CaesarNeRF: Calibrated Semantic Representation for Few-Shot Generalizable Neural Rendering

Haidong Zhu^{1,*}

Tianyu Ding 2,*,† Ram Nevatia¹ L

Tianyi Chen² Luming Liang^{2,†} Ilya Zharkov²

¹University of Southern California ²Microsoft {haidongz,nevatia}@usc.edu {tianyuding,tiachen,zharkov,lulian}@microsoft.com



Fig. 1: Novel view synthesis for novel scenes using **ONE** reference view on Shiny [67], LLFF [42], and MVImgNet [77] (top to bottom). Each triplet of images corresponds to the results from GNT [60] (left), CaesarNeRF (middle) and groundtruth (right).

Abstract. Generalizability and few-shot learning are key challenges in Neural Radiance Fields (NeRF), often due to the lack of a holistic understanding in pixel-level rendering. We introduce CaesarNeRF, an end-to-end approach that leverages scene-level **CA**libratEd SemAntic Representation along with pixel-level representations to advance few-shot, generalizable neural rendering, facilitating a holistic understanding without compromising high-quality details. CaesarNeRF explicitly models pose differences of reference views to combine scene-level semantic representations, providing a calibrated holistic understanding. This calibration process aligns various viewpoints with precise location and is further enhanced by sequential refinement to capture varying details. Extensive experiments on public datasets, including LLFF, Shiny, mip-NeRF 360, and MVImgNet, show that CaesarNeRF delivers state-of-the-art performance across varying numbers of reference views, proving effective even with a single reference image.

Keywords: Few-view rendering · Generalizable NeRF · Neural rendering

*Equal contribution. [†]Corresponding author.

This work was done when Haidong Zhu was an intern at Microsoft.

1 Introduction

Rendering a scene from a novel camera position is essential in view synthesis [5,11,62]. The recent advancement of Neural Radiance Field (NeRF) [44] has shown impressive results in creating photo-realistic images from novel viewpoints. However, conventional NeRF methods are either typically scene-specific, necessitating retraining for novel scenes [15, 16, 44, 67, 74], or require a large number of reference views as input for generalizing to novel scenarios [6, 58, 60, 76]. These constraints highlight the complexity of the few-shot generalizable neural rendering, which aims to render unseen scenes from novel viewpoints with a limited number of reference images.

Generalizing NeRF to novel scenes often involves using pixel-level feature embeddings encoded from input images, as seen in existing methods [59, 76]. These methods adapt NeRF to novel scenes by separating the scene representation from the model through an image encoder. However, relying solely on pixel-level features has its drawbacks: it requires highly precise epipolar geometry and often overlooks occlusion in complex scenes. Moreover, employing pixel-level features ignores the inherent interconnections within objects in the scene, treating the prediction of each pixel independently. Prior attempts to utilize scene-level representation either suffers from style mismatches [35] during scene-wide rendering or have been limited to specific object categories [23, 39, 69]. With few input reference under generalizable settings, these limitations become more pronounced, exacerbating the ambiguity in predictions due to biases from camera viewpoints. Although recent diffusion-based models [18, 56] attempt to address these issues through generative approaches, they struggle to effectively use the input views as contextual references for specific scenes.

We present CaesarNeRF, a method that advances the generalizability of NeRF by incorporating calibrated semantic representation. This enables rendering from novel viewpoints using as few as one input reference view, as depicted in Figure 1. Our approach combines semantic scene-level representation with per-pixel features, enhancing consistency across different views of the same scene. The encoder-generated scenelevel representations capture both semantic features and biases linked to specific camera poses. When reference views are limited, these biases can introduce uncertainty in the rendered images. To counter this, CaesarNeRF integrates camera pose transformations into the semantic representation, hence the term calibrated. By isolating pose-specific information from the scene-level representation, our model harmonizes features across input views, mitigating view-specific biases and, in turn, reducing ambiguity. In addition, CaesarNeRF introduces a sequential refinement process, which equips the model with varying levels of detail needed to enhance the semantic features. Extensive experiments on datasets such as LLFF [42], Shiny [67], mip-NeRF 360 [4], and the newly released MVImgNet [77] demonstrate that CaesarNeRF outperforms current state-ofthe-art methods with limited reference views available, proving effective in generalizable settings with as few as one reference view. The project can be found here.

In summary, our contributions are as follows:

 We introduce CaesarNeRF, which utilizes scene-level calibrated semantic representation to achieve few-shot, generalizable neural rendering. This innovation leads to coherent and high-quality renderings. Calibrated Semantic Representation for Few-Shot Generalizable Neural Rendering

- We integrate semantic scene context with pixel-level details, in contrast to existing
 methods that rely solely on pixel-level features. We also address view-specific biases by modeling camera pose transformations and enhance the scene understanding through the sequential refinement of semantic features.
- We demonstrate through extensive experiments that CaesarNeRF consistently outperforms state-of-the-art generalizable NeRF methods across a variety of datasets.
 Furthermore, integrating the Caesar pipeline into other baseline methods leads to consistent performance gains, highlighting its effectiveness and adaptability.

2 Related work

Neural Radiance Field (NeRF) implicitly captures the density and appearance of points within a scene or object [43, 44] and enables rendering from novel camera positions. In recent years, NeRF has witnessed improvements in a wide range of applications, such as photo-realistic novel view synthesis for large-scale scenes [40,68,80], dynamic scene decomposition and deformation [24,30,34,46–48,50,84,85], occupancy or depth estimation [66,71,83], scene generation and editing [1,21,31,33,41,49,72,73,79], and so on. Despite these advances, most methods still rely on the original NeRF and require retraining or fine-tuning for novel scenes not covered in the training data.

Generalizable NeRF aims to adapt a single NeRF model to multiple scenes by separating the scene representation from the model. This field has seen notable advancements, with efforts focused on avoiding the need for retraining [8,20,38,64,65,75,76]. PixelNeRF [76] and GRF [59] pioneered the application of an image encoder to transform images into per-pixel features, with NeRF functioning as a decoder for predicting density and color from these features. MVSNeRF [6] introduces the use of a cost volume from MVSNet [75] to encode 3-D features from multiple views. Recognizing the intrinsic connection between points along a ray, IBRNet [64] employs self-attention to enhance point density predictions. Transformer-based [61] networks like GNT [10,60], GeoNeRF [25], and GPNR [58] are explored as alternatives to volume rendering, concentrating on pixel and patch-level representations. Additionally, InsertNeRF [2] utilizes hypernet modules to adapt parameters for novel scenes efficiently.

These methods primarily depend on image encoders to extract pixel-aligned features from reference views. As a result, many of them lack a comprehensive understanding of the entire scene. Furthermore, with few reference views, the features become intertwined with view-specific details, compromising the quality of the rendering results.

