

NeRFVS: Neural Radiance Fields for Free View Synthesis via Geometry Scaffolds

Chen Yang^{1*}, Peihao Li³, Zanwei Zhou¹, Shanxin Yuan², Bingbing Liu²,
 Xiaokang Yang¹, Weichao Qiu², Wei Shen^{1†}

¹ MoE Key Lab of Artificial Intelligence, AI Institute, Shanghai Jiao Tong University

² Huawei Noah's Ark Lab ³Tsinghua University

{ycyangchen, SJTU19zzw, xkyang, wei.shen}@sjtu.edu.cn; lph21@mails.tsinghua.edu.cn;

{shanxin.yuan, liu.bingbing, qiuweichaol}@huawei.com



Figure 1. **Illustration of the free view synthesis task.** Free view synthesis aims at synthesizing photo-realistic images on both interpolation and extrapolation setting. In this paper, we propose a novel method, NeRFVS, based on the neural radiance field (NeRF) [12] to achieve indoor scene free view synthesis. Our NeRFVS significantly reduces the distortions and floaters (as also evidenced by the PSNR numbers), producing high-quality images on interpolation and extrapolation setting.

Abstract

We present NeRFVS, a novel neural radiance fields (NeRF) based method to enable free navigation in a room. NeRF achieves impressive performance in rendering images for novel views similar to the input views while suffering for novel views that are significantly different from the training views. To address this issue, we utilize the holistic priors, including pseudo depth maps and view coverage information, from neural reconstruction to guide the learning of implicit neural representations of 3D indoor scenes. Concretely, an off-the-shelf neural reconstruction method is leveraged to generate a geometry scaffold. Then, two loss functions based on the holistic priors are proposed to improve the learning of NeRF: 1) A robust depth loss that can tolerate the error of the pseudo depth map to guide the ge-

ometry learning of NeRF; 2) A variance loss to regularize the variance of implicit neural representations to reduce the geometry and color ambiguity in the learning procedure. These two loss functions are modulated during NeRF optimization according to the view coverage information to reduce the negative influence brought by the view coverage imbalance. Extensive results demonstrate that our NeRFVS outperforms state-of-the-art view synthesis methods quantitatively and qualitatively on indoor scenes, achieving high-fidelity free navigation results.

1. Introduction

Reconstructing an indoor scene from a collection of images and enabling users to navigate inside it freely is a core component for many downstream applications. It is the most challenging novel-view-synthesis (NVS) task, since it requires high fidelity synthesis from *any view*, including not

* Work done during an internship at Huawei Noah's Ark Lab.

† Corresponding Author.

only views similar to the training views (interpolation), but also views that are significantly different from input views (extrapolation), as shown in Fig. 1. To clarify its difference to other NVS tasks, we term it as free-view-synthesis (FVS). The difficulties of FVS lie in not only the common obstacles in scene reconstruction, including low-texture areas, complex scene geometry, and illumination change, but also view imbalance, *e.g.*, casual photos usually cover the scene unevenly, with hundreds of frames for one table and a few for the floor and wall, as shown in Fig. 3.

Recently, NeRF has emerged as a promising technique for 3D reconstruction and novel view synthesis. Although NeRF can achieve impressive interpolation performance, its extrapolation ability is relatively poor [35], especially for low-texture and few-shot regions. In contrast, some neural reconstruction methods can recover the holistic scene geometry successfully with various priors [9, 22, 25, 34], while the synthesized images from these methods contain plenty of artifacts and are over-smoothed. Inspired by the phenomena, we demonstrate that equipping the NeRF with the scene priors of the geometry captured from neural reconstruction is a potential solution for indoor FVS.

Extending NeRF to enable FVS with geometry from neural reconstruction methods is a non-trivial task with two main challenges. 1) *Depth error*. The reconstructed geometry might contain some failures, including holes, depth shifting, and floaters. The optimization of NeRF relies on the multi-view color consistency, while these failures may conflict with the multi-view color consistency, resulting in artifacts. 2) *Distribution ambiguity*. The depth from NeRF is a weighted sum of sampling distance. Merely supervising the depth expectation leads to arbitrary radiance distribution, especially in low-texture and few-shot regions. This ambiguous distribution leads to floaters and blur among rendered images, as shown in Fig. 5.

In this paper, we propose a novel method which exploits the holistic priors, including pseudo depth maps and view coverage information, outputted from a geometry scaffold to guide NeRF optimization, significantly improving quality on low-texture and few-shot regions. Specifically, to address the *depth error*, we propose a robust depth loss that can tolerate the error from the pseudo depth maps, reducing the negative impact of inaccurate geometry. As for the *distribution ambiguity*, it mainly happens in the low-texture and rarely observed areas, *e.g.* ceilings. We propose a variance loss to regularize the variance of the density and color distribution to decrease the ambiguity of these areas. The weights of these two losses are further adjusted according to the view coverage sufficiency to reduce the negative influence brought by the view imbalance. With the geometry priors and variance regularization, our method can significantly reduce the floaters and distortions among low-texture and few-shot regions, achieving high-fidelity extrapolation

performance.

Experiments on synthetic and real-world datasets demonstrate that our method performs high-fidelity extrapolation by removing the distortions and floaters, significantly outperforming other view synthesis methods. Considering the rendering quality and 3D consistency among interpolation and extrapolation, our NeRFVS achieves new state-of-the-art performance on indoor scene FVS.

