS dataset [33], which is a real scene dataset consisting of 103 scenes. Each scene contains one or multiple objects placed on a table, while images and incomplete depths are collected by the camera and structured light scanner mounted on an industrial robot arm. Corresponding camera poses are provided.




Figure 5. **View synthesis results from FWD.** We show the view synthesis results with 3 input views on DTU dataset from FWD-D (row. 1), FWD (row. 2) and FWD-U (row. 3). Our methods synthesize high-quality and geometrically correct novel views in real time.

As stated in [89], this dataset is challenging since it consists of complex real scenes without apparent semantic similarities across scenes. Also, images are taken under varying lighting conditions with distinct color inconsistencies between views. Moreover, with only under 100 scenes available for training, it is prone to overfitting in training.

We follow the same training and evaluation pipelines as PixelNeRF [89] for all methods to give a fair comparison. The data consists of 88 training and 15 test scenes, between which there are no shared or highly similar scenes. Images are down-sampled to a resolution of 300×400 . For training, three input views are randomly sampled, with the rest as target views. For inference, we choose three fixed informative input views and synthesize other views of the scene. **Baselines.** We evaluate a set of representatives of generalizable NeRF and IBR methods in two different scenarios: with RGB or RGB-D available as inputs during inference.

PixelNeRF [89], IBRNet [78] and MVSNeRF [9] are the SOTA generalizable NeRF variants, taking RGB as inputs. We use the official PixelNeRF model trained on DTU MVS and carefully retrain IBRNet and MVSNeRF with the same 3-input-view settings. PixelNeRF-DS is also included as reported in [16], which is PixelNeRF supervised with depths. Please note that our settings are very different from evaluations used in original papers of IBRNet and MVSNeRF.

A series of IBR methods are also evaluated. Since COLMAP [66] fails to give reasonable outputs with sparse input images, methods using COLMAP like FVS [61], DeepBlending [25] cannot estimate scene geometry in this setting. For these methods, we use depths captured by sensors as estimated depths, which should give upper-bound performance of these methods. To better cope with missing regions, we add our refinement model to DeepBlending [25] and retrain it on DTU dataset, termed Blending-R.

For fairness, we evaluate all methods using the same protocol, distinct from some of their original settings. Although we try our best to adopt these methods, our reported results may still not perfectly reflect their true capacity.

Qualitative Results. Synthesis results are shown in Fig-

ure 5, where high-quality and geometrically correct novel views are synthesized in real-time (over 35 FPS) under significant viewpoint changes. Our refinement module faithfully inpaints invisible regions; also, synthesized images have good shadows, light reflections, and varying appearances across views, showing the efficacy of view-dependent MLP. With sensor depths, results can be further improved.

We show comparisons to baselines in Figure 6. Our methods provide noticeably better results than baselines across different depth settings. For models without depths in test, IBRNet and PixelNeRF give blurry results in areas of high detail such as the buildings in the top row, while our FWD-U and FWD give more realistic and sharper images. With sensor depths in test, baseline Blending-R produces more cogent outputs, but still struggles to distinguish objects from the background, such as in the middle row, while FWD-D gives faithfully synthesis and clear boundaries.

Quantitative Results. We evaluate synthesis quality quantitatively by user study following a standard A/B paradigm. Workers choose the closest to a ground truth image between competing methods, and are monitored using a qualifier and sentinel examples. All views in the test set (690 in total) are evaluated, and each view is judged by three workers.

In Figure 7, user study results support qualitative observations. Among all baselines with and without test depths, users choose our method as more closely matching ground truth images than others most of the time. FWD-U is selected over PixelNeRF in 65.6% of examples, and 77.8% compared to IBRNet. Also, over 90% workers prefer FWD-D to FWD, showing advantage of using sensor depths.

We show automated view synthesis metrics and speeds in Table 3. Across all depth availability settings, our method is competitive with the SOTA baselines while significantly faster. FWD-D runs in real-time and gives substantially better image quality than others. FWD has competitive metrics to PixelNeRF-DS while $1000\times$ faster. Notably, NeRF variants such as PixelNeRF, IBRNet, MVSNeRF, and PixelNeRF-DS are at least two orders of magnitude slower.