Few-shot Neural Radiance Field aims to render novel views using a limited number of reference images. To this end, various methods have been developed, incorporating additional information such as normalization-flow [78], semantic constraints [17, 22], depth cues [13, 52], geometry consistency [3, 29, 45, 63, 70], and frequency content [74]. Others [6, 9] emphasize pretraining on large-scale datasets.

While these methods offer reasonable reconstructions with few inputs, they typically still require training or fine-tuning for specific scenes. Moreover, these methods usually require at least three reference images. With fewer than three, view-specific biases lead to ambiguity, complicating the rendering. Diffusion-based [12, 35, 57, 82] and other generative methods [27, 86] have been explored for single-view synthesis or gen-

erative rendering, yet they are mostly limited to single-object rendering and generally fall short for complex scenes, which often result in a style change.

CaesarNeRF confronts the above challenges by leveraging calibrated semantic representations that exploit scene geometry and variations in camera viewpoints. As a result, CaesarNeRF overcomes the limitations of pixel-level features and reduces dependency on external data or extensive pretraining, delivering high-quality renderings in few-shot and generalizable settings.

3 The proposed method

We first outline the general framework of existing generalizable NeRF in Section 3.1. Then, we present our proposed CaesarNeRF, as illustrated in Figure 2. This model integrates elements of semantic representation, calibration, and sequential refinement, detailed in Section 3.2, 3.3, and 3.4, respectively. The training objective is given in 3.5.

3.1 NeRF and generalizable NeRF

Neural Radiance Field (NeRF) [43, 44] aims to render 3D scenes by predicting both the density and RGB values at points where light rays intersect the radiance field. For a query point $x \in \mathbb{R}^3$ and a viewing direction d on the unit sphere \mathbb{S}^2 in 3D space, the NeRF model \mathcal{F} is defined as:

$$\sigma, \boldsymbol{c} = \boldsymbol{\mathcal{F}}(\boldsymbol{x}, \boldsymbol{d}). \tag{1}$$

Here, $\sigma \in \mathbb{R}$ and $c \in \mathbb{R}^3$ denote the density and the RGB values, respectively. After computing these values for a collection of discretized points along each ray, volume rendering techniques are employed to calculate the final RGB values for each pixel, thus reconstructing the image.

However, traditional NeRF models \mathcal{F} are limited by their requirement for scenespecific training, making it unsuitable for generalizing to novel scenes. To overcome this, generalizable NeRF models, denoted by \mathcal{F}_G , are designed to render images of novel scenes without per-scene training. Given N reference images $\{I_n\}_{n=1}^N$, an encoderbased generalizable NeRF model \mathcal{F}_G decouples the object representation from the original NeRF by using an encoder to extract per-pixel feature maps $\{F_n\}_{n=1}^N$ from the input images. To synthesize a pixel associated with a point x along a ray in direction d, it projects $\{F_n\}_{n=1}^N$ from nearby views and aggregates this multi-view pixel-level information using techniques such as average pooling [76] or cost volumes [6]. This results in a fused feature embedding \tilde{F} , allowing \mathcal{F}_G to predict density σ and RGB values c for each point along the ray, as expressed by:

$$\sigma, \boldsymbol{c} = \boldsymbol{\mathcal{F}}_{\boldsymbol{G}}(\boldsymbol{x}, \boldsymbol{d}, \boldsymbol{F}). \tag{2}$$

In our method, we adopt the recently introduced fully attention-based generalizable NeRF method, GNT [60], as both the backbone and the baseline. GNT shares a similar paradigm with (2) but employs transformers [61] to aggregate pixel-level features into \tilde{F} . It uses a *view transformer* to fuse projected pixel-level features from reference



Fig. 2: Overview of CaesarNeRF. CaesarNeRF employs a shared encoder to capture two types of features from input views, including scene-level semantic representation $\{S_n\}$ and pixel-level feature representation $\{F_n\}$. We use the same encoder for both the scene-level semantic representation and the pixel-level embeddings. Following calibration and aggregation of $\{S_n\}$ from various views, we concatenate it with the pixel-level fused feature, processed by the view transformer. Subsequent use of the ray-transformer, coupled with sequential refinement, enables us to render the final RGB values for each pixel in the target view. The output features serve as the input for the next stage, indicated by matching line colors.

views, and a *ray transformer* to combine features from different points along a ray, eliminating the need for volume rendering. Further details about GNT can be found in [60]. We also demonstrate that our approach can be extended to other generalizable NeRF models, as discusses in Section 4.2.

3.2 Scene-level semantic representation

Both encoder-based generalizable NeRF models [6, 59, 76] and their attention-based counterparts [58,60] mainly rely on pixel-level feature representations. While effective, this approach restricts their capability for a holistic scene understanding, especially when reference views are scarce. This limitation also exacerbates challenges in resolving depth ambiguities between points along the rays, a problem that becomes more pronounced with fewer reference views.

To address these challenges, we introduce semantic representations aimed at enriching the scene-level understanding. We utilize a shared CNN encoder and apply a Global Average Pooling (GAP) to its *C*-dimensional output feature map, generating *N* global feature vectors $\{S_n\}_{n=1}^N$ corresponding to each input view. These feature vectors are then averaged to form a unified scene-level representation S, *i.e.*,

$$\boldsymbol{S} = \frac{1}{N} \sum_{n=1}^{N} \boldsymbol{S}_n \in \mathbb{R}^C.$$
(3)

In GNT [60], which uses a view transformer to aggregate pixel-level features into an *L*-dimensional vector \widetilde{F} , we extend this by concatenating \widetilde{F} with S to construct a



Fig. 3: An illustration of conflicting semantic meanings from multiple viewpoints of the same object. When observing the cup from distinct angles, the features extracted after pooling retain spatial information but are inconsistent in the scene-level semantic understanding, leading to conflicts across various reference images after aggregation.

global-local embedding E, as formulated by:

$$\boldsymbol{E} = \operatorname{Concat}(\widetilde{\boldsymbol{F}}, \boldsymbol{S}) \in \mathbb{R}^{L+C}.$$
(4)

This combined embedding E is then subjected to the standard self-attention mechanism used in GNT [60]. This approach enables the scene-level semantic representation (S)to integrate with per-point features (\tilde{F}) , offering a more nuanced understanding at both levels. It also allows each point to selectively draw from the scene-level information. To maintain dimensional consistency across the input and output layers of multiple transformer modules, we employ a two-layer MLP to project the enhanced features back to the original dimension L of the per-point embedding \tilde{F} .

3.3 Calibration of semantic representation

The integration of the scene-level semantic representation S, generated through simple averaging of global feature vectors as in (3), improves rendering quality. However, this approach has limitations when dealing with multiple views. As illustrated in Figure 3, viewing the same object from distinct angles may retain spatial attributes but can lead to conflicting semantic meanings. Merely averaging these global feature vectors without accounting for camera positions can result in a distorted scene-level understanding.

