



ViP-NeRF: Visibility Prior for Sparse Input Neural Radiance Fields

Nagabhushan Somraj
Indian Institute of Science
Bengaluru, India
nagabhushans@iisc.ac.in

Rajiv Soundararajan
Indian Institute of Science
Bengaluru, India
rajivs@iisc.ac.in

ABSTRACT

Neural radiance fields (NeRF) have achieved impressive performances in view synthesis by encoding neural representations of a scene. However, NeRFs require hundreds of images per scene to synthesize photo-realistic novel views. Training them on sparse input views leads to overfitting and incorrect scene depth estimation resulting in artifacts in the rendered novel views. Sparse input NeRFs were recently regularized by providing dense depth estimated from pre-trained networks as supervision, to achieve improved performance over sparse depth constraints. However, we find that such depth priors may be inaccurate due to generalization issues. Instead, we hypothesize that the visibility of pixels in different input views can be more reliably estimated to provide dense supervision. In this regard, we compute a visibility prior through the use of plane sweep volumes, which does not require any pre-training. By regularizing the NeRF training with the visibility prior, we successfully train the NeRF with few input views. We reformulate the NeRF to also directly output the visibility of a 3D point from a given viewpoint to reduce the training time with the visibility constraint. On multiple datasets, our model outperforms the competing sparse input NeRF models including those that use learned priors. The source code for our model can be found on our project page: <https://nagabhushansn95.github.io/publications/2023/ViP-NeRF.html>.

CCS CONCEPTS

• **Computing methodologies** → **Rendering; Volumetric models; Computer vision; Computational photography**; 3D imaging; Reconstruction.

KEYWORDS

neural rendering, novel view synthesis, sparse input NeRF, visibility prior, plane sweep volumes

ACM Reference Format:

Nagabhushan Somraj and Rajiv Soundararajan. 2023. ViP-NeRF: Visibility Prior for Sparse Input Neural Radiance Fields. In *Special Interest Group on Computer Graphics and Interactive Techniques Conference Conference Proceedings (SIGGRAPH '23 Conference Proceedings)*, August 06–10, 2023, Los Angeles, CA, USA. ACM, New York, NY, USA, 11 pages. <https://doi.org/10.1145/3588432.3591539>

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.
SIGGRAPH '23 Conference Proceedings, August 06–10, 2023, Los Angeles, CA, USA
© 2023 Copyright held by the owner/author(s). Publication rights licensed to ACM.
ACM ISBN 979-8-4007-0159-7/23/08...\$15.00
<https://doi.org/10.1145/3588432.3591539>

1 INTRODUCTION

The goal of novel view synthesis is to synthesize a scene from novel viewpoints given RGB images of a few other viewpoints and their relative camera poses. By representing the scene implicitly using multi-layer perceptrons (MLP) and employing volume rendering, neural radiance fields (NeRF) [Barron et al. 2021, 2022; Liu et al. 2022; Mildenhall et al. 2020] have achieved impressive view synthesis performance. Such superior performance is usually achieved when a large number of views is input to train the NeRF. However, in multiple applications such as virtual or augmented reality, telepresence, robotics, and autonomous driving, very few input images may be available for training [Niemeyer et al. 2022]. In such settings, external sensors or a pre-calibrated fixed camera array may be employed to obtain accurate camera poses. Thus, there is a need to train NeRFs with few input views referred to as the sparse input NeRF problem.

The key challenge with sparse input images is that the volume rendering equations in NeRF are under-constrained, leading to solutions that overfit the input views. This results in uncertain and inaccurate depth in the learned representation. Synthesized novel views in such cases contain extreme distortions such as blur, ghosting, and floater artifacts [Niemeyer et al. 2022; Roessle et al. 2022]. Recent works have proposed different approaches to constrain the training of NeRF to output visually pleasing novel views. While a few recent works [Yang et al. 2022; Zhang et al. 2021; Zhou and Tulsiani [n. d.]] focus on training NeRF models on a specific category of objects such as chairs or airplanes, we focus on training category agnostic sparse input NeRF models [Niemeyer et al. 2022]. Such prior work can be broadly classified into conditional NeRF models and other regularization approaches.

The conditional NeRF models employ a latent representation of the scene obtained by pre-training on a large dataset of different scenes [Chen et al. [n. d.]; Hamdi et al. [n. d.]; Johari et al. 2022; Wang et al. 2021; Yu et al. 2021] to condition the NeRF. The latent prior helps overcome the limitation on the number of views by enabling the NeRF model to effectively understand the scene. Such an approach is popular even when only a single image of the scene is available as input to the NeRF [Cai et al. 2022; Lin et al. 2023; Xu et al. 2022]. Different from the above, MetaNeRF [Tancik et al. 2021] learns the latent information as initial weights of the NeRF MLPs by employing meta-learning. However, the pre-trained latent prior could suffer from poor generalization on a given target scene [Niemeyer et al. 2022]. Thus, we believe that there is a need to study the sparse-input NeRF without conditioning the NeRF on latent representations.

On the other hand, regularization based approaches constrain the NeRF training with novel loss functions to yield better solutions. DS-NeRF [Deng et al. 2022] uses sparse depth provided by a structure from motion (SfM) model as additional supervision for the NeRF. To

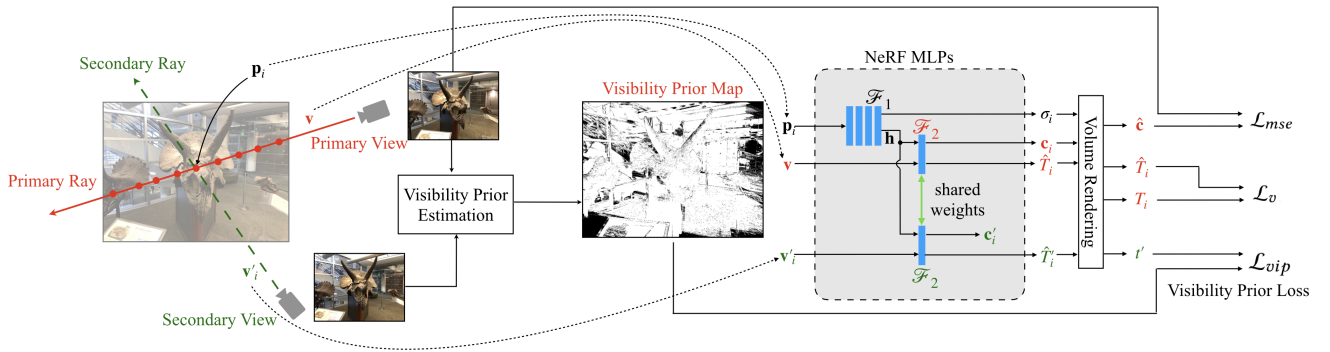


Figure 1: Overview of ViP-NeRF architecture. Given the images from **primary** and **secondary** views, we estimate a visibility prior map in the primary view and use it to supervise the visibility of pixels as predicted by the NeRF. Specifically, we cast a ray through a randomly selected pixel in the primary view and sample 3D points along the ray. For every point p_i , we use the NeRF MLPs to obtain its visibility in primary and secondary views, along with volume density σ_i and color c_i . Volume rendering outputs visibility t' of the chosen pixel in the secondary view which is supervised by the visibility prior. \mathcal{L}_v constrains the visibilities \hat{T}_i output by network and T_i computed using volume rendering to be consistent with each other.

provide richer dense supervision, DDP-NeRF [Roessle et al. 2022] completes the sparse depth map using a pre-trained convolutional neural network (CNN). However, the requirement of pre-training on a large dataset of scenes is cumbersome and the dense depth prior may suffer from generalization errors. RegNeRF [Niemeyer et al. 2022] and InfoNeRF [Kim et al. 2022] impose constraints to promote depth smoothness and reduce depth uncertainty respectively. However, in our experiments, we observe that these methods are still inferior to DS-NeRF on popular datasets. This motivates the exploration of other reliable features for dense supervision to constrain the NeRF in addition to sparse depth supervision.