The exception to highly competitive performance is




Figure 6. **Qualitative Comparison.** We compare synthesis results from different methods with 3 input views (one of them shown in figure). Our methods give geometrically consistent and visually appealing results, while other results suffering shaking artifacts at some views. Unlike other methods, FWD-D and Blending+R get access sensor depths as inputs during inference.




Figure 7. **User study on DTU.** We conduct a user study by asking subjects to select the results most similar to the ground truth. The numbers indicate the percentage of preference. Methods are grouped based whether using depths during test. We also report FWD vs. FWD-D showing the advantages of sensor depths.

weaker PSNR and SSIM of our unsupervised FWD-U against PixelNeRF and IBRNet. However, FWD-U has noticeably better perceptual quality with the best LPIPS, and human raters prefer it to other methods in A/B tests. The visual quality in figure 6 also illustrates the disparity between comparisons using PSNR and LPIPS. Meanwhile, FWD-U is above $1000\times$ faster than PixelNeRF and above $100\times$ faster than IBRNet. Depth estimations, rendering and CNN would introduce tiny pixel shiftings, which harm the PSNR of our method. NeRF-like methods are trained to optimize L2 loss for each pixel independently, leading to blur results.

Among all methods without test depths, FWD has the best results. Although it uses a pretrained MVS module, we think this comparison is still reasonable since pretrained depth module is easy to get. Also, training depths can be easily calculated from training images since they are dense.

Baseline comparisons also show that IBR methods are fast, but do not give images that are competitive with our method. Our method outperforms them in both perceptual quality and standard metrics, showing the efficacy of proposed methods. Note that Blending+R doesn’t support variable number inputs and our refinement module improves its results significantly. We also compare FWD-U with SynSin [82] which only receives a single input image, showing the benefits of using multi-view inputs in NVS.

Table 3. **Quantitative comparison on DTU real images.** We compare our method with representatives of generalizable NeRF variants and IBR methods for image quality and rendering speed. Our method achieves significantly better speed-quality tradeoff, indicating the effectiveness and efficiency of our design. † Unlike other methods, SynSin receives only one image as input.

| Test | Train | Model | PSNR \uparrow | SSIM \uparrow | LPIPS \downarrow | FPS \uparrow |
|-------|-------|-------------------|-----------------|-----------------|--------------------|----------------|
| RGB | RGB | PixelNeRF [89] | 19.24 | 0.687 | 0.399 | 0.03 |
| | | IBRNet [78] | 18.86 | 0.695 | 0.387 | 0.27 |
| | | MVSNeRF [9] | 17.13 | 0.611 | 0.444 | 0.34 |
| | | SynSin [82] † | 15.66 | 0.564 | 0.388 | 51.8 |
| | | FWD-U | 17.42 | 0.598 | 0.341 | 35.4 |
| RGB | RGB-D | PixelNeRF-DS [16] | 19.87 | 0.710 | 0.370 | 0.03 |
| | | FWD | 20.15 | 0.721 | 0.259 | 35.4 |
| RGB-D | RGB-D | Blending-R [25] | 16.98 | 0.661 | 0.351 | 41.8 |
| | | FVS [61] | 15.92 | 0.733 | 0.267 | 9.70 |
| | | FWD-D | 21.98 | 0.791 | 0.208 | 43.2 |

4.3. Ablations and Analysis

We evaluate the effectiveness of our designs and study depth in more detail through ablation experiments.

Effects of Fusion Transformer. We design a model without Transformer, which concatenates point clouds across views into a bigger one for later rendering and refinement. Its results in FWD-U settings are shown in Figure 8. The ablated version is vulnerable to inaccurate depths learned in unsupervised manner and synthesizes “ghost objects” since points with bad depths occlude other views’ points.

We repeat the same ablation in FWD-D settings, shown in Table 4, which settings give much better depth estimations with sensor depths. The ablated model has notably worse results for all metrics, indicating that the proposed method is not only powerful to tackle inaccurate depth estimations, but also fuse semantic features effectively.

Effects of View Dependent MLP. For ablation, we remove the view-dependent feature MLP and report its results in Table 4. Removing this module reduces model’s ability to produce view-dependent appearances, leading to worse per-




Figure 8. **Ablation on Fusion Transformer.** We show results for FWD-U with and without Transformed-based fusion.