In conclusion, our contribution can be summarized as follows:

- A novel approach enabling neural radiance fields to perform free view synthesis on real-world scenes at room scale.
- A robust depth loss to address the inaccuracy of the neural-reconstructed geometry.
- A flexible variance loss with view coverage based adjustment to improve the rendering quality among low-texture and few-shot regions.

2 Related Work

Indoor Scene Rendering A classical type of method is image-based rendering (IBR) [3], which fuses multi-view images to output a target view image. FVS [16] and SVS [17] created a coarse geometric scaffold via Structure from Motion (SfM) [20] and applied an encoder-decoder structure to extract features and map them to target views for synthesis. However, these methods fail when few views are provided. SIBRNet [23] proposed to recover geometry from sparse depth by depth completion. Some methods focus on indoor lighting. Given a 3D mesh obtained by multi-view stereo (MVS) reconstruction, Philip et al. [15] employed an implicit representation of scene materials and illumination to enable relighting and free-viewpoint navigation. PhotoScene [32] inferred spatially-varying procedural materials and scene illumination from source images. DSRNet [31] extracted a global mesh from input images and used a two-layer geometric representation to encode RGB-D and reflection information. Recently, some methods combined volume rendering into their pipeline. IBRNet [26] fed features from source images into an MLP backbone and produced point-wise color and density values for volume rendering. Scalable neural scene [29] allocated tiles on a global mesh proxy. Each tile encoded surface and reflection representations, which were summed via volume rendering to produce the final rendering results. Compared with previous methods, our method achieves indoor scene free navigation without ground truth mesh nor RGB-D sensors.

NeRF with Depth and Regularization

NeRF [12] generates accurate geometry under dense multi-view supervision, but struggles with sparse views and unbounded regions, leading to inaccurate depth values and “floaters” in renderings. Previous methods have proposed

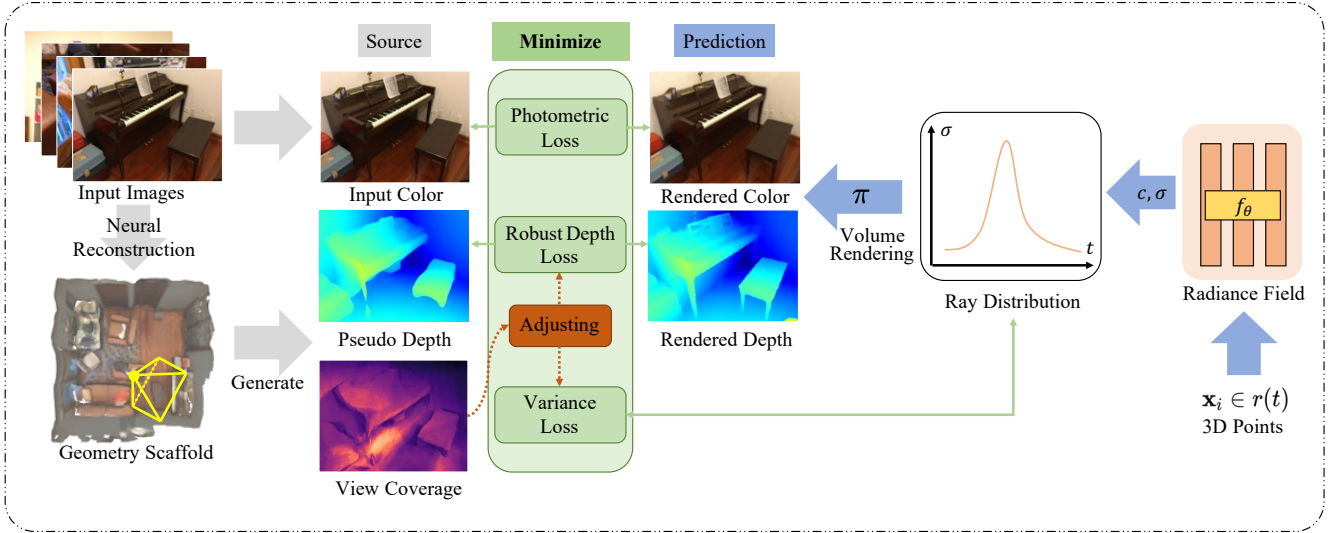


Figure 2. **Overview of NeRFVS.** From left to right, we first generate a geometry scaffold via an off-the-shelf neural reconstruction method, then we get the pseudo depth and view coverage map from the scaffold. From right to left, we optimize a NeRF with regularization from color, depth, and variance distribution along with adjustments based on view coverage.

regularization strategies to address this issue, such as explicit density constraints in InfoNeRF [11] and weight consolidation in Mip-NeRF 360 [2], or depth priors in RegNeRF [14], SinNeRF [30], DS-NeRF [6], and Dense Depth Priors [18]. Our method constrains both depth expectation and density distribution for more consistent visual effects.