In our work, we explore the use of regularization in terms of visibility of any pixel from a pair of viewpoints. Here visibility of a pixel refers to whether the corresponding object is seen in both the viewpoints. For example, foreground objects are typically visible in multiple views whereas the background objects may be partially occluded. The visibility of a pixel in different views relies more on the relative depth of the scene objects than the absolute depth. We hypothesize that, given sparse input views, it may be easier to estimate the relative depth and visibility instead of the absolute depth. Thus, the key idea of our work is to regularize the NeRF with a dense visibility prior estimated using the given sparse input views. This allows the NeRF to learn better scene representation. We refer to our Visibility Prior regularized NeRF model as ViP-NeRF.

To obtain the visibility prior, we employ the plane sweep volumes (PSV) [Collins 1996] that have successfully been used in depth estimation [Gallup et al. 2007; Ha et al. 2016; Im et al. [n. d.]; Yang and Pollefeys 2003] and view synthesis models [Zhou et al. 2018]. We create the PSV by warping one of the images to the view of the other at different depths (or planes) and compare them to obtain error maps. We determine a binary visibility map for each pixel based on the corresponding errors in the PSV. We regularize the NeRF training by using such a map as supervision for every pair of input views. We use the visibility prior in conjunction with the depth prior from DS-NeRF [Deng et al. 2022], where the former

provides a dense prior on relative depth while the latter provides a sparse prior on absolute depth. Note that the estimation of our visibility prior does not require any pre-training on a large dataset.

Regularizing the NeRF with a dense visibility prior is computationally intensive and can lead to impractical training times. We reformulate the NeRF to directly and additionally output visibility to impose the regularization in a computationally efficient manner. We conduct experiments on two popular datasets to demonstrate the efficacy of the visibility prior for sparse input NeRF.

The main contributions of our work are as follows.

- We introduce visibility regularization to train the NeRF with sparse input views and refer to our model as ViP-NeRF.
- We estimate the dense visibility prior reliably using plane sweep volumes.
- We reformulate the NeRF MLP to output visibility thereby significantly reducing the training time.
- We achieve the state-of-the-art performance of sparse input NeRFs on multiple datasets.

2 RELATED WORK

2.1 Novel View Synthesis

Novel view synthesis methods typically use one or more input views to synthesize the scene from novel viewpoints. Recent pieces of work focus on obtaining volumetric 3D representations of the scene that can be computed once to render any viewpoint later. Zhou et al. [2018] propose multi-plane image (MPI) representations for view synthesis. Srinivasan et al. [2019] further extend this by infilling the occluded regions in the MPIs. Wiles et al. [2020] study an extreme case with a single input image and generate novel views by employing a monocular depth estimation network for scene re-projection. In contrast to the above explicit representations, neural radiance fields [Mildenhall et al. 2020] use an implicit representation through coordinate-based neural networks. Although NeRFs achieve excellent performance, they require dense input views for

training. In this work, we focus on solving this problem, i.e. to train a NeRF given very few input views.

2.2 Sparse Input NeRF

Several recent works have studied sparse input NeRF by regularizing the NeRF with various priors. One of the early works, Diet-NeRF [Jain et al. 2021], hallucinates novel viewpoints during training and constrains the NeRF to generate novel views similar to the input images in the CLIP [Radford et al. 2021] representation space. DS-NeRF improves the performance by using fine-grained supervision at the pixel level using a sparse depth estimated by an SfM model [Deng et al. 2022]. DDP-NeRF further completes the sparse depth using a pre-trained network to obtain dense depth along with uncertainty estimates. Uncertainty modeling allows DDP-NeRF to relax the depth supervision at locations where the dense depth estimation is not confident. However, the completed depth may contain errors that may adversely affect the performance. Diffu-NeRF [Wynn and Turmukhambetov [n. d.]] instead employs a pre-trained denoising diffusion model to regularize the distribution of RGB-D patches in novel viewpoints. In contrast, our work uses a more reliable visibility prior which can be estimated without the use of sophisticated CNNs and does not require pre-training on a large dataset of scenes.

Instead of depth estimates as priors, RegNeRF [Niemeyer et al. 2022] regularizes the NeRF using depth smoothness constraints on the rendered patches in the hallucinated viewpoints. Different from depth regularization models, InfoNeRF [Kim et al. 2022] tries to circumvent overfitting by encouraging concentration of volume density along a ray. In addition, it also minimizes the variation of volume density distributions along rays of two nearby viewpoints. Although these constraints are meaningful, our visibility prior imposes constraints across multiple views and can exploit the structure of the problem more effectively.

2.3 Single Image NeRF

Recently, there is increased interest in training NeRFs with a single input image [Lin et al. 2023; Xu et al. 2022]. A common thread in single image NeRF models is to use an encoder to obtain a latent representation of the input image. A NeRF based decoder conditioned on the representation, outputs volume density and color at given 3D points. For example, pix2NeRF [Cai et al. 2022] combines π -GAN [Chan et al. 2021] with NeRF to render photo-realistic images of objects or human faces. Gao et al. [[n. d.]] focus on human faces alone and use a more structured approach by exploiting facial geometry. MINE [Li et al. 2021] combines NeRF with MPI by replacing the MLP based implicit representation with an MPI based explicit representation in the decoder. Lin et al. [2023] obtain a richer latent representation by fusing global and local features obtained using a vision transformer and CNN respectively. Different from the above models, Wimbauer et al. [[n. d.]] use the MLP decoder to predict volume density alone and obtain the color by directly sampling from the given images. However, a common drawback of these models is the need for pre-training. Thus the performance may be inferior when testing on a generic scene.




Figure 2: A toy example to illustrate the computation of visibility prior. The scene contains a blue sphere and a brown box and the relative pose between the views is a translation in x direction. The secondary view image is warped to the primary view at different depth planes to create a PSV and compared with the primary view image to obtain error maps. We observe that the brown square and the blue circle are matched better in the second and third planes respectively leading to lower error (denoted as white) in the respective error maps. The minimum error across all the planes is thresholded to obtain the visibility prior map corresponding to the primary view image. The right portion of the sphere which is occluded in the secondary view image is denoted in black in the visibility map.