Table 4. **Ablation Studies.** We show the effectiveness of Transformer Fusion and View-dependent MLP by ablation study on FWD-D. These designs improve synthesis quality noticeably while maintaining real-time rendering speed.

| Model | PSNR | SSIM | LPIPS | FPS |
|---------------------|--------------|--------------|--------------|-------------|
| Full model | 21.98 | 0.791 | 0.208 | 43.2 |
| w/o Transformer | 20.95 | 0.748 | 0.241 | 48.4 |
| w/o View dependence | 21.16 | 0.769 | 0.212 | 44.0 |

formance for all metrics. We show more results in Supp.

Depth Analysis and Ablations. We visualize depths in Figure 9. Estimating depths from sparse inputs is challenging and gives less accurate results because of the huge baselines between inputs. We show estimated depths from PatchmatchNet here, filtered based on the confidence scores. Therefore, refinement is essential in our design to propagate multi-view geometry cues to the whole image. Our end-to-end model learns it by synthesis losses.

We ablate the depth network in Table 5 and report depth error δ_{3cm} , which is the percentage of estimated depths within 3 cm of sensor depths. MVS module is critical (row 2) to give geometrically consistent depths. U-Net further refines depths and improves the synthesis quality (row 3). PatchmatchNet has its own shallow refinement layer, already giving decent refinements. Learning unsupervised MVS and NVS jointly from scratch is challenging (row 4), and training depth network without supervision [14] first may give a good initialization for further jointly training.

5. Conclusion

We propose a real-time and generalizable method for NVS with sparse inputs by using explicit depths. This method inherits the core idea of SynSin while extending it to multi-view input settings, which is more challenging. Our experiments show that estimating depths can give impressive results with a real-time speed, outperforming existing methods. Moreover, the proposed method could utilize sensor depths seamlessly and improve synthesis quality significantly. With the increasing availability of mobile depth sensors, we believe our method has exciting real-world 3D applications. We acknowledge there could be the potential




Figure 9. **Depth visualizations.** We visualize the normalized inverse depths involved in our method. Sensor depths are incomplete because of hardware limitations and MVS estimated depths are inaccurate, where many predictions have low confidence. This demonstrates the necessity of depth completion and refinement.

Table 5. **Depths network ablation and error.** We ablate depth network and compute δ_{3cm} as error, which is the percentage of predicted depths within 3 cm of sensor depths.

| Test | Train | Model | PSNR | SSIM | LPIPS | δ_{3cm} |
|------|-------|------------|-------|-------|-------|----------------|
| RGB | RGB-D | FWD | 20.15 | 0.721 | 0.259 | 79.07 |
| RGB | RGB-D | -w/o MVS | 16.69 | 0.594 | 0.357 | 61.62 |
| RGB | RGB-D | -w/o U-Net | 19.10 | 0.702 | 0.285 | 73.62 |
| RGB | RGB | FWD-U | 17.42 | 0.598 | 0.341 | 54.27 |

for the technology to be used for negative purposes by nefarious actors, like synthesizing fake images for cheating.

There are also challenges and limitations yet to be explored. 1) Although using explicit depths gives tremendous speedups, it potentially inflicts depth reliance on our model. We designed a hybrid depth regressor to improve the quality of depth by combining MVS and single image depth estimations. We also employed an effective fusion and refinement module to reduce the degrades caused by inaccurate depths. Despite these designs, the depth estimator may still work poorly in some challenging settings (like very wide camera baselines), and it would influence the synthesis results. Exploring other depth estimation methods like MiDaS [59] could be an interesting direction for future work.

2) The potential capacity of our method is not fully explored. Like SynSin [82], our model (depth/feature network and refinement module especially) is suitable and beneficial from large-scale training data, while the DTU MVS dataset is not big enough and easy to overfit during training. Evaluating our method on large-scale datasets like Hypersim [63] would potentially reveal more advantages of our model, which dataset is very challenging for NeRF-like methods.

3) Although our method gives more visually appealing results, our PSNR and SSIM are lower than NeRF-like methods. We hypothesize that our refinement module is not perfectly trained to decode RGB colors from feature vectors because of limited training data. Also, tiny misalignments caused during the rendering process may also harm the PSNR, although it is not perceptually visible.