3 Preliminaries

We make a quick review of NeRF’s rendering pipeline. A NeRF, optimized by multi-view consistency, represents a 3D scene as the volume density and the directional emitted radiance for any point in space. Given a camera ray parameterized as $\mathbf{r}(t) = \mathbf{o} + t\mathbf{d}$ that passes through the 3D scene, the color and depth are rendered as follows:

$$\hat{\mathbf{C}}(\mathbf{r}) = \int_{t_n}^{t_f} w(\mathbf{r}(t))\mathbf{c}(\mathbf{r}(t), \mathbf{d})dt, \quad (1)$$

$$\hat{D}(\mathbf{r}) = \int_{t_n}^{t_f} w(\mathbf{r}(t))tdt, \quad (2)$$

where $w(\mathbf{r}(t)) = T(t)\sigma(\mathbf{r}(t))$, t_n and t_f represent the near and far bounds, respectively. The weights $w(t)$ are computed by the accumulation of transmittance $T(t) = \exp(-\int_{t_n}^t \sigma(\mathbf{r}(s))ds)$, which stands for the accumulated transmittance from t_n to t . σ and \mathbf{c} represent the volume density and emitted color respectively. We refer to the loss function used by NeRF as the photometric loss $\mathcal{L}_{\text{color}}$:

$$\mathcal{L}_{\text{color}}(\mathbf{r}) = \|\hat{\mathbf{C}}(\mathbf{r}) - \mathbf{C}(\mathbf{r})\|_2^2, \quad (3)$$

where the $\mathbf{C}(\mathbf{r})$ represents the color of the ray \mathbf{r} .

Recently, instant-NGP [13] proposed to represent the whole 3D space with multi-resolution grids stored in a hash

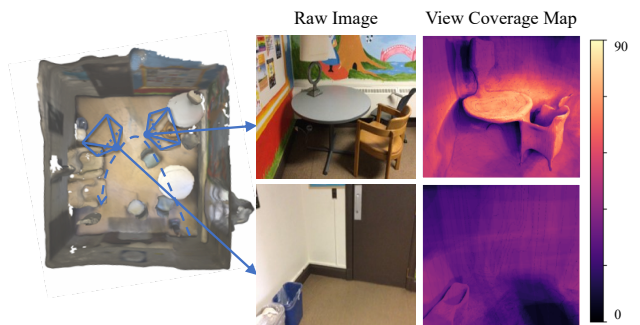


Figure 3. **Imbalanced view coverage for an indoor scene.** The table is captured far more times than the floor and wall.

table. Volume rendering is also applied to sum discrete values in each ray and generate the rendering results. Thanks to this efficient structure, instant-NGP accelerates the training and inference stage of NeRF by a large margin without obvious performance degradation. Our method can be easily integrated into instant-NGP and boost its extrapolation performance on indoor scene rendering as well.

4 Method

Our method facilitates room-scale free view synthesis from a human captured collection of RGB images $\{I_i\}_{i=0}^{N-1}$, along with cameras’ intrinsic parameters $\mathbf{K}_i \in \mathbb{R}^{3 \times 3}$ and poses $\mathbf{p}_i \in \mathbb{R}^6$. We first generate a geometry scaffold from off-the-shelf geometry reconstruction methods [9] using the images and camera parameters (Sec. 4.1). Then the scene priors from the scaffold are integrated into the optimization of a neural radiance field with a robust depth loss to alleviate the influence of scaffold error (Sec. 4.2). To decrease the

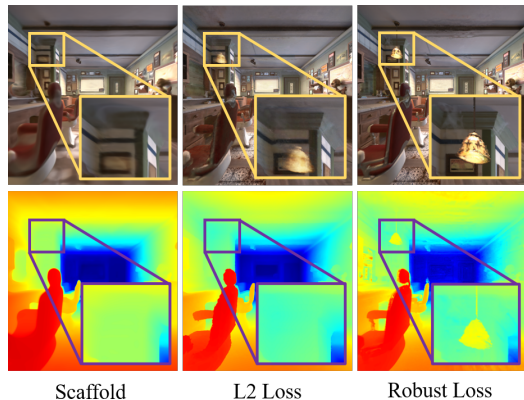


Figure 4. **Effectiveness of our robust depth loss.** Using the proposed robust loss, our method can recover more accurate geometry even when the depth map is inaccurate, *e.g.*, it helps to recover the ceiling lamp from the wrong geometry scaffold.

ambiguity of low-texture areas, we regularize the variance of predicted density distribution and color values along per ray (Sec. 4.3). Finally, we propose a training strategy based on the view coverage of each input view to adjust the intensity of the depth and variance loss accordingly (Sec. 4.4). Fig. 2 shows an overview of our method.

4.1. Geometry Scaffold and View Coverage

Our method applies a 3D geometric scaffold to guiding the optimization of NeRF. To construct this scaffold, we use the geometry from neural geometry reconstruction methods [9, 21, 37] as they can reconstruct a complete and smooth global mesh containing holistic priors. Specifically, we apply the geometry produced by [9] which incorporates Manhattan world assumptions on the structure of the scene into the optimization process. [9] achieves smooth and coherent geometry reconstruction with semantic information, especially in planar regions.

We can render depth maps and compute the view coverage information with the geometry scaffold. We convert the scene priors from the scaffold into the depth priors by rendering distance maps $\{\mathbf{D}_i\}_{i=0}^{N-1}, \mathbf{D}_i \in \mathbb{R}_+^{H \times W \times 1}$. Every pixel’s value in the distance map represents the distance, rather than depth, to the camera. Besides, we generate view coverage maps $\{\mathbf{V}_i\}_{i=0}^{N-1}, \mathbf{V}_i \in [1, N]^{H \times W \times 1}$ by ray casting and shadow mapping. Concretely, we cast rays among test views and get the corresponding positions. Then we re-project these points into training views and use the shadow mapping technique to consider occlusion. In this manner, every pixel’s value of \mathbf{V}_i represents the observed times of this pixel among the training set.