3 NERF PRELIMINARIES

We first provide a brief introduction to NeRF and define the notations for subsequent use. A neural radiance field is an implicit representation of a scene using two multi-layer perceptrons (MLP). Given a set of images of a scene with corresponding camera poses, a pixel \mathbf{q} is selected at random, and a ray \mathbf{r} is passed from the camera center \mathbf{o} through \mathbf{q} . Let $\mathbf{p}_1, \mathbf{p}_2, \dots, \mathbf{p}_N$ be N randomly sampled 3D points along \mathbf{r} . If \mathbf{d} is the direction vector of \mathbf{r} and z_i is the depth of a 3D point \mathbf{p}_i , $i \in \{1, 2, \dots, N\}$, then $\mathbf{p}_i = \mathbf{o} + z_i \mathbf{d}$. An MLP \mathcal{F}_1 is trained to predict the volume density σ_i at \mathbf{p}_i as

$$\sigma_i, \mathbf{h}_i = \mathcal{F}_1(\mathbf{p}_i), \quad (1)$$

where \mathbf{h}_i is a latent representation. A second MLP \mathcal{F}_2 then predicts the color using \mathbf{h}_i and the viewing direction $\mathbf{v} = \mathbf{d}/\|\mathbf{d}\|$ as

$$\mathbf{c}_i = \mathcal{F}_2(\mathbf{h}_i, \mathbf{v}). \quad (2)$$

Let the distance between two consecutive samples \mathbf{p}_i and \mathbf{p}_{i+1} be $\delta_i = z_{i+1} - z_i$. The visibility or transmittance of \mathbf{p}_i is then given by

$$T_i = \exp\left(-\sum_{j=1}^{i-1} \delta_j \sigma_j\right). \quad (3)$$

The weight or contribution of \mathbf{p}_i in rendering the color $\hat{\mathbf{c}}$ of pixel \mathbf{q} is computed as

$$w_i = T_i (1 - \exp(-\delta_i \sigma_i)) \quad (4)$$

to obtain

$$\hat{\mathbf{c}} = \sum_{i=1}^N w_i \mathbf{c}_i. \quad (5)$$

The MLPs are trained using mean squared error loss with the true color \mathbf{c} of \mathbf{q} as

$$\mathcal{L}_{mse} = \|\mathbf{c} - \hat{\mathbf{c}}\|^2. \quad (6)$$

Table 1: Quantitative results on RealEstate-10K dataset.

Model	learned prior	2 views			3 views			4 views		
		LPIPS ↓	SSIM ↑	PSNR ↑	LPIPS ↓	SSIM ↑	PSNR ↑	LPIPS ↓	SSIM ↑	PSNR ↑
InfoNeRF		0.6796	0.4653	12.30	0.6979	0.4024	11.15	0.6745	0.4298	11.52
DietNeRF	✓	0.5730	0.6131	15.90	0.5365	0.6190	16.60	0.5337	0.6282	16.89
RegNeRF		0.5307	0.5709	16.14	0.4675	0.6096	17.38	0.4831	0.6068	17.46
DS-NeRF		0.4273	0.7223	21.40	0.3930	0.7554	23.73	0.3961	0.7575	24.24
DDP-NeRF	✓	0.2527	0.7890	21.44	0.2240	0.8223	23.10	0.2190	0.8270	24.17
ViP-NeRF		0.1704	0.8087	24.48	0.1441	0.8505	27.21	0.1386	0.8588	28.13

Table 2: Quantitative results on NeRF-LLFF dataset.

Model	learned prior	2 views			3 views			4 views		
		LPIPS ↓	SSIM ↑	PSNR ↑	LPIPS ↓	SSIM ↑	PSNR ↑	LPIPS ↓	SSIM ↑	PSNR ↑
InfoNeRF		0.7561	0.2095	9.23	0.7679	0.1859	8.52	0.7701	0.2188	9.25
DietNeRF	✓	0.7265	0.3209	11.89	0.7254	0.3297	11.77	0.7396	0.3404	11.84
RegNeRF		0.4402	0.4872	16.90	0.3800	0.5600	18.62	0.3446	0.6056	19.83
DS-NeRF		0.4548	0.5068	17.06	0.4077	0.5686	19.02	0.3825	0.6016	20.11
DDP-NeRF	✓	0.4223	0.5377	17.21	0.4178	0.5610	17.90	0.3821	0.5999	19.19
ViP-NeRF		0.4017	0.5222	16.76	0.3750	0.5837	18.92	0.3593	0.6085	19.57

4 METHOD

We illustrate the outline of our model in Fig. 1. The core idea of our work is that when only a few multiview images are available for NeRF training, the visibility of a pixel in different views can be more reliably densely estimated as compared to its absolute depth. In this regard, we introduce visibility regularization to train the NeRF with sparse input views in Sec. 4.1. To impose the visibility regularization, we obtain a binary visibility prior map for every pair of input training images, which we explain in Sec. 4.2. Finally, to reduce the training time, we design a method to efficiently predict the visibility of a given pixel in different views in Sec. 4.3. Sec. 4.4 summarizes the various loss functions used in training our model.

4.1 Visibility Regularization

Recall from Sec. 3 that NeRF trains MLPs by picking a random pixel \mathbf{q} and predicting the color of \mathbf{q} using the MLPs and volume rendering. Without loss of generality, we refer to the view corresponding to the ray \mathbf{r} passing through \mathbf{q} as the primary view and choose any other view as a secondary view. NeRF then samples N candidate 3D points, $\mathbf{p}_1, \mathbf{p}_2, \dots, \mathbf{p}_N$, along \mathbf{r} . Let T'_i be the visibility of \mathbf{p}_i from the secondary view, computed similar to Eq. (3). We define the visibility of pixel \mathbf{q} in the secondary view, $t'(\mathbf{q})$, as the weighted visibilities of all the candidate 3D points \mathbf{p}_i analogous to Eq. (5) as

$$t'(\mathbf{q}) = \sum_{i=1}^N w_i T'_i \in [0, 1], \quad (7)$$

where w_i are obtained through Eq. (4). We omit the dependence of w_i and T'_i on \mathbf{q} in the above equation for ease of reading. We obtain a prior $\tau'(\mathbf{q}) \in \{0, 1\}$ on the visibility $t'(\mathbf{q})$ as described in Sec. 4.2. We constrain the visibility $t'(\mathbf{q})$ to match the prior $\tau'(\mathbf{q})$. However, we find that the prior may be unreliable at pixels where $\tau' = 0$, as

we describe in Sec. 4.2. Hence, we do not impose any visibility loss on such pixels and formulate our visibility prior loss as

$$\mathcal{L}_{vip}(\mathbf{q}) = \max(\tau'(\mathbf{q}) - t'(\mathbf{q}), 0). \quad (8)$$

Note that our loss function constrains the NeRF across pairs of views, unlike previous works which regularize [Niemeyer et al. 2022; Roessle et al. 2022] in a given view alone. We believe that this leads to a better regularization for synthesizing novel views.

4.2 Visibility Prior

Given primary and secondary views, our goal is to estimate whether every pixel in the primary view is also visible in the secondary view through a binary visibility prior $\tau'(\mathbf{q})$. We employ plane sweep volumes to compute the visibility prior. We illustrate the computation of the visibility prior with a toy example in Fig. 2. Here, we warp the image in the secondary view to the primary view using the camera parameters at different depths varying between the near depth z_{\min} and far depth z_{\max} . We sample D depths uniformly in inverse depth similar to StereoMag [Zhou et al. 2018]. The set of warped images is referred to as plane sweep volume (PSV) [Huang et al. 2018].

