

Entangled View-Epipolar Information Aggregation for Generalizable Neural Radiance Fields

Zhiyuan Min¹ Yawei Luo^{1,*} Wei Yang² Yuesong Wang² Yi Yang¹
¹Zhejiang University, China ²Huazhong University of Science and Technology, China
 {minzhiyuan, yaweiluo, yangyics}@zju.edu.cn {weiyangcs, yuesongwang}@hust.edu.cn

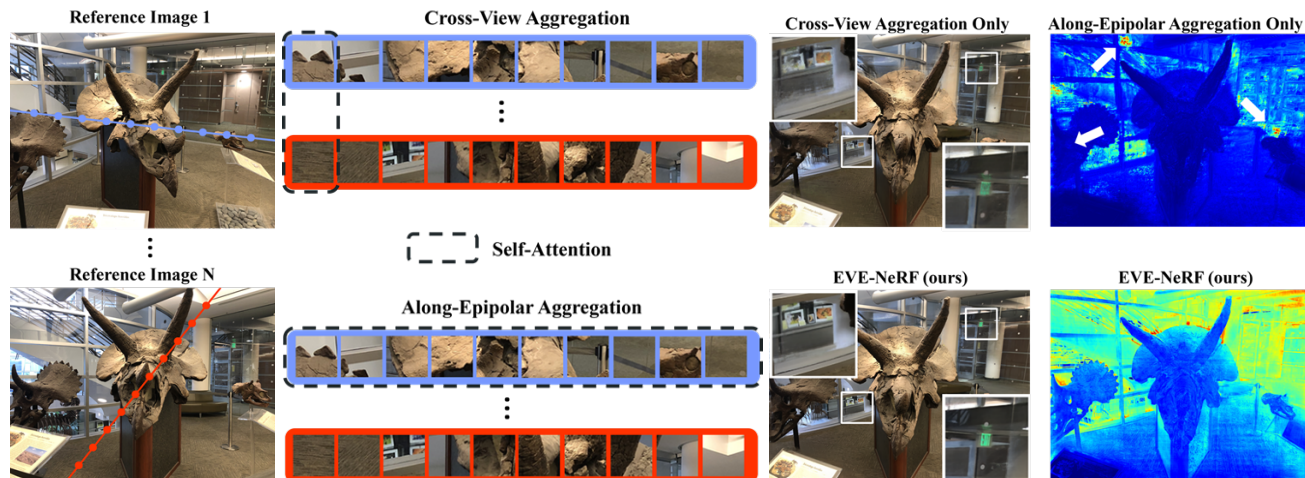


Figure 1. Given the sampling points along a target ray that are re-projected on the epipolar lines in each source view, existing approaches [43, 49] employ attention mechanism to aggregate the cross-view features for each sampling point and perform epipolar aggregation of sampling points along the epipolar lines within individual views, either sequentially or circularly. However, our investigation reveals the limitations in existing strategies: exclusively aggregating cross-view information results in rendering artifacts, stemming from the absence of **appearance continuity between adjacent depth** provided by epipolar cues. Conversely, relying solely on epipolar information leads to depth map discontinuities due to the absence of **geometry consistency across multiple views**. Our proposed EVE-NeRF harnesses both cross-view and along-epipolar information in an entangled manner and effectively addresses the above issues.

Abstract

Generalizable NeRF can directly synthesize novel views across new scenes, eliminating the need for scene-specific re-training in vanilla NeRF. A critical enabling factor in these approaches is the extraction of a generalizable 3D representation by aggregating source-view features. In this paper, we propose an Entangled View-Epipolar Information Aggregation method dubbed **EVE-NeRF**. Different from existing methods that consider cross-view and along-epipolar information independently, EVE-NeRF conducts the view-epipolar feature aggregation in an entangled manner by injecting the scene-invariant appearance continuity and geometry consistency priors to the aggregation process. Our approach effectively mitigates the potential lack of inherent geometric and appearance constraints resulting from one-dimensional interactions, thus further boosting the 3D

representation generalizability. EVE-NeRF attains state-of-the-art performance across various evaluation scenarios. Extensive experiments demonstrate that, compared to prevailing single-dimensional aggregation, the entangled network excels in the accuracy of 3D scene geometry and appearance reconstruction. Our code is publicly available at <https://github.com/tatakai1/EVENeRF>.