4.2. Robust Depth Loss

In contrast to sparse and confident depth priors, we optimize NeRF with dense while inaccurate depth maps. If

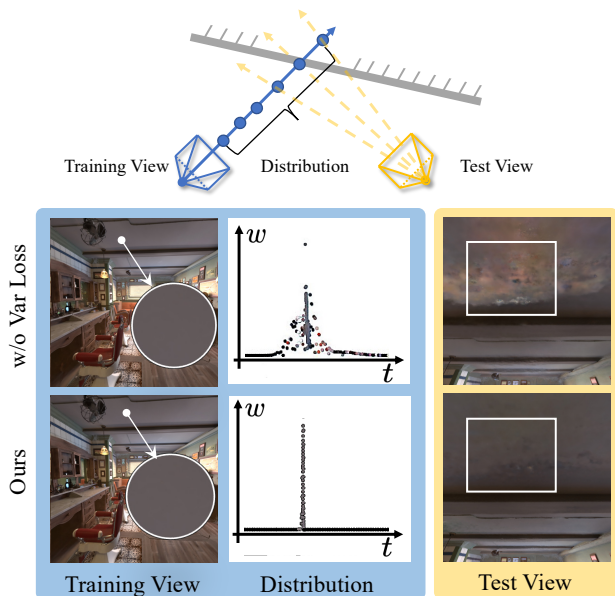


Figure 5. **Variance regularization.** We show that even with the same expected depth and color, the density and color distribution can vary significantly. This arbitrary distribution may not incur errors among interpolation tasks, while causing notable artifacts in extrapolation. Our NeRFVS uses variance regularization to reduce this ambiguity and produces significantly better visual quality.

we directly adopt the optimization strategy from previous methods combining NeRF and geometry, the geometry error would degrade the extrapolation performance, as shown in Fig. 4. Inspired by [8, 10], we propose a robust depth loss to address this issue. Given a ray $\mathbf{r}(t) = \mathbf{o} + t\mathbf{d}$ and its distance value $D(\mathbf{r})$ from the geometry scaffold, our robust loss is defined as:

$$\mathcal{L}_{\text{robust}}(\mathbf{r}) = \begin{cases} 0.5\Delta D(\mathbf{r})^2, & \text{if } \Delta D(\mathbf{r}) < \beta \\ \beta^2(0.5 + \log(\frac{\Delta D(\mathbf{r})}{\beta})), & \text{otherwise} \end{cases} \quad (4)$$

where $\Delta D(\mathbf{r}) = |\hat{D}(\mathbf{r}) - D(\mathbf{r})|$ represents the absolute depth difference, β is a constant set to 0.1 throughout the experiments, and $\hat{D}(\mathbf{r})$ is the depth expectation along the camera axis of $\mathbf{r}(t)$.

The region where the absolute depth difference is larger than β makes the robust loss less sensitive to outliers than L2, with corresponding gradient gradually decreasing to zero. The $\mathcal{L}_{\text{robust}}$ encourages NeRF to terminate rays around the pseudo depth. At the same time, NeRF can retain some freedom for inaccurate geometry since the gradient gets negligible when the absolute depth difference is approaching zero or very large.

4.3. Variance Regularization

To decrease the ambiguity of low-texture and few-shot regions, we introduce the regularization on the variance of weight distribution $w(\mathbf{r})$ and color distribution $\mathbf{c}(\mathbf{r}, \mathbf{d})$. $\hat{D}(\mathbf{r})$ from volume rendering is a weighted sum of sample distance t . Thus, without sufficient constraint signal, even if the $\hat{D}(\mathbf{r})$ is identical, the weight distribution can vary significantly, as shown in Fig. 5. Unfortunately, low texture and few-shot regions are common in indoor scene data. Thus, we propose to further regularize the variance of density and color distribution to inhibit the ambiguity of these regions. Our variance regularization on weights is defined as follows:

$$\mathcal{L}_{\text{var}}^w(\mathbf{r}) = \int_{t_n}^{t_f} w(\mathbf{r}(t)) \left(t - \hat{D}(\mathbf{r}) \right)^2 dt \quad (5)$$

With the same weight distribution $w(\mathbf{r})$ and rendered pixel color $\hat{\mathbf{C}}(\mathbf{r})$, the color of each discrete point can still vary significantly. This issue has little influence on the interpolation performance since the rays emitted at similar angles of the training set result in similar RGB values. However, when the novel view is significantly different from the training views, the derived color value $\mathbf{c}(\mathbf{r}(t), \mathbf{d})$ is messy and will cause unexpected artifacts, as shown in Fig. 5. To address this issue, we design a variance loss on color termed $\mathcal{L}_{\text{var}}^c(\mathbf{r})$, which is defined as follows:

$$\mathcal{L}_{\text{var}}^c(\mathbf{r}) = \int_{t_n}^{t_f} \mathcal{K}(w(\mathbf{r}(t))) \left(\mathbf{c}(\mathbf{r}(t), \mathbf{d}) - \hat{\mathbf{C}}(\mathbf{r}) \right)^2 dt \quad (6)$$

where $\mathcal{K}(\cdot)$ means stopping gradient.

4.4. View Coverage Adjustment

Performing indoor scene FVS encounters various regions with different attributes including texture, transparency, reflection and view coverage. For instance, when modeling the ceiling which is low texture, Lambertian surface and with few observations, we ought to rely more on the pseudo depth from the geometry scaffold and increase the variance constraint to reduce the ambiguity under photometric loss.