Let $I^{(1)}$ be the image in the primary view and $I_k^{(2)}$ be the set of D warped images, where $k \in \{0, 1, \dots, D-1\}$ denotes the plane index. We then compute the error map E_k of the warped secondary image with the primary image at each plane k of the PSV as

$$E_k = \|I^{(1)} - I_k^{(2)}\|_1, \quad (9)$$

where the norm is computed across the color channels. We determine the visibility prior τ' for pixel \mathbf{q} by thresholding the minimum

error across all the planes as

$$e(\mathbf{q}) = \min_k E_k(\mathbf{q}),$$

$$\tau'(\mathbf{q}) = \mathbb{1}_{\{\exp(-e(\mathbf{q})/\gamma) > 0.5\}}, \quad (10)$$

where γ is a hyper-parameter.

Intuitively, for a given pixel \mathbf{q} , a lower error in any of the planes indicates the presence of a matching pixel in the secondary view, i.e. \mathbf{q} is visible in the secondary view. Note that this holds true when the intensity of pixels does not change significantly across views, which is typical for most of the objects in real-world scenes [Li et al. 2021]. Consequently, the absence of a matching point across all the planes may indicate that \mathbf{q} is not visible in the secondary view or \mathbf{q} belongs to a highly specular object whose color varies significantly across different viewpoints. Thus, our prior is used to regularize the NeRF only in the first case above i.e. the pixels for which we find a match. Following the above procedure, we obtain the visibility prior for every pair of images obtained from the training set, by treating either image in the pair as the primary or the secondary view.

4.3 Efficient Prediction of Visibility

Recall that imposing \mathcal{L}_{vip} in Eq. (8) requires computing visibility T'_i in the secondary view for every \mathbf{p}_i . A naive approach to compute T'_i involves sampling up to N points along a secondary ray from the secondary view camera origin to \mathbf{p}_i and querying the NeRF MLP \mathcal{F}_1 for each of these points. Thus, obtaining $\tau'(\mathbf{q})$ in Eq. (7) requires upto N^2 MLP queries, which increases the training time making it computationally prohibitive. We overcome this limitation by reformulating the NeRF MLP \mathcal{F}_2 to also output a view-dependent visibility of a given 3D point as,

$$c_i, \hat{T}_i = \mathcal{F}_2(\mathbf{h}_i, \mathbf{v}); \quad c'_i, \hat{T}'_i = \mathcal{F}_2(\mathbf{h}_i, \mathbf{v}'_i), \quad (11)$$

where \mathbf{v}'_i is the viewing direction of the secondary ray. We use the MLP output \hat{T}'_i instead of T'_i in Eq. (7).

Note that to output \hat{T}'_i , we need not query \mathcal{F}_1 again and can reuse \mathbf{h}_i obtained from Eq. (1). We only need to query \mathcal{F}_2 additionally and since \mathcal{F}_2 is a single layer MLP and significantly smaller than \mathcal{F}_1 , the additional computational burden is negligible. Thus, directly obtaining the secondary visibility \hat{T}'_i of \mathbf{p}_i through Eq. (11) allows us to compute $\tau'(\mathbf{q})$ in Eq. (7) using only N queries of the MLP \mathcal{F}_1 , as opposed to N^2 queries in the naive approach.

However, the use of \hat{T}'_i in place of T'_i regularizes the NeRF training only if the two quantities are close to each other. Thus, we introduce an additional loss to constrain the visibility \hat{T}_i output by \mathcal{F}_2 to be consistent with the visibility T_i computed using Eq. (3) as

$$\mathcal{L}_v = \sum_{i=1}^N \left(\left(\text{SG}(T_i) - \hat{T}_i \right)^2 + \left(T_i - \text{SG}(\hat{T}_i) \right)^2 \right), \quad (12)$$

where $\text{SG}(\cdot)$ denotes the stop-gradient operation. The first term in the above loss function uses T_i as a target and brings \hat{T}_i closer to it. On the other hand, since \hat{T}_i gets additionally updated directly based on the visibility prior, the second term helps transfer such updates to \mathcal{F}_1 more efficiently than backpropagation through \mathcal{F}_2 .

Table 3: Comparison of reliability of priors used in different models. The reference visibility is obtained using NeRF trained with dense input views.

model	RealEstate-10K			NeRF-LLFF		
	Prec. \uparrow	Rec. \uparrow	F1 \uparrow	Prec. \uparrow	Rec. \uparrow	F1 \uparrow
ViP-NeRF	0.97	0.83	0.89	0.82	0.85	0.83
DDP-NeRF	0.98	0.53	0.66	0.86	0.33	0.47

4.4 Overall Loss

Similar to DS-NeRF [Deng et al. 2022], we also use the sparse depth given by an SfM model to supervise the NeRF as

$$\mathcal{L}_{sd} = \|z - \hat{z}\|^2, \quad (13)$$

where z is the depth provided by the SfM model, $\hat{z} = \sum_i w_i z_i$ is the depth estimated by NeRF and w_i are obtained in Eq. (4). Our overall loss for ViP-NeRF is a linear combination of the losses obtained in Eq. (6), Eq. (8), Eq. (12) and Eq. (13) as

$$\mathcal{L} = \lambda_1 \mathcal{L}_{mse} + \lambda_2 \mathcal{L}_{sd} + \lambda_3 \mathcal{L}_{vip} + \lambda_4 \mathcal{L}_v, \quad (14)$$

where $\lambda_1, \lambda_2, \lambda_3$ and λ_4 are hyper-parameters. We note that \mathcal{L}_{vip} is always employed in conjunction with \mathcal{L}_v to make the learning computationally tractable.

5 EXPERIMENTS

5.1 Evaluation Setup

We conduct experiments on two different datasets, namely RealEstate-10K and NeRF-LLFF. We evaluate all the models in the more challenging setup of 2, 3, or 4 input views, unlike prior work which use 9–18 input views [Jain et al. 2021; Roessle et al. 2022]. The test set is retained to be the same across all different settings for both datasets.

RealEstate-10K [Zhou et al. 2018] dataset is commonly used to evaluate view synthesis models [Han et al. 2022; Tucker and Snavely 2020] and contains videos of camera motion, both indoor and outdoor. The dataset also provides the camera intrinsics and extrinsics for all the frames. For our experiments, we choose 5 scenes from the test set, each containing 50 frames with a spatial resolution of 1024×576 . In each scene, we reserve every 10th frame for training and use the remaining 45 frames for testing. Please refer to the supplementary for more details on the choice of scenes.

NeRF-LLFF [Mildenhall et al. 2019] dataset is used to evaluate the performance of various NeRF Models including sparse input NeRF models. It consists of 8 forward-facing scenes with a variable number of frames per scene at a spatial resolution of 1008×756 . Following RegNeRF [Niemeyer et al. 2022], we use every 8th frame for testing. For training, we pick 2, 3 or 4 frames uniformly among the remaining frames following RegNeRF [Niemeyer et al. 2022].

Evaluation measures. We quantitatively evaluate the methods using LPIPS [Zhang et al. 2018], structural similarity (SSIM) [Wang et al. 2004], and peak signal to noise ratio (PSNR) measures. For LPIPS, we use the v0.1 release with the AlexNet [Krizhevsky et al. 2012] backbone as suggested by the authors.