To mitigate this inconsistency, we propose a semantic calibration technique using feature rotation. This adjustment aligns the semantic representation across different camera poses. Our inspiration comes from the use of camera pose projection in computing the fused pixel-level feature \tilde{F} and is further motivated by [51], which demonstrates that explicit rotation operations in feature spaces are feasible. Unlike point clouds in [51] that inherently lack a defined canonical orientation, NeRF explicitly encodes differences between camera viewpoints, thereby enabling precise calibration between the reference and target images.

Building on this observation, we calculate calibrated semantic representations $\{S_n\}_{n=1}^N$ from the N original semantic representations $\{S_n\}_{n=1}^N$ derived from the reference views. We accomplish this by leveraging their respective rotation matrices $\{T_n\}_{n=1}^N$



Fig. 4: Visualization of decoded feature maps for "orchid" in LLFF dataset, produced by *ray transformers* [60] at different stages. From left to right, the transformer stages increase in depth.

to model the rotational variations between each input view and the target view. The alignment of the original semantic features is performed as follows:

$$\widetilde{\boldsymbol{S}}_n = \boldsymbol{\mathcal{P}}(\boldsymbol{T}_n \cdot \boldsymbol{\mathcal{P}}^{-1}(\boldsymbol{S}_n)), \text{ where } \boldsymbol{T}_n = \boldsymbol{T}_{\text{out}}^{\text{w2c}} \cdot \boldsymbol{T}_n^{\text{c2w}}.$$
 (5)

Here, T_n^{c2w} is the inverse of the extrinsic matrix used for I_n , and T_{out}^{w2c} is the extrinsic matrix for the target view. $\mathcal{P}(\cdot)$ and $\mathcal{P}^{-1}(\cdot)$ are the flattening and inverse flattening operations, which reshape the feature to a 1D vector of shape 1-by-C and a 2D matrix of shape 3-by- $\frac{C}{3}$, respectively.

Note that for the extrinsic matrix, we consider only the top-left 3×3 submatrix that accounts for rotation. Using GAP to condense feature maps of various sizes into a 1-by-C feature vector eliminates the need for scaling parameters in the semantic representation. As a result, modeling the intrinsic matrix is unnecessary, assuming no skewing, making our approach adaptable to different camera configurations.

With the calibrated semantic features $\{\hat{S}_n\}_{n=1}^N$ for each reference view, we average these, similar to (3), to obtain the calibrated scene-level semantic representation \tilde{S} , *i.e.*,

$$\widetilde{\boldsymbol{S}} = \frac{1}{N} \sum_{n=1}^{N} \widetilde{\boldsymbol{S}}_n \in \mathbb{R}^C.$$
(6)

Finally, akin to (4), we concatenate the pixel-level fused feature \tilde{F} with the calibrated scene-level semantic representation \tilde{S} to form the final global-local embedding \tilde{E} :

$$\widetilde{\boldsymbol{E}} = \operatorname{Concat}(\widetilde{\boldsymbol{F}}, \widetilde{\boldsymbol{S}}) \in \mathbb{R}^{L+C}.$$
 (7)

This unified embedding then feeds into ray transformers, passing through standard selfattention mechanisms. In the original GNT [60], multiple view transformers and ray transformers are stacked alternately for sequential feature processing. The last ray transformer integrates features from multiple points along a ray to yield the final RGB value. We denote the corresponding feature representations at stage k as $\tilde{F}^{(k)}$ and $\tilde{E}^{(k)}$. Notably, the calibrated semantic representation \tilde{S} remains constant across these stages.

3.4 Sequential refinement

While leveraging \hat{S} improves consistency, a single, uniform \hat{S} may not be adequate for deeper layers that demand more nuanced details. In fact, we find that deeper transform-

ers capture finer details compared to shallower ones, as shown in Figure 4. To address this limitation, we introduce a sequential semantic feature refinement module that progressively enriches features at each stage. Specifically, we learn the residual $\Delta^{(k)}$ to update \tilde{S} at each stage k as follows:

$$\widetilde{\boldsymbol{S}}^{(k+1)} \leftarrow \widetilde{\boldsymbol{S}}^{(k)} + \boldsymbol{\Delta}^{(k)}.$$
 (8)

Here, $\Delta^{(k)}$ is calculated by first performing specialized cross-attentions between $\widetilde{S}^{(k)}$ and the original, uncalibrated per-frame semantic features $\{S_n\}_{n=1}^N$ (see Figure 2), followed by their summation. Our goal is to fuse information from different source views to enrich the scene-level semantic representation with features from each reference frame. With this sequential refinement, we combine $\widetilde{S}^{(k)}$ with $\widetilde{F}^{(k)}$ at each stage, yielding a stage-specific global-local embedding $\widetilde{E}^{(k)}$, which completes our approach.

Discussion. In scenarios with few reference views, especially when limited to just one, the primary issue is inaccurate depth estimation, resulting in depth ambiguity [12]. This compromises the quality of images when rendered from novel viewpoints. Despite this, essential visual information generally remains accurate across different camera poses. Incorporating our proposed scene-level representation improves the understanding of the overall scene layout [7], distinguishing our approach from existing generalizable NeRF models that predict pixels individually. The advantage of our approach is its holistic view; the semantic representation enriches per-pixel predictions by providing broader context. This semantic constraint ensures that fewer abrupt changes between adjacent points. Consequently, it leads to more reliable depth estimations, making the images rendered from limited reference views more plausible.

3.5 Training objectives

During training, we employ three different loss functions:

MSE loss. The Mean Square Error (MSE) loss is the standard photometric loss used in NeRF [43]. It computes the MSE between the actual and predicted pixel values.

Central loss. Since we project the view-level features from different input camera poses to the shared target view, we introduce a central loss to ensure frame-wise calibrated semantic features $\{\tilde{S}_n\}_{n=1}^N$ are consistent when projected onto the same target view, which is defined as:

$$\mathcal{L}_{\text{central}} = \frac{1}{N} \sum_{n=1}^{N} \left\| \widetilde{\boldsymbol{S}}_n - \widetilde{\boldsymbol{S}} \right\|_1.$$
(9)

Point-wise perceptual loss. During the rendering of a batch of pixels in a target view, we inpaint the ground-truth image by replacing the corresponding pixels with the predicted ones. Then, a perceptual loss [26] is computed between the inpainted image and the target image to guide the training process at the whole-image level.

The final loss function is formulated as follows:

$$\mathcal{L} = \mathcal{L}_{\text{MSE}} + \lambda_1 \mathcal{L}_{\text{central}} + \lambda_2 \mathcal{L}_{\text{perc}}.$$
 (10)

Empirically, we set $\lambda_1 = 1$ and $\lambda_2 = 0.001$, following [32].

Table 1: Results for generalizable scene rendering on LLFF with few reference views. GeoNeRF, MatchNeRF, and MVSNeRF necessitate variance as input, defaulting to 0 for single-image cases, hence their results are not included for 1-view scenarios.