Acknowledgement. Toyota Research Institute provided funds to support this work. We thank Dandan Shan, Hao Ouyang, Jiaxin Xie, Linyi Jin, Shengyi Qian for helpful discussions.

References

- [1] Henrik Aanæs, Rasmus Ramsbøl Jensen, George Vogiatzis, Engin Tola, and Anders Bjorholm Dahl. Large-scale data for multiple-view stereopsis. *IJCV*, 120(2):153–168, 2016.
- [2] Alex Yu and Sara Fridovich-Keil, Matthew Tancik, Qinhong Chen, Benjamin Recht, and Angjoo Kanazawa. Plenoxels: Radiance fields without neural networks, 2021.
- [3] Andrew Brock, Jeff Donahue, and Karen Simonyan. Large scale GAN training for high fidelity natural image synthesis. In *ICLR*, 2019.
- [4] Nicolas Carion, Francisco Massa, Gabriel Synnaeve, Nicolas Usunier, Alexander Kirillov, and Sergey Zagoruyko. End-to-end object detection with transformers. In *ECCV*, pages 213–229. Springer, 2020.
- [5] Angel Chang, Angela Dai, Thomas Funkhouser, Maciej Halber, Matthias Niessner, Manolis Savva, Shuran Song, Andy Zeng, and Yinda Zhang. Matterport3d: Learning from rgb-d data in indoor environments. *International Conference on 3D Vision (3DV)*, 2017.
- [6] Angel X Chang, Thomas Funkhouser, Leonidas Guibas, Pat Hanrahan, Qixing Huang, Zimo Li, Silvio Savarese, Manolis Savva, Shuran Song, Hao Su, et al. Shapenet: An information-rich 3d model repository. *Technical Report arXiv:1512.03012*, 2015.
- [7] Gaurav Chaurasia, Sylvain Duchene, Olga Sorkine-Hornung, and George Drettakis. Depth synthesis and local warps for plausible image-based navigation. *ACM Transactions on Graphics (TOG)*, 32(3):1–12, 2013.
- [8] Anpei Chen, Zexiang Xu, Andreas Geiger, , Jingyi Yu, and Hao Su. Tensorf: Tensorial radiance fields, 2022.
- [9] 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.
- [10] Weifeng Chen, Zhao Fu, Dawei Yang, and Jia Deng. Single-image depth perception in the wild. In *NeurIPS*, volume 29, 2016.
- [11] Wenzheng Chen, Jun Gao, Huan Ling, Edward J Smith, Jaakko Lehtinen, Alec Jacobson, and Sanja Fidler. Learning to predict 3d objects with an interpolation-based differentiable renderer. 2019.
- [12] Inchang Choi, Orazio Gallo, Alejandro Troccoli, Min H Kim, and Jan Kautz. Extreme view synthesis. In *ICCV*, pages 7781–7790, 2019.
- [13] Angela Dai, Angel X. Chang, Manolis Savva, Maciej Halber, Thomas Funkhouser, and Matthias Nießner. Scannet: Richly-annotated 3d reconstructions of indoor scenes. In *Proc. Computer Vision and Pattern Recognition (CVPR), IEEE*, 2017.
- [14] Yuchao Dai, Zhidong Zhu, Zhibo Rao, and Bo Li. Mvs2: Deep unsupervised multi-view stereo with multi-view symmetry. In *2019 International Conference on 3D Vision (3DV)*, pages 1–8, 2019.