Figure 3: Qualitative examples on RealEstate-10K dataset with two input views. We observe that the predictions of ViP-NeRF are close to the ground truth, while those of other models suffer from various distortions. In particular, DDP-NeRF blurs regions of the frame near the left door and contains black floater artifacts.




Figure 4: Qualitative examples on RealEstate-10K and NeRF-LLFF dataset with two, three, and four input views. We observe that ViP-NeRF models specular regions better as the number of input views increases. For example, in the first row, the reflection of the chair is better reconstructed as the number of views increases.

Table 4: Evaluation of depth estimated by different models with two input views. The reference depth is obtained using NeRF trained with dense input views. The depth RMSE on the two datasets are of different orders on account of different depth ranges.

model	RealEstate-10K		NeRF-LLFF	
	RMSE ↓	SROCC ↑	RMSE ↓	SROCC ↑
ViP-NeRF	1.6411	0.7702	45.6314	0.6184
DDP-NeRF	1.7211	0.7544	46.6268	0.6136

5.2 Comparisons and Implementation Details

We compare the performance of our model with other sparse input NeRF models such as DDP-NeRF [Roessle et al. 2022] and DietNeRF [Jain et al. 2021] which use learned priors to constrain the NeRF training. We also compare with DS-NeRF [Deng et al. 2022], InfoNeRF [Kim et al. 2022], and RegNeRF [Niemeyer et al. 2022] that

Table 5: Ablation experiments on both the datasets with two input views.

model	RealEstate-10K		NeRF-LLFF	
	LPIPS ↓	SSIM ↑	LPIPS ↓	SSIM ↑
ViP-NeRF	0.1704	0.8087	0.4017	0.5222
w/o sparse depth	0.2754	0.7588	0.5056	0.4631
w/o dense visibility	0.4273	0.7223	0.4548	0.5068

do not use learned priors. We train the models for 50k iterations on both datasets using the code provided by the respective authors.

For ViP-NeRF, we use Adam optimizer with a learning rate of $5e-4$ that exponentially decays to $5e-6$ following NeRF [Mildenhall et al. 2020]. We set the loss weights such that the magnitudes of all the losses are of similar order after scaling. Specifically, we set $\lambda_1 = 1, \lambda_2 = 0.1, \lambda_3 = 0.001$ and $\lambda_4 = 0.1$. For visibility prior estimation, we set $D = 64$ and $\gamma = 10$. Since we require \hat{T}_i' to be close to T_i' while using \hat{T}_i' to compute \mathcal{L}_{vip} , we impose \mathcal{L}_{vip} after 20k iterations. We train our models on a single NVIDIA RTX A4000 16GB GPU.

5.3 Results

We show the quantitative performance of ViP-NeRF and other competing models on RealEstate-10K and NeRF-LLFF datasets in Tabs. 1 and 2. Our model outperforms all the competing models, particularly in terms of the perceptual metric, LPIPS. ViP-NeRF even outperforms models such as DDP-NeRF and DietNeRF which involve pre-training on a large dataset. Fig. 3 shows qualitative comparisons on a scene from the RealEstate-10K dataset, where we observe significantly better synthesis by our model as compared to the competing models. We show more qualitative comparisons in Figs. 7 to 10 in the figure only pages at the end of this manuscript. In these samples, we find that ViP-NeRF removes most of the floater artifacts and successfully retains the shapes of objects.

In Fig. 4, we qualitatively compare the predictions of our model with different numbers of input views. We observe that ViP-NeRF




Figure 5: Estimated depth map on RealEstate-10K dataset with two input views. We find that ViP-NeRF is better in both frame synthesis and depth estimation compared to the competing models. For example, in the first row, the depth estimated by DDP-NeRF is smooth which may be leading to a loss of sharpness in synthesizing the shrubs. In contrast, ViP-NeRF predictions are sharper. For better visualization, we show inverse depth and normalize it to set the maximum value to unity.




Figure 6: Visualization of the visibility map predicted by ViP-NeRF. White indicates the regions of the ‘Primary View’ which are visible in the ‘Secondary View’ and black indicates the occluded regions. From the primary and secondary views, we observe that the left part of the fortress and the neighboring portion of the wood are hidden in the secondary view. ViP-NeRF is able to reasonably determine the visible and occluded regions.

estimates the geometry reasonably well with even two input views. However, with more input views, the performance of ViP-NeRF improves in reflective or specular regions. Fig. 6 visualizes the visibility map predicted by ViP-NeRF, where we observe that it is able to accurately predict the regions in the primary image which are visible and occluded in the secondary image.

Dense depth vs dense visibility. The key idea of our paper is that it may be possible to reliably estimate dense visibility than dense depth. From Tab. 1, we find that ViP-NeRF outperforms DDP-NeRF consistently, which indicates that the dense visibility prior we compute without any pre-training is superior to the learned dense depth prior used by DDP-NeRF. Further from Tab. 2, we observe that ViP-NeRF consistently improves over DS-NeRF in terms of LPIPS and SSIM, whereas DDP-NeRF does not. This may be due to the domain shift between the training dataset of DDP-NeRF and the LLFF dataset, resulting in no performance improvement over DS-NeRF. Thus, we conclude that augmenting sparse depth with dense visibility leads to better view synthesis performance than dense completion of the sparse depth. We further validate this conclusion by comparing the two priors in the following.

Validating priors. We compare the reliability of the dense visibility prior used in our model against the dense depth prior from DDP-NeRF. For this comparison, we convert the dense depth to visibility and compare it with the visibility prior of our approach. Specifically, we warp the image in the secondary view to the primary view using the dense depth prior and compute the visibility map similar to Eq. (10). We compare the visibility maps obtained using dense depth and our approach with the visibility map predicted by a NeRF model trained with dense input views. We evaluate the visibility maps in terms of precision, recall, and F1 score.

From Tab. 3, we observe that our approach significantly outperforms DDP-NeRF prior in terms of the recall and F1 score, while performing similarly in terms of precision. A high precision of our prior indicates that it makes very few mistakes when imposing \mathcal{L}_{vip} . On the other hand, a high recall shows that our prior is able to capture most of the visible regions where \mathcal{L}_{vip} needs to be imposed. On the contrary, a low recall for the DDP-NeRF prior indicates that large regions that are actually visible in the secondary view are marked as occluded by the dense depth prior. Consequently, this indicates the presence of a large number of pixels with inaccurate depth in the prior of DDP-NeRF. Thus, we conclude that our visibility prior is more reliable than the dense depth prior from DDP-NeRF for training the NeRF.

As discussed in Sec. 1, visibility is related to relative depth, and thus a prior on visibility only constrains the relative depth ordering of the objects. On the other hand, the dense depth prior constrains the absolute depth, perhaps incorrectly. Thus the visibility prior provides more freedom to the NeRF in reconstructing the 3D geometry and is also more reliable compared to the depth prior. This may explain the superior performance of visibility regularization over dense depth regularization.