To address the view imbalance, we design an optimizing strategy based on the view coverage to adjust the intensity of the depth and variance loss. Specifically, we first use ray tracing and shadow mapping to obtain the view coverage map of each image from the geometry scaffold. Then we apply different regularization weights accordingly based on the view coverage. The optimizing strategy is defined as follows:

$$\lambda(\mathbf{r}) = \begin{cases} 1, & \text{if } V(\mathbf{r}) > \alpha \\ 1 + \frac{\lambda_{max}-1}{\alpha-1}(\alpha - V(\mathbf{r})), & \text{otherwise} \end{cases} \quad (7)$$

where $V(\mathbf{r})$ is the view coverage of the ray \mathbf{r} , $\lambda(\mathbf{r})$ represents the loss term of the depth and variance regularization, α and λ_{max} are constants for view coverage adjustment. We strongly regularize the depth and variance on areas whose number of observations is less than α , serving as a powerful supplement to the photometric loss (Eq. 3).

In summary, the total loss for optimizing each ray is:

$$\mathcal{L}(\mathbf{r}) = \mathcal{L}_{\text{color}}(\mathbf{r}) + \lambda(\mathbf{r})(\lambda_d \mathcal{L}_{\text{depth}}(\mathbf{r}) + \lambda_w \mathcal{L}_{\text{var}}^w(\mathbf{r}) + \lambda_c \mathcal{L}_{\text{var}}^c(\mathbf{r})) \quad (8)$$

which is a linear combination of all losses presented above with loss weights λ_d , λ_w and λ_c .

5 Experiments

We conduct a series of experiments to evaluate our method and to test whether the proposed modules enable better FVS performance among indoor scenes.

5.1. Datasets

Barbershop. Our FVS task aims to enable 6-DOF navigation in indoor scenes, but evaluating FVS requires considering all candidate view directions and positions, which is difficult to achieve with real-world scenes. To address this, we created the Barbershop [1] FVS dataset using Blender’s Cycles pathtracer [4], including interpolation and extrapolation views to enable evaluation. Barbershop consists of 543 images captured in a human-like trajectory. For interpolation, we used one frame from every six frames. For extrapolation, we grid-sampled camera positions and assigned 24 evenly distributed directions to each position.

ScanNet. ScanNet [5] is a large-scale indoor dataset with 1613 scenes, including ground truth camera intrinsics and camera poses. We choose the scenes chosen in [9] for experiments. In particular, we select the middle frames between training frames as nvs frames, which can be seen as interpolations of the training set. For our extrapolation setting, we select several trajectories different from that in the training set to build up a more difficult test set. The ratio of the training, interpolation and extrapolation is 2:1:1. We show a real trajectory from “scene0050_00” in Fig. 7 (b) for further illustration.

5.2. Implementation Details and Metrics

Our implementation is based on nerf-pytorch [33] and torch-NGP. The training is performed on one NVIDIA RTX3090 GPU. In experiments, we normalize all cameras to be inside a unit sphere so that the scene is bounded within $[-1, 1]^3$. The image resolution in ScanNet is 468×624 after resizing and cropping dark borders, following [19]. We train our model using Adam optimizer with an initial learning rate of 1e-3 and train the network with batches of 1024 rays for 20k iterations. α is set to 9, and λ_{max} is set to 5




Figure 6. **Qualitative results.** From top to bottom, we show the ground truth and extrapolation rendering results obtained by different methods on the ScanNet and Barbershop datasets.

Extrapolation setting	Scene0050_00			Scene0084_00			Scene0580_00			Scene0616_00		
	↑PSNR	↑SSIM	↓LPIPS	↑PSNR	↑SSIM	↓LPIPS	↑PSNR	↑SSIM	↓LPIPS	↑PSNR	↑SSIM	↓LPIPS
NeRF [12]	21.625	0.741	0.450	20.017	0.870	0.426	22.421	0.773	0.466	18.967	0.781	0.481
Instant-NGP [13]	20.930	0.584	0.425	15.176	0.691	0.534	19.525	0.579	0.480	18.073	0.620	0.487
DS-NeRF [7]	22.848	0.764	0.429	20.220	0.858	0.455	22.665	0.778	0.466	19.796	0.794	0.483
Dense Depth Priors [19]	22.800	0.614	0.426	19.232	0.781	0.418	22.399	0.634	0.460	19.986	0.684	0.442
NerfingMVS [28]	16.893	0.613	0.600	13.975	0.699	0.645	18.106	0.662	0.606	14.569	0.645	0.663
SVS [17]	21.231	0.732	0.314	16.878	0.859	0.367	21.575	0.783	0.319	17.530	0.795	0.406
Ours(NGP)	22.479	0.623	0.357	20.553	0.723	0.451	22.043	0.634	0.390	19.352	0.636	0.473
Ours(NeRF)	24.261	0.788	0.388	24.452	0.899	0.401	23.964	0.798	0.434	22.180	0.832	0.433
Interpolation setting	↑PSNR	↑SSIM	↓LPIPS	↑PSNR	↑SSIM	↓LPIPS	↑PSNR	↑SSIM	↓LPIPS	↑PSNR	↑SSIM	↓LPIPS
NeRF [12]	25.215	0.818	0.370	26.953	0.927	0.356	25.672	0.828	0.430	23.278	0.840	0.407
Instant-NGP [13]	24.505	0.675	0.318	21.742	0.800	0.378	23.657	0.677	0.379	21.718	0.681	0.391
DS-NeRF [7]	25.678	0.822	0.368	27.879	0.927	0.365	26.129	0.832	0.431	24.124	0.839	0.428
Dense Depth Priors [19]	24.971	0.668	0.371	22.580	0.820	0.384	24.995	0.682	0.432	21.322	0.697	0.415
NerfingMVS [28]	22.020	0.751	0.459	22.789	0.856	0.486	22.384	0.766	0.511	19.824	0.754	0.530
SVS [17]	23.937	0.830	0.239	18.899	0.901	0.302	24.308	0.844	0.252	19.443	0.829	0.341
Ours(NGP)	25.319	0.699	0.286	25.705	0.820	0.347	25.303	0.709	0.336	22.960	0.702	0.374
Ours(NeRF)	25.840	0.828	0.347	28.004	0.929	0.360	26.157	0.835	0.412	24.362	0.850	0.407