Method	1 reference view			2 reference views			3 reference views		
Wiethou	PSNR (†)	LPIPS (\downarrow)	SSIM (†)	PSNR (†)	LPIPS (\downarrow)	SSIM (†)	PSNR (†)	LPIPS (\downarrow)	SSIM (†)
PixelNeRF [76]	9.32	0.898	0.264	11.23	0.766	0.282	11.24	0.671	0.486
GPNR [58]	15.91	0.527	0.400	18.79	0.380	0.575	21.57	0.288	0.695
NeuRay [38]	16.18	0.584	0.393	17.71	0.336	0.646	18.26	0.310	0.672
GeoNeRF [25]	-	-	-	18.76	0.473	0.500	23.40	0.246	0.766
MatchNeRF [8]	-	-	-	21.08	0.272	0.689	22.30	0.234	0.731
MVSNeRF [6]	-	-	-	19.15	0.336	0.704	19.84	0.314	0.729
IBRNet [64]	16.85	0.542	0.507	21.25	0.333	0.685	23.00	0.262	0.752
GNT [60]	16.57	0.500	0.424	20.88	0.251	0.691	23.21	0.178	0.782
Ours	18.31	0.435	0.521	21.94	0.224	0.736	23.45	0.176	0.794

4 Experiments

In this section, we first discuss our experimental setups and other details required for our experiments in Section 4.1, followed by our results for different experimental settings, along with ablation studies and analysis in Section 4.2.

4.1 Experimental setups

Datasets. Firstly, following [60], we construct our training data from both synthetic and real data. This collection includes scanned models from Google Scanned Objects [14], RealEstate10K [81], and handheld phone captures [64]. For evaluation, we utilize real data encompassing complex scenes from sources such as LLFF [42], Shiny [67], and mip-NeRF 360 [4]. Additionally, we train and test our model using the recently released MVImgNet dataset [77]. We adhere to the official split, focusing on examples from the *containers* category, and select 2,500 scenes for training. During inference, we choose 100 scenes, using their first images as target views and the spatially nearest images as references. Since MVImgNet does not provide camera poses, we utilize COLMAP [54, 55] to deduce the camera positions within these scenes.

Implementation details. CaesarNeRF is built upon GNT [60], for which we maintain the same configuration, setting the *ray* and *view* transformers stack number (K) to 8 for generalizable setting and 4 for single-scene setting. The feature encoder extracts bottleneck features, applies GAP, and then uses a fully connected (FC) layer to reduce the input dimension C to 96. Training involves 500,000 iterations using the Adam optimizer [28], with learning rate set at 0.001 for the feature encoder and 0.0005 for CaesarNeRF, halving them every 100,000 iterations. Each iteration samples 4,096 rays from a single scene. In line with [60], we randomly choose between 8 to 10 reference views for training, and 3 to 7 views when using the MVImgNet [77].

Baseline methods. We compare CaesarNeRF with several state-of-the-art methods suited for generalizable NeRF applications, including earlier works such as MVS-NeRF [6], PixelNeRF [76], and IBRNet [64], alongside more recent ones, including GPNR [58], NeuRay [38], GNT [60], GeoNeRF [25] and MatchNeRF [8].

Table 2: Results for generalizable scene rendering on Shiny with few reference views.

Method	1 reference view			2:	reference vie	ws	3 reference views		
Wethod	PSNR (†)	LPIPS (\downarrow)	SSIM (†)	PSNR (†)	LPIPS (\downarrow)	SSIM (†)	PSNR (†)	LPIPS (\downarrow)	SSIM (†)
MatchNeRF [8]	-	-	-	20.28	0.278	0.636	20.77	0.249	0.672
MVSNeRF [6]	-	-	-	17.25	0.416	0.577	18.55	0.343	0.645
IBRNet [64]	14.93	0.625	0.401	18.40	0.400	0.595	21.96	0.281	0.710
GNT [60]	15.99	0.548	0.400	20.42	0.327	0.617	22.47	0.247	0.720
Ours	17.57	0.467	0.472	21.47	0.293	0.652	22.74	0.241	0.723

Table 3: Results for generalizable scene rendering on mip-NeRF 360 with few reference views.

Method	1 reference view			2 1	reference vie	ws	3 reference views		
	PSNR (†)	LPIPS (\downarrow)	SSIM (†)	PSNR (†)	LPIPS (\downarrow)	SSIM (†)	PSNR (†)	LPIPS (\downarrow)	SSIM (†)
MatchNeRF [8]	-	-	-	17.00	0.566	0.392	17.26	0.551	0.407
MVSNeRF [6]	-	-	-	14.23	0.681	0.366	14.29	0.674	0.406
IBRNet [64]	14.12	0.682	0.283	16.24	0.618	0.360	17.70	0.555	0.420
GNT [60]	13.48	0.630	0.314	15.21	0.559	0.370	15.59	0.538	0.395
Ours	15.20	0.592	0.350	17.05	0.538	0.403	17.55	0.512	0.430

4.2 Results and analysis

We compare results in two settings: a generalizable setting, where the model is trained on multiple scenes without fine-tuning during inference for both few and all reference view cases, and a single-scene setting where the model is trained and evaluated on just one scene. Following these comparisons, we conduct ablation studies and test the generalizability of our method with other state-of-the-art approaches.

Generalizable rendering. In the generalizable setting, we adopt two training strategies. First, we train the model on multiple datasets as described in Section 4.1 and evaluate on LLFF [42], Shiny [67] and mip-NeRF 360 [4] datasets. In addition, the model is trained and tested on the MVImgNet [77] for object-centric generalizability.

(a) LLFF, Shiny, and mip-NeRF 360. The results for few-reference view scenarios on these datasets are shown in Tables 1, 2 and 3, respectively. Methods like MatchN-eRF [8], MVSNeRF [75], and GeoNeRF [25] require at least two reference views. On the LLFF dataset, all methods experience a performance decline as the number of views decreases. CaesarNeRF, however, consistently outperforms others across varying reference view numbers, with the performance gap becoming more significant with fewer views. For example, with 3 views, while IBRNet [64] and GNT [60] have comparable PSNRs, CaesarNeRF demonstrates a more substantial lead in LPIPS and SSIM metrics.

Similar patterns are observed on the Shiny [67] and mip-NeRF 360 [4] datasets. We apply the highest-performing methods from the LLFF evaluations and report the results

Table 4: Results on MVImgNet across varying numbers of reference views. 'C.' represents the use of calibration before averaging.