- [15] Paul E Debevec, Camillo J Taylor, and Jitendra Malik. Modeling and rendering architecture from photographs: A hybrid geometry-and image-based approach. In *Proceedings of the 23rd annual conference on Computer graphics and interactive techniques*, pages 11–20, 1996.
- [16] Kangle Deng, Andrew Liu, Jun-Yan Zhu, and Deva Ramanan. Depth-supervised nerf: Fewer views and faster training for free. *arXiv preprint arXiv:2107.02791*, 2021.
- [17] Alexey Dosovitskiy, Lucas Beyer, Alexander Kolesnikov, Dirk Weissenborn, Xiaohua Zhai, Thomas Unterthiner, Mostafa Dehghani, Matthias Minderer, Georg Heigold, Sylvain Gelly, et al. An image is worth 16x16 words: Transformers for image recognition at scale. *arXiv preprint arXiv:2010.11929*, 2020.
- [18] Stephan J Garbin, Marek Kowalski, Matthew Johnson, Jamie Shotton, and Julien Valentin. Fastnerf: High-fidelity neural rendering at 200fps. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 14346–14355, 2021.
- [19] Kyle Genova, Forrester Cole, Avneesh Sud, Aaron Sarna, and Thomas Funkhouser. Local deep implicit functions for 3d shape. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 4857–4866, 2020.
- [20] Steven J Gortler, Radek Grzeszczuk, Richard Szeliski, and Michael F Cohen. The lumigraph. In *Proceedings of the 23rd annual conference on Computer graphics and interactive techniques*, pages 43–54, 1996.
- [21] Pengsheng Guo, Miguel Angel Bautista, Alex Colburn, Liang Yang, Daniel Ulbricht, Joshua M Susskind, and Qi Shan. Fast and Explicit Neural View Synthesis. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision*, pages 3791–3800, 2022.
- [22] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *CVPR*, pages 770–778, 2016.
- [23] Peter Hedman and Johannes Kopf. Instant 3d photography. *TOG*, 37(4):1–12, 2018.
- [24] Peter Hedman, Julien Philip, True Price, Jan-Michael Frahm, George Drettakis, and Gabriel Brostow. Deep blending for free-viewpoint image-based rendering. *ACM Transactions on Graphics (TOG)*, 37(6):1–15, 2018.
- [25] Peter Hedman, Julien Philip, True Price, Jan-Michael Frahm, George Drettakis, and Gabriel Brostow. Deep blending for free-viewpoint image-based rendering. 37(6):257:1–257:15, 2018.
- [26] Peter Hedman, Tobias Ritschel, George Drettakis, and Gabriel Brostow. Scalable inside-out image-based rendering. *ToG*, 35(6):1–11, 2016.
- [27] Peter Hedman, Pratul P Srinivasan, Ben Mildenhall, Jonathan T Barron, and Paul Debevec. Baking neural radiance fields for real-time view synthesis. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 5875–5884, 2021.
- [28] Ronghang Hu, Nikhila Ravi, Alexander C Berg, and Deepak Pathak. Worldsheet: Wrapping the world in a 3d sheet for view synthesis from a single image. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 12528–12537, 2021.