Table 1. **Quantitative results on ScanNet dataset.** We compare our method with NeRF-like and other rendering methods. Our method can significantly improve the performance of both NeRF and instant-NGP.

Barbershop	Extrapolation			Interpolation		
	↑PSNR	↑SSIM	↓LPIPS	↑PSNR	↑SSIM	↓LPIPS
NeRF [12]	24.367	0.868	0.294	32.416	0.957	0.149
Instant-NGP [13]	19.712	0.753	0.397	36.121	0.964	0.056
DS-NeRF [7]	24.105	0.866	0.313	32.081	0.952	0.174
Dense Depth Priors [19]	20.700	0.723	0.470	29.074	0.837	0.284
NerfingMVS [28]	15.577	0.616	0.635	27.265	0.897	0.247
SVS [17]	21.725	0.878	0.210	32.997	0.988	0.019
Ours(NGP)	24.276	0.833	0.235	33.756	0.959	0.052
Ours(NeRF)	26.946	0.891	0.268	32.766	0.958	0.146

Table 2. **Quantitative results on Barbershop dataset.**

throughout all the experiments except for relative ablation. Besides, we prepare a relaxing stage which occupies 10% of training iterations to fine-tune the model. For further implementation details, please refer to the supplementary.

We use visual quality assessment metrics to evaluate the performance of Free View Synthesis, including the PSNR, the SSIM [27], and the LPIPS [36].

5.3. Comparisons with the State-of-the-art Methods

We construct our NeRFVS based on vanilla NeRF [12], term as Ours(NeRF) and instant-NGP [13], term as Ours(NGP) and make comparisons with NeRF-like and other rendering methods, including a) Neural volume rendering methods: NeRF and instant-NGP; b) a method with sparse depth input: DS-NeRF [7] which uses colmap [20] to acquire sparse depth priors; c) methods with depth completion: Dense Depth Priors [19] and NerfingMVS [28], which get dense depth priors from sparse point clouds with completion networks; d) a method with on-surface image feature aggregation: Stable View Synthesis (SVS) [17], which maps view-dependent features extracted from input images to surfaces and generates images from surface features.

In Table 1 and Table 2, we report the interpolation

and extrapolation results across the test images in two datasets. Qualitative results are shown in Figure 6. Comparing the baselines and ours, our method significantly improves the extrapolation ability of both NeRF and instant-NGP. Ours(NeRF) can synthesize high-fidelity images with a consistent 3D layout compared to the baseline. Integrating holistic priors and regularization on radiance distribution into the learning of a neural radiance field efficiently reduces the floaters and distortions, which is a common problem in applying NeRF with indoor scenes, resulting in clean and reasonable FVS performances. Considering Ours(NeRF) produces competitive performances among interpolation settings and significantly outperforms other methods among extrapolation settings considering both PSNR and SSIM, our method achieves new state-of-art performance among indoor scene FVS tasks. The extrapolation ability of instant-NGP is weaker than NeRF, as shown in Table 1, while Ours(NGP) still achieves competitive performance. Considering Ours(NGP) trains and inferences significantly faster than NeRF, about $8\times$ speedup, a spend-quality trade-off exists.

DS-NeRF leverages sparse depth with re-projection error as confidence, which is error-prone compared with the holistic priors since the sparse depth may omit components that rarely appear in training views (example 4, 5, Fig. 6). NerfingMVS and Dense leverage depth completion network to generate per-frame depth maps, which are not 3D consistent compared with our geometry scaffold. Besides, the view imbalance property makes the depth-completion performance vary dramatically, leading to unstable optimization (example 1, 2, Fig. 6). Our method takes the view imbalance and ray regularization into consideration, resulting in better performance on few-shot regions.




Figure 7. Left(a): Performance along the view coverage. Right(b): Trajectory with train, interpolation, and extrapolation viewpoints.

SVS performs well on LPIPS while synthesizing noticeable artifacts and distortions (example 3, 4, 5, Fig. 6), which significantly degrades the experience of indoor roaming. The superior LPIPS performance may be attributed to the direct supervision of an LPIPS-like perceptual loss, as discussed in [2], while our method only minimizes per-pixel loss. SVS heavily relies on the quality of the geometry scaffold. When the geometry scaffold is inaccurate, the renderings will be distorted. Besides, SVS runs COLMAP MVS on training and testing images together to get their geometry scaffold, which gives them an advantage. Our method only relies on training images and applies a robust depth loss to alleviate the influence of the inaccurate geometry, leading to more consistent and high-fidelity results.