Mathod	1 reference view		2 reference views		3 reference views			4 reference views			5 reference views				
weniou	PSNR (†)	LPIPS (\downarrow)	SSIM (†)	PSNR (†)	LPIPS (↓)	SSIM (†)	PSNR (†)	LPIPS (\downarrow)	SSIM (†)	PSNR (†)	LPIPS (\downarrow)	SSIM (†)	PSNR (†)	LPIPS (\downarrow)	SSIM (†)
IBRNet	19.14	0.458	0.595	24.38	0.266	0.818	25.53	0.203	0.858	25.99	0.190	0.867	26.12	0.188	0.867
GNT	22.22	0.433	0.678	26.94	0.236	0.850	27.41	0.206	0.870	27.51	0.197	0.875	27.51	0.194	0.876
Ours w/o C.	23.61	0.371	0.718	26.34	0.274	0.817	27.10	0.228	0.850	27.30	0.210	0.862	27.34	0.203	0.865
Ours	24.28	0.334	0.747	27.34	0.215	0.856	27.82	0.190	0.875	27.92	0.181	0.881	27.92	0.179	0.882

Table 5: Results of per-scene optimization on LLFF, in comparison with state-of-the-art methods.

Method	LLFF [42]	NeRF [44]	NeX [67]	GNT [60]	Ours
PSNR (†)	23.27	26.50	27.26	27.24	27.64
LPIPS (\downarrow)	0.212	0.250	0.179	0.087	0.081
SSIM (†)	0.798	0.811	0.904	0.889	0.904

Table 6: Results on LLFF for few-shot generalization after adapting Caesar to other baselines.

Method	1 reference view			2 1	eference vie	ws	3 reference views		
	PSNR (†)	SSIM (†)	LPIPS (\downarrow)	PSNR (†)	SSIM (†)	LPIPS (\downarrow)	PSNR (†)	SSIM (†)	LPIPS (\downarrow)
MatchNeRF [8] + Caesar	-	-	-	20.59 21.55	0.775 0.782	0.276 0.268	22.43 22.98	0.805 0.824	0.244 0.242
IBRNet [64] + Caesar	16.85 17.76	0.507 0.543	0.542 0.500	21.25 22.39	0.685 0.740	0.333 0.275	23.00 23.67	0.752 0.772	0.262 0.242

for those that produce satisfactory outcomes with few reference views. CaesarNeRF maintains superior performance throughout. Notably, for complex datasets like mip-NeRF 360 [4], which have sparse camera inputs, the quality of rendered images generally decreases with fewer available reference views. Nonetheless, CaesarNeRF shows the most robust performance compared to the other methods.

(b) MVImgNet. We extend our comparison of CaesarNeRF with GNT [60] and IBRNet [64] on the MVImgNet dataset, focusing on object-centric scenes, as shown in Table 4. We examine a variant of CaesarNeRF where semantic calibration is substituted with simple feature averaging from multiple frames. While the performance of all methods improves with more views, CaesarNeRF consistently outperforms GNT and IBRNet. Notably, CaesarNeRF with feature averaging surpasses GNT in 1-view case but lags with additional views, implying that the absence of calibration lead to ambiguities when rendering from multiple views.

Per-scene optimization. Beyond the multi-scene generalizable setting, we demonstrate per-scene optimization results in Table 5. We calculate the average performance over 8 categories from the LLFF dataset [42]. CaesarNeRF consistently outperforms

	Model Variations		PSNR (†)	I PIPS (1)	SSIM (†)	
$oldsymbol{R}$ len.	Seq.	Cali.	15000 (1)	EI II 5 (_{\p})	551M ()	
	(Baseline GNT)		20.93	0.185	0.731	
Ext.			20.85	0.173	0.735	
+32			21.43	0.152	0.763	
+64			21.49	0.149	0.766	
+96			21.46	0.150	0.766	
+128			21.49	0.147	0.763	
+ 96	1		21.53	0.146	0.770	
+ 96		1	21.51	0.147	0.769	
+ 96	1	1	21.67	0.139	0.781	

Table 7: Ablations on the semantic representation length R, sequential refinement (Seq.) and calibration (Cali.). 'Ext.' denotes the extension of per-pixel feature to a length of 64 in GNT.



(b) Using two images as reference view.

Fig. 5: Comparative visualization of our proposed method against other state-of-the-art methods.

nearly all state-of-the-art methods in the comparison, across all three metrics, showing a significant improvement over our baseline method, GNT [60].

Adaptability. To test the adaptability of our Caesar pipeline, we apply it to two other state-of-the-art methods that use *view transformers*, namely MatchNeRF [8] and IBRNet [64]. We demonstrate in Table 6 that our enhancements in scene-level semantic understanding significantly boost the performance of these methods across all metrics. This indicates that the Caesar framework is not only beneficial in our CaesarNeRF, which is based on GNT [60], but can also be a versatile addition to other NeRF pipelines with view transformers to aggregate different input views.

Ablation analysis. We conduct ablation studies on the "orchid" scene from the LLFF dataset, with findings detailed in Table 7. Testing variations in representation and the impact of the sequential refinement and calibration modules, we find that increasing the latent size in GNT yields marginal benefits. However, incorporating even a modest semantic representation size distinctly improves results. The length of the semantic representation has a minimal impact on quality. Our ablation studies indicate that while sequential refinement and calibration each offer slight performance gains, their combined effect is most significant. In a single-scene context, semantic information is effectively embedded within the representation, making the benefits of individual modules subtler. Together, however, they provide a framework where sequential refinement can leverage calibrated features for deeper insights.





Fig. 6: Depth estimation prediction using one reference view (first row) and two reference views (second row) as input from LLFF comparing CaesarNeRF with GNT.



Fig. 7: Largest and smallest distances for two examples from LLFF test split when matching with training scenes. Numbers $(\times 10^{-2})$ denote the L-2 distance to the source image within the semantic feature space.

Visualizations. We present our visualization results in Figure 5, where we compare our method with others using one or two views from the LLFF dataset. Additional visual comparisons are provided in the supplementary materials. These visualizations highlight that in scenarios with few views, our method significantly surpasses other generalizable NeRF models, particularly excelling when only a single view is available. CaesarNeRF demonstrates rendering with sharper boundaries and more distinct objects.

Depth estimation. We extend our evaluation to depth prediction within the LLFF [42] dataset, focusing on challenges presented by few reference views, such as scenarios with just one or two images. In a comparison between CaesarNeRF and GNT [60], we observe in Figure 6 that GNT struggles to accurately capture the relative positions of objects when reference images are sparse. For instance, with only a single view of a flower, CaesarNeRF precisely indicates the flower's proximity to the camera compared to the leaves in the background, a distinction that GNT fails to make. Furthermore, the depth estimations provided by CaesarNeRF are consistently more reliable. In the horn example involving two views, CaesarNeRF offers better boundary delineation, showing particular strength in handling reflective surfaces such as grass in the background.

Semantic analysis for S. To assess whether the calibrated semantic representation \tilde{S} truly captures semantic details of the scene, we analyze the highest and lowest



(a) leaves

(b) orchid

Fig. 8: Synthetic results for two examples from LLFF [42], "leaves" and "orchid", using zero 1-to-3 [35] with one reference image as input and 2 degrees of vertical shift. The left image of each pair is the input, and the right one is the output of zero 1-to-3 [35].

response based on the L-2 distance between features from two different scenes, interpreting smaller distances as greater similarity. We consider the first image from each category in the LLFF dataset as a reference and extract the scene-level representation for the ten closest views to this reference image using CaesarNeRF.