- [29] Po-Han Huang, Kevin Matzen, Johannes Kopf, Narendra Ahuja, and Jia-Bin Huang. Deepmvs: Learning multi-view stereopsis. *Conference on Computer Vision and Pattern Recognition (CVPR)*, 2018.
- [30] Rui Huang, Wanyue Zhang, Abhijit Kundu, Caroline Pantofaru, David A Ross, Thomas Funkhouser, and Alireza Fathi. An lstm approach to temporal 3d object detection in lidar point clouds. *arXiv preprint arXiv:2007.12392*, 2020.
- [31] 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*, pages 5885–5894, 2021.
- [32] Krishna Murthy Jatavallabhula, Edward Smith, Jean-Francois Lafleche, Clement Fuji Tsang, Artem Rozantsev, Wenzheng Chen, Tommy Xiang, Rev Lebededian, and Sanja Fidler. Kaolin: A pytorch library for accelerating 3d deep learning research. *arXiv preprint arXiv:1911.05063*, 2019.
- [33] Rasmus Jensen, Anders Dahl, George Vogiatzis, Engil Tola, and Henrik Aanæs. Large scale multi-view stereopsis evaluation. In *2014 IEEE Conference on Computer Vision and Pattern Recognition*, pages 406–413. IEEE, 2014.
- [34] Chiyu Jiang, Avneesh Sud, Ameesh Makadia, Jingwei Huang, Matthias Nießner, Thomas Funkhouser, et al. Local implicit grid representations for 3d scenes. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 6001–6010, 2020.
- [35] Yue Jiang, Dantong Ji, Zhizhong Han, and Matthias Zwicker. Sdfdiff: Differentiable rendering of signed distance fields for 3d shape optimization. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 1251–1261, 2020.
- [36] Tero Karras, Samuli Laine, Miika Aittala, Janne Hellsten, Jaakko Lehtinen, and Timo Aila. Analyzing and improving the image quality of stylegan. In *CVPR*, pages 8110–8119, 2020.
- [37] Diederik P Kingma and Jimmy Ba. Adam: A method for stochastic optimization. 2015.
- [38] C. Ledig, L. Theis, F. Huszár, J. Caballero, A. Cunningham, A. Acosta, A. Aitken, A. Tejani, J. Totz, Z. Wang, and W. Shi. Photo-realistic single image super-resolution using a generative adversarial network. In *CVPR*, pages 105–114, 2017.
- [39] Marc Levoy and Pat Hanrahan. Light field rendering. In *Proceedings of the 23rd annual conference on Computer graphics and interactive techniques*, pages 31–42, 1996.
- [40] Andrew Liu, Richard Tucker, Varun Jampani, Ameesh Makadia, Noah Snavely, and Angjoo Kanazawa. Infinite nature: Perpetual view generation of natural scenes from a single image. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 14458–14467, 2021.
- [41] Lingjie Liu, Jiatao Gu, Kyaw Zaw Lin, Tat-Seng Chua, and Christian Theobalt. Neural sparse voxel fields. *Advances in Neural Information Processing Systems*, 33:15651–15663, 2020.
- [42] Shichen Liu, Weikai Chen, Tianye Li, and Hao Li. Soft rasterizer: Differentiable rendering for unsupervised single-view mesh reconstruction. *ICCV*, 2019.
- [43] 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 *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 2019–2028, 2020.
- [44] Stephen Lombardi, Tomas Simon, Jason Saragih, Gabriel Schwartz, Andreas Lehrmann, and Yaser Sheikh. Neural volumes: Learning dynamic renderable volumes from images. *arXiv preprint arXiv:1906.07751*, 2019.
- [45] Xuan Luo, Jia-Bin Huang, Richard Szeliski, Kevin Matzen, and Johannes Kopf. Consistent video depth estimation. *ACM Transactions on Graphics (TOG)*, 39(4):71–1, 2020.
- [46] 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.
- [47] Lars Mescheder, Michael Oechsle, Michael Niemeyer, Sebastian Nowozin, and Andreas Geiger. Occupancy networks: Learning 3d reconstruction in function space. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 4460–4470, 2019.
- [48] 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.
- [49] Thomas Müller, Alex Evans, Christoph Schied, and Alexander Keller. Instant neural graphics primitives with a multiresolution hash encoding. *arXiv:2201.05989*, Jan. 2022.
- [50] Mahyar Najibi, Guangda Lai, Abhijit Kundu, Zhichao Lu, Vivek Rathod, Thomas Funkhouser, Caroline Pantofaru, David Ross, Larry S Davis, and Alireza Fathi. Dops: learning to detect 3d objects and predict their 3d shapes. In *CVPR*, pages 11913–11922, 2020.
- [51] Thomas Neff, Pascal Stadlbauer, Mathias Parger, Andreas Kurz, Joerg H Mueller, Chakravarty R Alla Chaitanya, Anton Kaplanyan, and Markus Steinberger. DONeRF: Towards real-time rendering of compact neural radiance fields using depth oracle networks. In *Computer Graphics Forum*, volume 40, pages 45–59. Wiley Online Library, 2021.
- [52] Michael Niemeyer, Lars Mescheder, Michael Oechsle, and Andreas Geiger. Differentiable volumetric rendering: Learning implicit 3d representations without 3d supervision. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 3504–3515, 2020.
- [53] Michael Niemeyer, Lars Mescheder, Michael Oechsle, and Andreas Geiger. Differentiable volumetric rendering: Learning implicit 3d representations without 3d supervision. In *Proc. IEEE Conf. on Computer Vision and Pattern Recognition (CVPR)*, 2020.
- [54] David Novotny, Ben Graham, and Jeremy Reizenstein. PerspectiveNet: A scene-consistent image generator for new view synthesis in real indoor environments. *Advances in Neural Information Processing Systems*, 32:7601–7612, 2019.
- [55] Jeong Joon Park, Peter Florence, Julian Straub, Richard Newcombe, and Steven Lovegrove. DeepSDF: Learning continuous signed distance functions for shape representation.