5.4. Analysis

We show that each proposed module is of critical importance to the final rendering performance in Table 3. We also discuss the influence of different geometry scaffold sources.

Without Depth Loss. In this experiment, we replace our robust depth loss in Eq. (4) with L2 loss. As aforementioned in Sec. 4.2, our robust depth loss can alleviate the influence of pseudo depth error. L2 loss struggles to deal with the inaccuracy of geometry, causing artifacts and disappearance of objects, *e.g.*, lamps, as shown in Fig. 4.

Without Var Loss. We omit the variance loss, including the $\mathcal{L}_{\text{var}}^w(\mathbf{r})$ and $\mathcal{L}_{\text{var}}^c(\mathbf{r})$ (Eq. (5), Eq. (6)). Even with consistent and dense depth priors, only depending on the depth expectation to guide the optimization of a NeRF over poorly textured regions or few-shot regions (*e.g.* ceiling) is particularly challenging since the density and color distribution is arbitrary. This results in various artifacts when performing extrapolation, as shown in Fig 5.

Without Adjustment. We remove the view coverage adjustment in this experiment, which means we treat all areas as fully observed and apply weak depth and variance regularization. In Table 3, the quantitative results show the degraded performance because of unbalanced regularization.

Performance Gain with View Coverage. We evaluated the effectiveness of NeRFVS in terms of view coverage, as shown in Fig. 7(a). The results demonstrate improvement over all extrapolation results, particularly in regions with insufficient view coverage, compared to NeRF.

Improvement over the scaffold. Experiments in Table. 3 shows that NeRFVS significantly outperforms DS-NeRF

	↑PSNR	↑SSIM	↓LPIPS
ManhattanSDF [9]	19.776	0.775	0.570
NeRF [12]	20.758	0.791	0.456
w/o Depth Loss	23.187	0.821	0.428
w/o Var Loss	23.435	0.822	0.431
w/o Adjustment	23.172	0.821	0.429
DS-NeRF [7] + Colmap MVS	20.906	0.790	0.485
DS-NeRF + Our depth	21.385	0.800	0.464
Ours + Colmap MVS	22.771	0.817	0.439
Full model	23.714	0.829	0.414

Table 3. **Ablation studies on ScanNet.** We report the mean image quality metrics among extrapolation setting with Ours(NeRF).

with our depth. Besides, DS-NeRF gains minor improvement by replacing the depth from Colmap MVS with our depth. These results prove that the improvement over prior work mainly depends on the method, rather than the choice of the scaffold. Though Colmap MVS is not suitable for our NeRFVS, ours with Colmap MVS still outperforms DS-NeRF with Colmap MVS, showing that our method is generic and can be applied to different scaffolds.

6 Conclusion

In this study, we introduce NeRFVS, a novel method enabling NeRF to perform free navigation for indoor scenes. Our method utilizes the holistic prior from a neural-reconstructed geometric scaffold to guide the optimization of NeRF and boosts the rendering accuracy by introducing the depth and variance constraints with dynamical adjustment. We show that the inaccuracy of the geometry scaffold degrades the extrapolation performance and propose a robust depth loss to address this issue. Variance regularization and view coverage-based training strategy improve the rendering quality on the few shot regions. Experiments show that NeRFVS significantly improves the extrapolation ability of NeRF-based methods. We believe our method leaps a meaningful step toward extending NeRF to free navigation.

Limitations. NeRFVS reduces geometry errors in regions with sufficient observations but relies on pseudo depth in few-shot regions. This may result in poor extrapolation if the geometry is exceptionally poor. The way we generate the view coverage map is an adequate approximation in plainer regions. While in some extreme cases, such as holes and thin objects, this could result in errors. Besides, our view coverage only considers the number of observations, regardless of the emitting directions, which may cause degradation in view-dependent appearance.

Acknowledgments. This work was supported by NSFC 62176159, Natural Science Foundation of Shanghai 21ZR1432200, Shanghai Municipal Science and Technology Major Project 2021SHZDZX0102, and the Fundamental Research Funds for the Central Universities. We thank Ruofan Liang for inspiring and valuable discussions. We thank MindSpore [24] for the partial support to this work, which is a new deep learning computing framework.