Evaluation of estimated depth. It is believed that better performance in synthesizing novel views is directly correlated with the accuracy of depth estimation [Deng et al. 2022]. Thus, we compare our model with DDP-NeRF on their ability to estimate absolute depth correctly using root mean squared error (RMSE). We also evaluate the models on their ability to estimate the relative depth of the scene correctly using spearman rank-order correlation coefficient (SROCC) [Corder and Foreman 2014], which computes the linear correlation between ranks of the estimated pixel depths with that of the ground truth depth. Due to the unavailability of ground truth depth on both the datasets, we train a NeRF model with dense input views and use its predicted depth as a pseudo ground truth. From Tab. 4, we observe that our model consistently outperforms DDP-NeRF both in terms of absolute and relative depth. Fig. 5 shows that the depth estimated by DDP-NeRF is smooth in textured regions, which may be leading to blur in the synthesized frame. In contrast, the dense visibility prior used in our model allows NeRF to predict sharp depth in such regions leading to sharper frame predictions.

Ablations. We analyze the contributions of dense visibility and sparse depth priors in ViP-NeRF, by disabling them one at a time. From Tab. 5 and Fig. 11a, we find that removing either priors leads to a drop in performance on both the datasets. This suggests that the dense visibility prior may be providing information that is complementary to the sparse depth prior. For a more fine-grained analysis, we compare the LPIPS scores on individual scenes in

Fig. 11b. We observe that the addition of dense visibility prior over sparse depth prior leads to an improvement in the performance on all the scenes. Further, we find that our model with dense visibility prior alone is able to achieve impressive performance, especially on the RealEstate-10K dataset.

5.4 Limitations and Future Work

Our visibility prior constrains only the regions visible in at least two of the input views. As a result, we observe inaccurate depth estimation in the regions that are visible in only one of the input images. However, such regions account for a very small portion of the scene and reduce further with three or four input views. Further, ViP-NeRF may fail to synthesize disoccluded regions that can occur in sparse-input view synthesis, similar to RegNeRF. It would be interesting to explore the use of generative NeRF models such as pix2NeRF [Cai et al. 2022] to synthesize such disocclusions.

Our approach to estimating the visibility prior may not account for significant color changes that can occur when the scene contains highly specular surfaces. We do not impose any loss on such pixels. It would be interesting to analyze if pre-training a network on a large dataset to estimate visibility can provide more supervision in specular regions. Moreover, it would be interesting to see if pre-training a network to predict dense visibility generalizes better when compared to depth completion. Also, we observe in Tab. 2 that adding a new view leads to a significant improvement in performance as compared to adding new regularizations. Thus, one could explore hallucinating new views using generative models and use the hallucinated views for additional supervision.

6 CONCLUSION

We study the problem of training NeRFs in sparse input scenarios, where the NeRF tends to overfit the input views and learn incorrect geometry. We propose a prior on the visibility of pixels in other viewpoints to regularize the training and mitigate such errors. The visibility prior obtained using a plane sweep volume is more reliable as compared to the depth prior estimated using pre-trained networks. We reformulate the NeRF MLPs to additionally output visibility to compute the visibility prior loss in a time-efficient manner. ViP-NeRF achieves state-of-the-art performance on two commonly used datasets for novel view synthesis.

ACKNOWLEDGMENTS

This work was supported in part by a grant from Qualcomm. The first author was supported by the Prime Minister’s Research Fellowship awarded by the Ministry of Education, Government of India. The authors would also like to thank Suhas Srinath and Nithin Babu for the valuable discussions.

REFERENCES

Jonathan T. Barron, Ben Mildenhall, Matthew Tancik, Peter Hedman, Ricardo Martin-Brualla, and Pratul P. Srinivasan. 2021. Mip-NeRF: A Multiscale Representation for Anti-Aliasing Neural Radiance Fields. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*.

Jonathan T. Barron, Ben Mildenhall, Dor Verbin, Pratul P. Srinivasan, and Peter Hedman. 2022. Mip-NeRF 360: Unbounded Anti-Aliased Neural Radiance Fields. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*.

Shengqu Cai, Anton Obukhov, Dengxin Dai, and Luc Van Gool. 2022. Pix2NeRF: Unsupervised Conditional p-GAN for Single Image to Neural Radiance Fields

Translation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*.

Eric R. Chan, Marco Monteiro, Petr Kellnhofer, Jiajun Wu, and Gordon Wetzstein. 2021. Pi-GAN: Periodic Implicit Generative Adversarial Networks for 3D-Aware Image Synthesis. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*.

Anpei Chen, Zexiang Xu, Fuqiang Zhao, Xiaoshuai Zhang, Fanbo Xiang, Jingyi Yu, and Hao Su. [n. d.]. MVSNeRF: Fast Generalizable Radiance Field Reconstruction from Multi-View Stereo. *arXiv e-prints*, Article arXiv:2103.15595 (n. d.).

R.T. Collins. 1996. A Space-Sweep Approach to True Multi-Image Matching. In *Proceedings of the IEEE Computer Society Conference on Computer Vision and Pattern Recognition (CVPR)*. <https://doi.org/10.1109/CVPR.1996.517097>

Gregory W Corder and Dale I Foreman. 2014. *Nonparametric statistics: A step-by-step approach*.

Kangle Deng, Andrew Liu, Jun-Yan Zhu, and Deva Ramanan. 2022. Depth-Supervised NeRF: Fewer Views and Faster Training for Free. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*.

David Gallup, Jan-Michael Frahm, Philippos Mordohai, Qingxiang Yang, and Marc Pollefeys. 2007. Real-Time Plane-Sweeping Stereo with Multiple Sweeping Directions. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*. <https://doi.org/10.1109/CVPR.2007.383245>

Chen Gao, Yichang Shih, Wei-Sheng Lai, Chia-Kai Liang, and Jia-Bin Huang. [n. d.]. Portrait Neural Radiance Fields from a Single Image. *arXiv e-prints*, Article arXiv:2012.05903 (n. d.).

Hyowon Ha, Sunghoon Im, Jaesik Park, Hae-Gon Jeon, and In So Kweon. 2016. High-Quality Depth From Uncalibrated Small Motion Clip. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*.

Abdullah Hamdi, Bernard Ghanem, and Matthias Nießner. [n. d.]. SPARF: Large-Scale Learning of 3D Sparse Radiance Fields from Few Input Images. *arXiv e-prints*, Article arXiv:2212.09100 (n. d.).

Yuxuan Han, Ruicheng Wang, and Jiaolong Yang. 2022. Single-View View Synthesis in the Wild with Learned Adaptive Multiplane Images. In *Proceedings of the ACM SIGGRAPH*.

Po-Han Huang, Kevin Matzen, Johannes Kopf, Narendra Ahuja, and Jia-Bin Huang. 2018. DeepMVS: Learning Multi-View Stereopsis. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*.

Sunghoon Im, Hae-Gon Jeon, Stephen Lin, and In So Kweon. [n. d.]. DPSNet: End-to-end Deep Plane Sweep Stereo. *arXiv e-prints*, Article arXiv:1905.00538 (n. d.).

Ajay Jain, Matthew Tancik, and Pieter Abbeel. 2021. Putting NeRF on a Diet: Semantically Consistent Few-Shot View Synthesis. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*.

Mohammad Mahdi Johari, Yann Lepoittevin, and François Fleuret. 2022. GeoNeRF: Generalizing NeRF With Geometry Priors. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*.

Mijeong Kim, Seonguk Seo, and Bohyung Han. 2022. InfoNeRF: Ray Entropy Minimization for Few-Shot Neural Volume Rendering. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*.