- In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 165–174, 2019.
- [56] Keunhong Park, Utkarsh Sinha, Jonathan T Barron, Sofien Bouaziz, Dan B Goldman, Steven M Seitz, and Ricardo Martin-Brualla. Nerfies: Deformable neural radiance fields. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 5865–5874, 2021.
- [57] Eric Penner and Li Zhang. Soft 3d reconstruction for view synthesis. *ToG*, 36(6):1–11, 2017.
- [58] Charles R Qi, Hao Su, Kaichun Mo, and Leonidas J Guibas. Pointnet: Deep learning on point sets for 3d classification and segmentation. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 652–660, 2017.
- [59] 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 (TPAMI)*, 2020.
- [60] Nikhila Ravi, Jeremy Reizenstein, David Novotny, Taylor Gordon, Wan-Yen Lo, Justin Johnson, and Georgia Gkioxari. Accelerating 3d deep learning with pytorch3d. *arXiv preprint arXiv:2007.08501*, 2020.
- [61] Gernot Riegler and Vladlen Koltun. Free view synthesis. In *ECCV*, pages 623–640. Springer, 2020.
- [62] Gernot Riegler and Vladlen Koltun. Stable view synthesis. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2021.
- [63] Mike Roberts, Jason Ramapuram, Anurag Ranjan, Atulit Kumar, Miguel Angel Bautista, Nathan Paczan, Russ Webb, and Joshua M Susskind. Hypersim: A photorealistic synthetic dataset for holistic indoor scene understanding. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 10912–10922, 2021.
- [64] Chris Rockwell, David F. Fouhey, and Justin Johnson. Pixelsynth: Generating a 3d-consistent experience from a single image. In *ICCV*, 2021.
- [65] Robin Rombach, Patrick Esser, and Björn Ommer. Geometry-free view synthesis: Transformers and no 3d priors. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 14356–14366, 2021.
- [66] Johannes Lutz Schönberger, Enliang Zheng, Marc Pollefeys, and Jan-Michael Frahm. Pixelwise view selection for unstructured multi-view stereo. In *European Conference on Computer Vision (ECCV)*, 2016.
- [67] Meng-Li Shih, Shih-Yang Su, Johannes Kopf, and Jia-Bin Huang. 3d photography using context-aware layered depth inpainting. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 8028–8038, 2020.
- [68] Nathan Silberman, Derek Hoiem, Pushmeet Kohli, and Rob Fergus. Indoor segmentation and support inference from rgb-d images. In *ECCV*, pages 746–760, 2012.
- [69] Vincent Sitzmann, Justus Thies, Felix Heide, Matthias Nießner, Gordon Wetzstein, and Michael Zollhofer. Deepvoxels: Learning persistent 3d feature embeddings. In *CVPR*, pages 2437–2446, 2019.
- [70] Vincent Sitzmann, Michael Zollhofer, and Gordon Wetzstein. Scene representation networks: Continuous 3d-structure-aware neural scene representations. In H. Wallach, H. Larochelle, A. Beygelzimer, F. d'Alché-Buc, E. Fox, and R. Garnett, editors, *Advances in Neural Information Processing Systems*, volume 32. Curran Associates, Inc., 2019.
- [71] Shuran Song, Samuel P Lichtenberg, and Jianxiong Xiao. Sun rgb-d: A rgb-d scene understanding benchmark suite. In *CVPR*, pages 567–576, 2015.
- [72] Pratul P Srinivasan, Richard Tucker, Jonathan T Barron, Ravi Ramamoorthi, Ren Ng, and Noah Snavely. Pushing the boundaries of view extrapolation with multiplane images. In *CVPR*, pages 175–184, 2019.
- [73] Karl Stelzner, Kristian Kersting, and Adam R Kosiorek. Decomposing 3d scenes into objects via unsupervised volume segmentation. *arXiv preprint arXiv:2104.01148*, 2021.
- [74] Maxim Tatarchenko, Alexey Dosovitskiy, and Thomas Brox. Multi-view 3d models from single images with a convolutional network. In *European Conference on Computer Vision*, pages 322–337. Springer, 2016.
- [75] Alex Trevithick and Bo Yang. Grf: Learning a general radiance field for 3d representation and rendering. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 15182–15192, 2021.
- [76] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N. Gomez, Lukasz Kaiser, and Illia Polosukhin. Attention is all you need. 2017.
- [77] Fangjinhua Wang, Silvano Galliani, Christoph Vogel, Pablo Speciale, and Marc Pollefeys. Patchmatchnet: Learned multi-view patchmatch stereo. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 14194–14203, 2021.
- [78] Qianqian Wang, Zhicheng Wang, Kyle Genova, Pratul P Srinivasan, Howard Zhou, Jonathan T Barron, Ricardo Martin-Brualla, Noah Snavely, and Thomas Funkhouser. Ibrnet: Learning multi-view image-based rendering. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 4690–4699, 2021.
- [79] Ting-Chun Wang, Ming-Yu Liu, Jun-Yan Zhu, Andrew Tao, Jan Kautz, and Bryan Catanzaro. High-resolution image synthesis and semantic manipulation with conditional gans. In *CVPR*, pages 8798–8807, 2018.
- [80] 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.
- [81] Zirui Wang, Shangzhe Wu, Weidi Xie, Min Chen, and Victor Adrian Prisacariu. Nerf-: Neural radiance fields without known camera parameters. *arXiv preprint arXiv:2102.07064*, 2021.
- [82] Olivia Wiles, Georgia Gkioxari, Richard Szeliski, and Justin Johnson. Synsin: End-to-end view synthesis from a single image. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 7467–7477, 2020.
- [83] Suttisak Wizadwongsa, Pakkapon Phongthawee, Jiraphon Yenphraphai, and Supasorn Suwajanakorn. Nex: Real-time