References

- [1] Barbershop. https://svn.blender.org/svnroot/bf-blender/trunk/lib/benchmarks/cycles/barbershop_interior/. 5
- [2] Jonathan T Barron, Ben Mildenhall, Dor Verbin, Pratul P Srinivasan, and Peter Hedman. Mip-nerf 360: Unbounded anti-aliased neural radiance fields. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 5470–5479, 2022. 3, 8
- [3] Chris Buehler, Michael Bosse, Leonard McMillan, Steven Gortler, and Michael Cohen. Unstructured lumigraph rendering. In *Proceedings of the 28th Annual Conference on Computer Graphics and Interactive Techniques, SIGGRAPH '01*, page 425–432, New York, NY, USA, 2001. Association for Computing Machinery. 2
- [4] Blender Online Community. Blender - a 3d modelling and rendering package, 2018. 5
- [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. 5
- [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. 3
- [7] 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. 7, 8
- [8] Ross Girshick. Fast r-cnn. In *Proceedings of the IEEE international conference on computer vision*, pages 1440–1448, 2015. 4
- [9] Haoyu Guo, Sida Peng, Haotong Lin, Qianqian Wang, Guofeng Zhang, Hujun Bao, and Xiaowei Zhou. Neural 3d scene reconstruction with the manhattan-world assumption. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 5511–5520, 2022. 2, 3, 4, 5, 8
- [10] Peter J Huber. Robust estimation of a location parameter. In *Breakthroughs in statistics*, pages 492–518. Springer, 1992. 4
- [11] Mijeong Kim, Seonguk Seo, and Bohyung Han. Infonerf: Ray entropy minimization for few-shot neural volume rendering. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 12912–12921, 2022. 3
- [12] 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. *Communications of the ACM*, 65(1):99–106, 2021. 1, 2, 7, 8
- [13] Thomas Müller, Alex Evans, Christoph Schied, and Alexander Keller. Instant neural graphics primitives with a multiresolution hash encoding. *arXiv preprint arXiv:2201.05989*, 2022. 3, 7
- [14] Michael Niemeyer, Jonathan T Barron, Ben Mildenhall, Mehdi SM Sajjadi, Andreas Geiger, and Noha Radwan. Regnerf: Regularizing neural radiance fields for view synthesis from sparse inputs. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 5480–5490, 2022. 3
- [15] Julien Philip, Sébastien Mordant, Michaël Gharbi, and George Drettakis. Free-viewpoint indoor neural relighting from multi-view stereo. *ACM Transactions on Graphics (TOG)*, 40(5):1–18, 2021. 2
- [16] Gernot Riegler and Vladlen Koltun. Free view synthesis. In *European Conference on Computer Vision*, pages 623–640. Springer, 2020. 2
- [17] Gernot Riegler and Vladlen Koltun. Stable view synthesis. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 12216–12225, 2021. 2, 7
- [18] 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. 3
- [19] 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. 5, 7
- [20] Johannes L Schonberger and Jan-Michael Frahm. Structure-from-motion revisited. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 4104–4113, 2016. 2, 7
- [21] Edgar Sucar, Shikun Liu, Joseph Ortiz, and Andrew J Davison. imap: Implicit mapping and positioning in real-time. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 6229–6238, 2021. 4
- [22] Jiaming Sun, Yiming Xie, Linghao Chen, Xiaowei Zhou, and Hujun Bao. Neuralrecon: Real-time coherent 3d reconstruction from monocular video. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 15598–15607, 2021. 2
- [23] Yuqi Sun, Shili Zhou, Ri Cheng, Weimin Tan, Bo Yan, and Lang Fu. Learning robust image-based rendering on sparse scene geometry via depth completion. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 7813–7823, 2022. 2
- [24] Huawei Technologies. Mindspore. <https://www.mindspore.cn/>. 8
- [25] Jiepeng Wang, Peng Wang, Xiaoxiao Long, Christian Theobalt, Taku Komura, Lingjie Liu, and Wenping Wang. Neuris: Neural reconstruction of indoor scenes using normal priors. In *Computer Vision—ECCV 2022: 17th European Conference, Tel Aviv, Israel, October 23–27, 2022, Proceedings, Part XXXII*, pages 139–155. Springer, 2022. 2
- [26] 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. 2
- [27] 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. 7
- [28] 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 *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 5610–5619, 2021. 7
- [29] Xiuchao Wu, Jiamin Xu, Zihan Zhu, Hujun Bao, Qixing Huang, James Tompkin, and Weiwei Xu. Scalable neural indoor scene rendering. *ACM Transactions on Graphics (TOG)*, 41(4):1–16, 2022. 2
- [30] Deji Xu, Yifan Jiang, Peihao Wang, Zhiwen Fan, Humphrey Shi, and Zhangyang Wang. Sinnerf: Training neural radiance fields on complex scenes from a single image. *arXiv preprint arXiv:2204.00928*, 2022. 3
- [31] Jiamin Xu, Xiuchao Wu, Zihan Zhu, Qixing Huang, Yin Yang, Hujun Bao, and Weiwei Xu. Scalable image-based indoor scene rendering with reflections. *ACM Transactions on Graphics (TOG)*, 40(4):1–14, 2021. 2
- [32] Yu-Ying Yeh, Zhengqin Li, Yannick Hold-Geoffroy, Rui Zhu, Zexiang Xu, Miloš Hašan, Kalyan Sunkavalli, and Manmohan Chandraker. Photoscene: Photorealistic material and lighting transfer for indoor scenes. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 18562–18571, 2022. 2
- [33] Lin Yen-Chen. Nerf-pytorch. <https://github.com/yenchenlin/nerf-pytorch/>, 2020. 5
- [34] Zehao Yu, Songyou Peng, Michael Niemeyer, Torsten Sattler, and Andreas Geiger. Monosdf: Exploring monocular geometric cues for neural implicit surface reconstruction. *ArXiv*, abs/2206.00665, 2022. 2
- [35] Jian Zhang, Yuanqing Zhang, Huan Fu, Xiaowei Zhou, Bowen Cai, Jinchu Huang, Rongfei Jia, Binqiang Zhao, and Xing Tang. Ray priors through reprojection: Improving neural radiance fields for novel view extrapolation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 18376–18386, 2022. 2
- [36] 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. 7
- [37] Zihan Zhu, Songyou Peng, Viktor Larsson, Weiwei Xu, Hujun Bao, Zhaopeng Cui, Martin R Oswald, and Marc Pollefeys. Nice-slam: Neural implicit scalable encoding for slam. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 12786–12796, 2022. 4