Alex Krizhevsky, Ilya Sutskever, and Geoffrey E Hinton. 2012. ImageNet Classification with Deep Convolutional Neural Networks. In *Proceedings of the Advances in Neural Information Processing Systems (NeurIPS)*.

Jiaxin Li, Zijian Feng, Qi She, Henghui Ding, Changhu Wang, and Gim Hee Lee. 2021. MINE: Towards Continuous Depth MPI With NeRF for Novel View Synthesis. In *Proceedings of the IEEE International Conference on Computer Vision (ICCV)*.

Kai-En Lin, Yen-Chen Lin, Wei-Sheng Lai, Tsung-Yi Lin, Yi-Chang Shih, and Ravi Ramamoorthi. 2023. Vision Transformer for NeRF-Based View Synthesis From a Single Input Image. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision (WACV)*.

Yuan Liu, Sida Peng, Lingjie Liu, Qianqian Wang, Peng Wang, Christian Theobalt, Xiaowei Zhou, and Wenping Wang. 2022. Neural Rays for Occlusion-Aware Image-Based Rendering. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*.

Ben Mildenhall, Pratul P. Srinivasan, Rodrigo Ortiz-Cayon, Nima Khademi Kalantari, Ravi Ramamoorthi, Ren Ng, and Abhishek Kar. 2019. Local Light Field Fusion: Practical View Synthesis with Prescriptive Sampling Guidelines. *ACM Transactions on Graphics (TOG)* 38, 4 (July 2019), 1–14. <https://doi.org/10.1145/3306346.3322980>

Ben Mildenhall, Pratul P. Srinivasan, Matthew Tancik, Jonathan T. Barron, Ravi Ramamoorthi, and Ren Ng. 2020. NeRF: Representing Scenes as Neural Radiance Fields for View Synthesis. In *Proceedings of the European Conference on Computer Vision (ECCV)*.

Michael Niemeyer, Jonathan T. Barron, Ben Mildenhall, Mehdi S. M. Sajjadi, Andreas Geiger, and Noha Radwan. 2022. RegNeRF: Regularizing Neural Radiance Fields for View Synthesis From Sparse Inputs. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*.

Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal, Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, Gretchen Krueger, and Ilya Sutskever. 2021. Learning Transferable Visual Models From Natural Language Supervision. In *Proceedings of the International Conference on Machine Learning (ICML)*.

- Barbara Roessle, Jonathan T. Barron, Ben Mildenhall, Pratul P. Srinivasan, and Matthias Nießner. 2022. Dense Depth Priors for Neural Radiance Fields From Sparse Input Views. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*.
- Pratul P. Srinivasan, Richard Tucker, Jonathan T. Barron, Ravi Ramamoorthi, Ren Ng, and Noah Snavely. 2019. Pushing the Boundaries of View Extrapolation With Multiplane Images. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*.
- Matthew Tancik, Ben Mildenhall, Terrance Wang, Divi Schmidt, Pratul P. Srinivasan, Jonathan T. Barron, and Ren Ng. 2021. Learned Initializations for Optimizing Coordinate-Based Neural Representations. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*.
- Richard Tucker and Noah Snavely. 2020. Single-View View Synthesis With Multiplane Images. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*.
- Qianqian Wang, Zhicheng Wang, Kyle Genova, Pratul P. Srinivasan, Howard Zhou, Jonathan T. Barron, Ricardo Martin-Brualla, Noah Snavely, and Thomas Funkhouser. 2021. IBRNet: Learning Multi-View Image-Based Rendering. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*.
- Zhou Wang, Alan C Bovik, Hamid R Sheikh, and Eero P Simoncelli. 2004. Image quality assessment: from error visibility to structural similarity. *IEEE Transactions on Image Processing (TIP)* 13, 4 (2004), 600–612.
- Olivia Wiles, Georgia Gkioxari, Richard Szeliski, and Justin Johnson. 2020. SynSin: End-to-End View Synthesis From a Single Image. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*.
- Felix Wimbauer, Nan Yang, Christian Rupprecht, and Daniel Cremers. [n. d.]. Behind the Scenes: Density Fields for Single View Reconstruction. *arXiv e-prints*, Article arXiv:2301.07668 ([n. d.]).
- Jamie Wynn and Daniyar Turmukhambetov. [n. d.]. DiffusioNeRF: Regularizing Neural Radiance Fields with Denoising Diffusion Models. *arXiv e-prints*, Article arXiv:2302.12231 ([n. d.]).
- Dejia Xu, Yifan Jiang, Peihao Wang, Zhiwen Fan, Humphrey Shi, and Zhangyang Wang. 2022. SinNeRF: Training Neural Radiance Fields on Complex Scenes from a Single Image. In *Proceedings of the European Conference on Computer Vision (ECCV)*.
- Ruigang Yang and M. Pollefeys. 2003. Multi-Resolution Real-Time Stereo on Commodity Graphics Hardware. In *Proceedings of the IEEE Computer Society Conference on Computer Vision and Pattern Recognition (CVPR)*. <https://doi.org/10.1109/CVPR.2003.1211356>
- Zhenpei Yang, Zhile Ren, Miguel Angel Bautista, Zaiwei Zhang, Qi Shan, and Qixing Huang. 2022. FvOR: Robust Joint Shape and Pose Optimization for Few-View Object Reconstruction. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*.
- Alex Yu, Vickie Ye, Matthew Tancik, and Angjoo Kanazawa. 2021. pixelNeRF: Neural Radiance Fields From One or Few Images. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*.
- Jason Zhang, Gengshan Yang, Shubham Tulsiani, and Deva Ramanan. 2021. NeRS: Neural Reflectance Surfaces for Sparse-view 3D Reconstruction in the Wild. In *Proceedings of the Advances in Neural Information Processing Systems (NeurIPS)*.
- Richard Zhang, Phillip Isola, Alexei A Efros, Eli Shechtman, and Oliver Wang. 2018. The Unreasonable Effectiveness of Deep Features as a Perceptual Metric. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*.
- Tinghui Zhou, Richard Tucker, John Flynn, Graham Fyffe, and Noah Snavely. 2018. Stereo Magnification: Learning View Synthesis Using Multiplane Images. *ACM Transactions on Graphics (TOG)* 37, 4 (July 2018).
- Zhizhuo Zhou and Shubham Tulsiani. [n. d.]. SparseFusion: Distilling View-conditioned Diffusion for 3D Reconstruction. *arXiv e-prints*, Article arXiv:2212.00792 ([n. d.]).



Figure 7: Qualitative examples on RealEstate-10K dataset with two input views. We observe sharp predictions by ViP-NeRF while predictions by other models suffer from blur and other artifacts. In particular, DDP-NeRF predictions contain blurred flowers (first row) and blurred tiles (second row).