In Figure 7, we present two examples, "room" and "fortress", from the LLFF dataset, representing the first image of these categories alongside images corresponding to the top-2 highest and lowest responses. This analysis reveals that the scene-level representation predominantly emphasizes structural information and objects of similar categories. For instance, with the source image "room", the highest responses correlate with images featuring table-like structures, indicating an object-centric focus. In contrast, the lowest responses include images of flowers or ponds found in open areas, which diverge significantly from the structural and object content in the source images.

Comparison with generative methods. Single-view scenarios are often addressed by generative methods [35–37] that employ diffusion models [19, 53]. While these models can produce reasonable results for object-centric renderings, they struggle with scene-level renderings from novel viewpoints. We show two LLFF examples in Figure 8 using zero 1-to-3 [35], where the left image is the input, and the right one is the output. Images rendered with zero 1-to-3 [35] suffer from style difference. Unlike NeRF-based approaches that reconstruct images from observed pixels, generative models synthesize an entire image from its semantic representation but may not maintain the style.

5 Conclusion and limitation

In this paper, we introduce CaesarNeRF, a few-shot and generalizable NeRF pipeline that combines scene-level semantic with per-pixel feature representations, aiding in rendering from novel camera positions with limited reference views. We calibrate the semantic representations across different input views and employ a sequential refinement network to offer distinct semantic representations at various levels. Our method has been extensively evaluated on a broad range of datasets, exhibiting state-of-the-art performance in both generalizable and single-scene settings.

Limitations and potential negative impact. Instead of a generative approach, CaesarNeRF relies on NeRF's scheme, using input images to render the target view. This approach restricts its ability to render parts of the scene that are not present in the reference images. CaesarNeRF may create fake images for authentication.

References

- Bao, C., Zhang, Y., Yang, B., Fan, T., Yang, Z., Bao, H., Zhang, G., Cui, Z.: Sine: Semanticdriven image-based nerf editing with prior-guided editing field. In: CVPR. pp. 20919–20929 (2023)
- Bao, Y., Ding, T., Huo, J., Li, W., Li, Y., Gao, Y.: Insertnerf: Instilling generalizability into nerf with hypernet modules. ICLR (2024)
- Bao, Y., Li, Y., Huo, J., Ding, T., Liang, X., Li, W., Gao, Y.: Where and how: Mitigating confusion in neural radiance fields from sparse inputs. arXiv preprint arXiv:2308.02908 (2023)
- Barron, J.T., Mildenhall, B., Verbin, D., Srinivasan, P.P., Hedman, P.: Mip-nerf 360: Unbounded anti-aliased neural radiance fields. In: CVPR. pp. 5470–5479 (2022)
- Buehler, C., Bosse, M., McMillan, L., Gortler, S., Cohen, M.: Unstructured lumigraph rendering. In: SIGGRAPH. pp. 425–432 (2001)
- Chen, A., Xu, Z., Zhao, F., Zhang, X., Xiang, F., Yu, J., Su, H.: Mvsnerf: Fast generalizable radiance field reconstruction from multi-view stereo. In: ICCV. pp. 14124–14133 (2021)
- Chen, W., Xu, X., Jia, J., Luo, H., Wang, Y., Wang, F., Jin, R., Sun, X.: Beyond appearance: a semantic controllable self-supervised learning framework for human-centric visual tasks. In: CVPR (2023)
- Chen, Y., Xu, H., Wu, Q., Zheng, C., Cham, T.J., Cai, J.: Explicit correspondence matching for generalizable neural radiance fields. arXiv preprint arXiv:2304.12294 (2023)
- Chibane, J., Bansal, A., Lazova, V., Pons-Moll, G.: Stereo radiance fields (srf): Learning view synthesis for sparse views of novel scenes. In: CVPR. pp. 7911–7920 (2021)
- Cong, W., Liang, H., Wang, P., Fan, Z., Chen, T., Varma, M., Wang, Y., Wang, Z.: Enhancing nerf akin to enhancing llms: Generalizable nerf transformer with mixture-of-view-experts. In: ICCV (2023)
- Debevec, P.E., Taylor, C.J., Malik, J.: Modeling and rendering architecture from photographs: A hybrid geometry-and image-based approach. In: SIGGRAPH. pp. 11–20 (1996)
- Deng, C., Jiang, C."., Qi, C.R., Yan, X., Zhou, Y., Guibas, L., Anguelov, D.: Nerdi: Singleview nerf synthesis with language-guided diffusion as general image priors. In: CVPR. pp. 20637–20647 (2023)
- Deng, K., Liu, A., Zhu, J.Y., Ramanan, D.: Depth-supervised nerf: Fewer views and faster training for free. In: CVPR. pp. 12882–12891 (2022)
- Downs, L., Francis, A., Koenig, N., Kinman, B., Hickman, R., Reymann, K., McHugh, T.B., Vanhoucke, V.: Google scanned objects: A high-quality dataset of 3d scanned household items. In: ICRA. pp. 2553–2560 (2022)
- Fridovich-Keil, S., Meanti, G., Warburg, F.R., Recht, B., Kanazawa, A.: K-planes: Explicit radiance fields in space, time, and appearance. In: CVPR. pp. 12479–12488 (2023)
- Fu, Y., Misra, I., Wang, X.: Multiplane nerf-supervised disentanglement of depth and camera pose from videos. In: ICML (2022)
- Gao, Y., Cao, Y.P., Shan, Y.: Surfelnerf: Neural surfel radiance fields for online photorealistic reconstruction of indoor scenes. In: CVPR. pp. 108–118 (2023)
- Gu, J., Liu, L., Wang, P., Theobalt, C.: Stylenerf: A style-based 3d-aware generator for highresolution image synthesis. arXiv preprint arXiv:2110.08985 (2021)
- Ho, J., Jain, A., Abbeel, P.: Denoising diffusion probabilistic models. NeurIPS pp. 6840– 6851 (2020)
- Irshad, M.Z., Zakharov, S., Liu, K., Guizilini, V., Kollar, T., Gaidon, A., Kira, Z., Ambrus, R.: Neo 360: Neural fields for sparse view synthesis of outdoor scenes. In: ICCV. pp. 9187–9198 (2023)
- Jain, A., Mildenhall, B., Barron, J.T., Abbeel, P., Poole, B.: Zero-shot text-guided object generation with dream fields. In: CVPR. pp. 867–876 (2022)