- view synthesis with neural basis expansion. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 8534–8543, 2021.
- [84] Wenqi Xian, Jia-Bin Huang, Johannes Kopf, and Changil Kim. Space-time neural irradiance fields for free-viewpoint video. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 9421–9431, 2021.
- [85] Jiaxin Xie, Chenyang Lei, Zhuwen Li, Li Erran Li, and Qifeng Chen. Video depth estimation by fusing flow-to-depth proposals. In *2020 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 10100–10107, 2020.
- [86] Chao Yang, Xin Lu, Zhe Lin, Eli Shechtman, Oliver Wang, and Hao Li. High-resolution image inpainting using multi-scale neural patch synthesis. In *CVPR*, pages 6721–6729, 2017.
- [87] Yao Yao, Zixin Luo, Shiwei Li, Tian Fang, and Long Quan. Mvsnet: Depth inference for unstructured multi-view stereo. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 767–783, 2018.
- [88] Alex Yu, Ruilong Li, Matthew Tancik, Hao Li, Ren Ng, and Angjoo Kanazawa. PlenOctrees for real-time rendering of neural radiance fields. In *arXiv*, 2021.
- [89] Alex Yu, Vickie Ye, Matthew Tancik, and Angjoo Kanazawa. pixelnerf: Neural radiance fields from one or few images. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 4578–4587, 2021.
- [90] Jiahui Yu, Zhe Lin, Jimei Yang, Xiaohui Shen, Xin Lu, and Thomas S Huang. Generative image inpainting with contextual attention. In *CVPR*, pages 5505–5514, 2018.
- [91] Han Zhang, Ian Goodfellow, Dimitris Metaxas, and Augustus Odena. Self-attention generative adversarial networks. In *ICML*, pages 7354–7363. PMLR, 2019.
- [92] Kai Zhang, Gernot Riegler, Noah Snaveley, and Vladlen Koltun. Nerf++: Analyzing and improving neural radiance fields. *arXiv preprint arXiv:2010.07492*, 2020.
- [93] Richard Zhang, Phillip Isola, Alexei A Efros, Eli Shechtman, and Oliver Wang. The unreasonable effectiveness of deep features as a perceptual metric. In *CVPR*, pages 586–595, 2018.
- [94] Tinghui Zhou, Richard Tucker, John Flynn, Graham Fyffe, and Noah Snaveley. Stereo magnification: Learning view synthesis using multiplane images. In *SIGGRAPH*, 2018.
- [95] Jun-Yan Zhu, Taesung Park, Phillip Isola, and Alexei A Efros. Unpaired image-to-image translation using cycle-consistent adversarial networks. In *ICCV*, 2017.