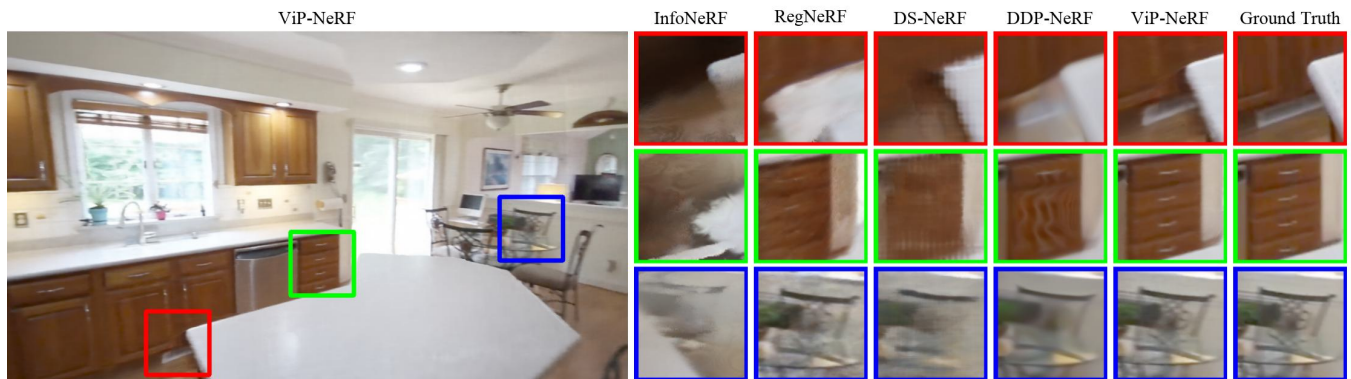


Figure 8: Qualitative examples on RealEstate-10K dataset with three input views. We find that ViP-NeRF is able to reconstruct novel views significantly better than the competing models. DDP-NeRF extends parts of the white table and fails to reconstruct the drawer handles accurately in the first and second examples. In the third example, DDP-NeRF fails to reconstruct thin objects in the chair.

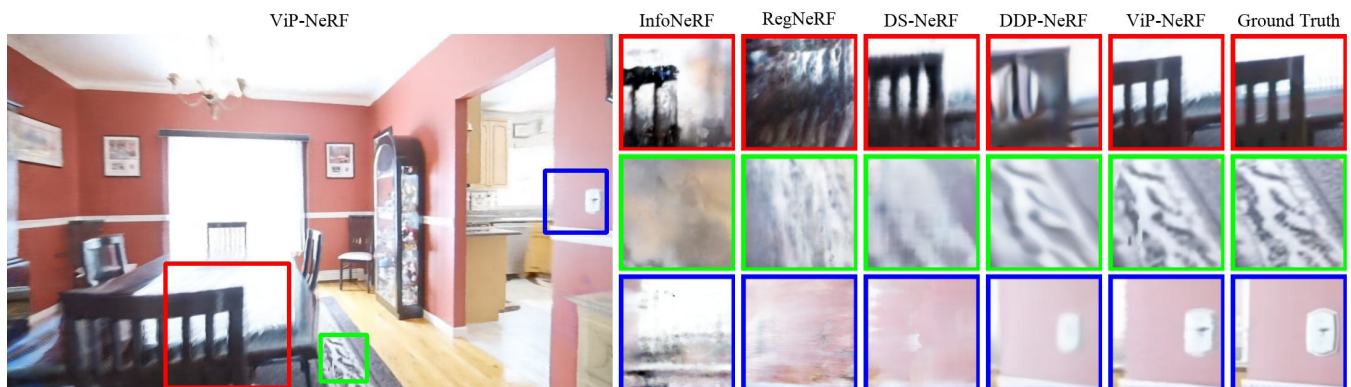


Figure 9: Qualitative examples on RealEstate-10K dataset with four input views. In the first example, DDP-NeRF fails to retain the structure of the chair while it blurs the texture of the carpet in the second example. We observe even more severe distortions among the predictions of other models.

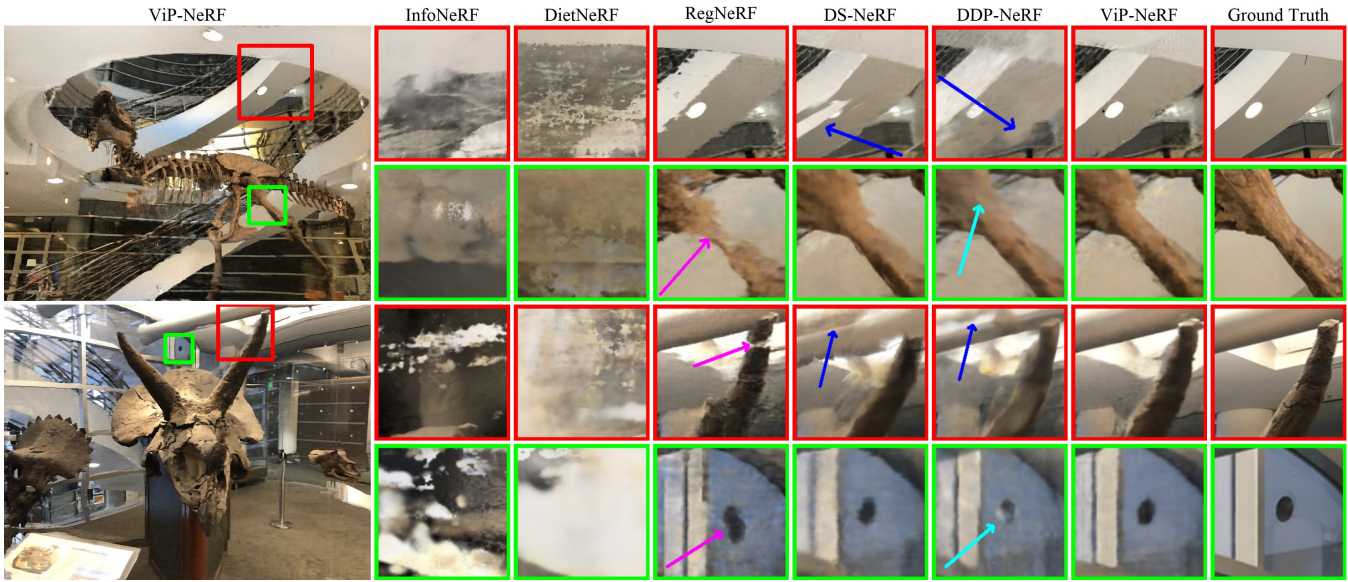




Figure 10: Qualitative examples on NeRF-LLFF dataset with two input views. In the first and third examples, we observe floater artifacts (blue arrows) in the predictions of DS-NeRF and DDP-NeRF, which are mitigated in the predictions of ViP-NeRF. We find that RegNeRF fails to capture thin t-rex bone in the second example and breaks the horn into two pieces in the third example (magenta arrows). Cyan arrows indicate color changes in the predictions of DDP-NeRF in the second and fourth examples. We note that predictions by our model do not suffer from the above described artifacts.



(a) Qualitative examples for ablations on RealEstate-10K and NeRF-LLFF dataset. We observe that the absence of dense visibility prior leads to significant blur in the predicted frames. While the reconstruction is reasonable without the sparse depth prior, we obtain the best reconstructions when using both the priors.



(b) Scene-wise LPIPS scores of ViP-NeRF and the ablated models. Note that lower LPIPS scores are better. ViP-NeRF performs better than both the ablated models in most cases leading to overall better performance.

Figure 11: Qualitative and quantitative comparisons of ablated models on both RealEstate-10K and NeRF-LLFF datasets.