- 16 H. Zhu and T. Ding et al.
- Jain, A., Tancik, M., Abbeel, P.: Putting nerf on a diet: Semantically consistent few-shot view synthesis. In: ICCV. pp. 5885–5894 (2021)
- Jang, W., Agapito, L.: Codenerf: Disentangled neural radiance fields for object categories. In: ICCV. pp. 12949–12958 (2021)
- Jiang, Y., Hedman, P., Mildenhall, B., Xu, D., Barron, J.T., Wang, Z., Xue, T.: Alignerf: Highfidelity neural radiance fields via alignment-aware training. In: CVPR. pp. 46–55 (2023)
- Johari, M.M., Lepoittevin, Y., Fleuret, F.: Geonerf: Generalizing nerf with geometry priors. In: CVPR. pp. 18365–18375 (2022)
- Johnson, J., Alahi, A., Fei-Fei, L.: Perceptual losses for real-time style transfer and superresolution. In: Computer Vision–ECCV 2016: 14th European Conference, Amsterdam, The Netherlands, October 11-14, 2016, Proceedings, Part II 14. pp. 694–711. Springer (2016)
- Kania, A., Kasymov, A., Zięba, M., Spurek, P.: Hypernerfgan: Hypernetwork approach to 3d nerf gan. arXiv preprint arXiv:2301.11631 (2023)
- Kingma, D.P., Ba, J.: Adam: A method for stochastic optimization. arXiv preprint arXiv:1412.6980 (2014)
- Kwak, M., Song, J., Kim, S.: Geconerf: Few-shot neural radiance fields via geometric consistency. arXiv preprint arXiv:2301.10941 (2023)
- Li, Z., Wang, Q., Cole, F., Tucker, R., Snavely, N.: Dynibar: Neural dynamic image-based rendering. In: CVPR. pp. 4273–4284 (2023)
- Lin, C.H., Gao, J., Tang, L., Takikawa, T., Zeng, X., Huang, X., Kreis, K., Fidler, S., Liu, M.Y., Lin, T.Y.: Magic3d: High-resolution text-to-3d content creation. In: CVPR. pp. 300– 309 (2023)
- Lin, H., Peng, S., Xu, Z., Yan, Y., Shuai, Q., Bao, H., Zhou, X.: Efficient neural radiance fields for interactive free-viewpoint video. In: SIGGRAPH Asia 2022 Conference Papers. pp. 1–9 (2022)
- Lin, Y., Bai, H., Li, S., Lu, H., Lin, X., Xiong, H., Wang, L.: Componerf: Text-guided multiobject compositional nerf with editable 3d scene layout. arXiv preprint arXiv:2303.13843 (2023)
- Liu, L., Habermann, M., Rudnev, V., Sarkar, K., Gu, J., Theobalt, C.: Neural actor: Neural free-view synthesis of human actors with pose control. TOC 40(6), 1–16 (2021)
- 35. Liu, R., Wu, R., Van Hoorick, B., Tokmakov, P., Zakharov, S., Vondrick, C.: Zero-1-to-3: Zero-shot one image to 3d object. In: ICCV (2023)
- Liu, X., Kao, S.h., Chen, J., Tai, Y.W., Tang, C.K.: Deceptive-nerf: Enhancing nerf reconstruction using pseudo-observations from diffusion models. arXiv preprint arXiv:2305.15171 (2023)
- Liu, Y., Lin, C., Zeng, Z., Long, X., Liu, L., Komura, T., Wang, W.: Syncdreamer: Generating multiview-consistent images from a single-view image. arXiv preprint arXiv:2309.03453 (2023)
- Liu, Y., Peng, S., Liu, L., Wang, Q., Wang, P., Theobalt, C., Zhou, X., Wang, W.: Neural rays for occlusion-aware image-based rendering. In: CVPR. pp. 7824–7833 (2022)
- Mariotti, O., Mac Aodha, O., Bilen, H.: Viewnerf: Unsupervised viewpoint estimation using category-level neural radiance fields. arXiv preprint arXiv:2212.00436 (2022)
- Martin-Brualla, R., Radwan, N., Sajjadi, M.S., Barron, J.T., Dosovitskiy, A., Duckworth, D.: Nerf in the wild: Neural radiance fields for unconstrained photo collections. In: CVPR. pp. 7210–7219 (2021)
- 41. Metzer, G., Richardson, E., Patashnik, O., Giryes, R., Cohen-Or, D.: Latent-nerf for shapeguided generation of 3d shapes and textures. In: CVPR. pp. 12663–12673 (2023)
- Mildenhall, B., Srinivasan, P.P., Ortiz-Cayon, R., Kalantari, N.K., Ramamoorthi, R., Ng, R., Kar, A.: Local light field fusion: Practical view synthesis with prescriptive sampling guidelines. TOG 38(4), 1–14 (2019)

Calibrated Semantic Representation for Few-Shot Generalizable Neural Rendering

17

- Mildenhall, B., Srinivasan, P.P., Tancik, M., Barron, J.T., Ramamoorthi, R., Ng, R.: Nerf: Representing scenes as neural radiance fields for view synthesis. In: ECCV (2020)
- Mildenhall, B., Srinivasan, P.P., Tancik, M., Barron, J.T., Ramamoorthi, R., Ng, R.: Nerf: Representing scenes as neural radiance fields for view synthesis. Communications of the ACM 65(1), 99–106 (2021)
- Niemeyer, M., Barron, J.T., Mildenhall, B., Sajjadi, M.S., Geiger, A., Radwan, N.: Regnerf: Regularizing neural radiance fields for view synthesis from sparse inputs. In: CVPR. pp. 5480–5490 (2022)
- Noguchi, A., Sun, X., Lin, S., Harada, T.: Neural articulated radiance field. In: ICCV. pp. 5762–5772 (2021)
- Park, K., Sinha, U., Barron, J.T., Bouaziz, S., Goldman, D.B., Seitz, S.M., Martin-Brualla, R.: Nerfies: Deformable neural radiance fields. In: ICCV. pp. 5865–5874 (2021)
- Peng, S., Dong, J., Wang, Q., Zhang, S., Shuai, Q., Zhou, X., Bao, H.: Animatable neural radiance fields for modeling dynamic human bodies. In: ICCV. pp. 14314–14323 (2021)
- Poole, B., Jain, A., Barron, J.T., Mildenhall, B.: Dreamfusion: Text-to-3d using 2d diffusion. arXiv preprint arXiv:2209.14988 (2022)
- Pumarola, A., Corona, E., Pons-Moll, G., Moreno-Noguer, F.: D-nerf: Neural radiance fields for dynamic scenes. In: CVPR. pp. 10318–10327 (2021)
- Qi, C.R., Su, H., Mo, K., Guibas, L.J.: Pointnet: Deep learning on point sets for 3d classification and segmentation. In: CVPR. pp. 652–660 (2017)
- Roessle, B., Barron, J.T., Mildenhall, B., Srinivasan, P.P., Nießner, M.: Dense depth priors for neural radiance fields from sparse input views. In: CVPR. pp. 12892–12901 (2022)
- Rombach, R., Blattmann, A., Lorenz, D., Esser, P., Ommer, B.: High-resolution image synthesis with latent diffusion models. In: CVPR. pp. 10684–10695 (2022)
- 54. Schönberger, J.L., Frahm, J.M.: Structure-from-motion revisited. In: CVPR (2016)
- 55. Schönberger, J.L., Zheng, E., Pollefeys, M., Frahm, J.M.: Pixelwise view selection for unstructured multi-view stereo. In: ECCV (2016)
- Schwarz, K., Liao, Y., Niemeyer, M., Geiger, A.: Graf: Generative radiance fields for 3daware image synthesis. NeurIPS 33, 20154–20166 (2020)
- Shue, J.R., Chan, E.R., Po, R., Ankner, Z., Wu, J., Wetzstein, G.: 3d neural field generation using triplane diffusion. In: CVPR. pp. 20875–20886 (2023)
- Suhail, M., Esteves, C., Sigal, L., Makadia, A.: Generalizable patch-based neural rendering. In: ECCV. pp. 156–174 (2022)
- Trevithick, A., Yang, B.: Grf: Learning a general radiance field for 3d representation and rendering. In: ICCV. pp. 15182–15192 (2021)
- Varma, M., Wang, P., Chen, X., Chen, T., Venugopalan, S., Wang, Z.: Is attention all that nerf needs? In: ICLR (2023)
- Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A.N., Kaiser, Ł., Polosukhin, I.: Attention is all you need. NeurIPS 30 (2017)
- Waechter, M., Moehrle, N., Goesele, M.: Let there be color! large-scale texturing of 3d reconstructions. In: ECCV. pp. 836–850 (2014)
- 63. Wang, G., Chen, Z., Loy, C.C., Liu, Z.: Sparsenerf: Distilling depth ranking for few-shot novel view synthesis. In: ICCV (2023)
- Wang, Q., Wang, Z., Genova, K., Srinivasan, P.P., Zhou, H., Barron, J.T., Martin-Brualla, R., Snavely, N., Funkhouser, T.: Ibrnet: Learning multi-view image-based rendering. In: CVPR. pp. 4690–4699 (2021)
- Wang, T., Zhang, B., Zhang, T., Gu, S., Bao, J., Baltrusaitis, T., Shen, J., Chen, D., Wen, F., Chen, Q., et al.: Rodin: A generative model for sculpting 3d digital avatars using diffusion. In: CVPR. pp. 4563–4573 (2023)
- Wei, Y., Liu, S., Rao, Y., Zhao, W., Lu, J., Zhou, J.: Nerfingmvs: Guided optimization of neural radiance fields for indoor multi-view stereo. In: CVPR. pp. 5610–5619 (2021)

- 18 H. Zhu and T. Ding et al.
- Wizadwongsa, S., Phongthawee, P., Yenphraphai, J., Suwajanakorn, S.: Nex: Real-time view synthesis with neural basis expansion. In: CVPR. pp. 8534–8543 (2021)
- Xiangli, Y., Xu, L., Pan, X., Zhao, N., Rao, A., Theobalt, C., Dai, B., Lin, D.: Bungeenerf: Progressive neural radiance field for extreme multi-scale scene rendering. In: ECCV (2022)
- Xie, C., Park, K., Martin-Brualla, R., Brown, M.: Fig-nerf: Figure-ground neural radiance fields for 3d object category modelling. In: 3DV. pp. 962–971 (2021)
- Xu, D., Jiang, Y., Wang, P., Fan, Z., Shi, H., Wang, Z.: Sinnerf: Training neural radiance fields on complex scenes from a single image. In: ECCV. pp. 736–753 (2022)
- Xu, Q., Xu, Z., Philip, J., Bi, S., Shu, Z., Sunkavalli, K., Neumann, U.: Point-nerf: Pointbased neural radiance fields. In: CVPR. pp. 5438–5448 (2022)
- Yang, B., Zhang, Y., Xu, Y., Li, Y., Zhou, H., Bao, H., Zhang, G., Cui, Z.: Learning objectcompositional neural radiance field for editable scene rendering. In: ICCV. pp. 13779–13788 (2021)
- Yang, H., Hong, L., Li, A., Hu, T., Li, Z., Lee, G.H., Wang, L.: Contranerf: Generalizable neural radiance fields for synthetic-to-real novel view synthesis via contrastive learning. In: CVPR. pp. 16508–16517 (2023)
- Yang, J., Pavone, M., Wang, Y.: Freenerf: Improving few-shot neural rendering with free frequency regularization. In: CVPR. pp. 8254–8263 (2023)
- Yao, Y., Luo, Z., Li, S., Fang, T., Quan, L.: Mvsnet: Depth inference for unstructured multiview stereo. In: ECCV. pp. 767–783 (2018)
- Yu, A., Ye, V., Tancik, M., Kanazawa, A.: pixelnerf: Neural radiance fields from one or few images. In: CVPR. pp. 4578–4587 (2021)
- Yu, X., Xu, M., Zhang, Y., Liu, H., Ye, C., Wu, Y., Yan, Z., Zhu, C., Xiong, Z., Liang, T., et al.: Mvimgnet: A large-scale dataset of multi-view images. In: CVPR. pp. 9150–9161 (2023)
- Zhang, J., Yang, G., Tulsiani, S., Ramanan, D.: Ners: Neural reflectance surfaces for sparseview 3d reconstruction in the wild. NeurIPS 34, 29835–29847 (2021)
- 79. Zhang, J., Li, X., Wan, Z., Wang, C., Liao, J.: Text2nerf: Text-driven 3d scene generation with neural radiance fields. arXiv preprint arXiv:2305.11588 (2023)
- Zhenxing, M., Xu, D.: Switch-nerf: Learning scene decomposition with mixture of experts for large-scale neural radiance fields. In: ICLR (2023)
- Zhou, T., Tucker, R., Flynn, J., Fyffe, G., Snavely, N.: Stereo magnification: Learning view synthesis using multiplane images. arXiv preprint arXiv:1805.09817 (2018)
- Zhou, Z., Tulsiani, S.: Sparsefusion: Distilling view-conditioned diffusion for 3d reconstruction. In: CVPR. pp. 12588–12597 (2023)
- Zhu, H., Sun, Y., Liu, C., Xia, L., Luo, J., Qiao, N., Nevatia, R., Kuo, C.H.: Multimodal neural radiance field. In: ICRA. pp. 9393–9399 (2023)
- Zhu, H., Zheng, Z., Zheng, W., Nevatia, R.: Cat-nerf: Constancy-aware tx2former for dynamic body modeling. In: CVPRW. pp. 6618–6627 (2023)
- Zhuang, Y., Zhu, H., Sun, X., Cao, X.: Mofanerf: Morphable facial neural radiance field. In: ECCV. pp. 268–285 (2022)
- Zimny, D., Trzciński, T., Spurek, P.: Points2nerf: Generating neural radiance fields from 3d point cloud. arXiv preprint arXiv:2206.01290 (2022)