1. Introduction

The neural radiance fields (NeRF) [35], along with its subsequent refinements [1, 2, 62], have exhibited remarkable efficacy in the domain of novel view synthesis [26, 27, 54, 55], showcasing immense potential for applications in smart education [28, 33], medical treatment [16, 36], and automatic driving [29–32]. Despite these advancements, the methods along this vein often pertain to the training scene thus necessitating re-training for synthesizing new scenes. Such

*Corresponding author

drawbacks severely constrain their practical applications.

More recently, the development of generalizable NeRF models [4, 52, 63] has emerged as a promising solution to address this challenge. These models can directly synthesize novel views across new scenes, eliminating the need for scene-specific re-training. A critical enabling factor in these approaches is the synthesis of a generalizable 3D representation by aggregating source-view features. Instead of densely aggregating every pixel in the source images, prior works draw inspiration from the epipolar geometric constraint across multiple views to aggregate view or epipolar information [21, 40, 43, 49, 51]. To capitalize on cross-view prior, specific methodologies [21, 49] interact with the re-projected feature information in the reference view at a pre-defined depth. On the along-epipolar aspect [43, 44], some methods employ self-attention mechanisms to sequentially obtain the entire epipolar line features in each reference view.

We posit that both view and epipolar aggregation are crucial for learning a generalizable 3D representation: cross-view feature aggregation is pivotal to capturing geometric information, as the features from different views that match tend to be on the surface of objects. Concurrently, epipolar feature aggregation contributes by extracting depth-relevant appearance features from the reference views associated with the target ray, thus achieving a more continuous appearance representation. Nevertheless, the prevailing methods often execute view and epipolar aggregation independently [21, 49] or in a sequential manner [43], thereby overlooking the simultaneous interaction of appearance and geometry information.

In this paper, we introduce a novel **Entangled View-Epipolar information aggregation network**, denoted as **EVE-NeRF**. EVE-NeRF is designed to enhance the quality of generalizable 3D representation through the simultaneous utilization of complementary appearance and geometry information. The pivotal components of EVE-NeRF are the **View-Epipolar Interaction Module (VEI)** and the **Epipolar-View Interaction Module (EVI)**. Both modules adopt a dual-branch structure to concurrently integrate view and epipolar information. On one hand, VEI comprises a view transformer in its first branch to engage with the features of sampling points re-projected on all source views. In the second branch, VEI is equipped with an **Along-Epipolar Perception** submodule to inject the appearance continuity prior to the view aggregation results. On the other hand, EVI consists of an epipolar transformer in its first branch to aggregate features from sampling points along the entire epipolar line in each source view. In the second branch, EVI utilizes a **Multi-View Calibration** submodule to incorporate the geometry consistency prior to the epipolar aggregation representation. The alternating organization of EVI and VEI results in a generalizable condition for predicting the color of target rays based on NeRF volumetric rendering.

Compared to the prevailing methods such as GNT [49] and GPNR [43], EVE-NeRF distinguishes itself in its ability to synthesize a target ray by entangling epipolar and view information. This capability serves to offset the appearance and geometry prior losses that typically arise from single-dimensional aggregation operations (see Figure 1). Our main contributions can be summarized as follows:

- Through extensive investigation, we have revealed the under-explored issues of prevailing cross-view and along-epipolar information aggregation methods for generalizable NeRF.
- We propose EVE-NeRF, which harnesses the along-epipolar and cross-view information in an entangled manner. EVE-NeRF complements the cross-view aggregation with appearance continuity prior and calibrates the along-epipolar aggregation with geometry consistency prior.
- EVE-NeRF produces more realistic novel-perspective images and depth maps for previously unseen scenes without any additional ground-truth 3D data. Experiments demonstrate that EVE-NeRF achieves state-of-the-art performance in various novel scene synthesis tasks.

2. Related Work

NeRF and Generalizable NeRF. Recently, NeRF [35] has made groundbreaking advancements in the field of novel view synthesis through a compact implicit representation based on differentiable rendering. Subsequent developments in the NeRF framework have explored various avenues, enhancing rendering quality [1–3, 18, 64], accelerating rendering speed [5, 8, 17, 39, 45, 62], applicability to both rigid and non-rigid dynamic scenes [15, 37, 38, 47, 48, 65], and extending its capabilities for editing [24, 46, 50, 57, 58].

The original NeRF and its subsequent improvements have achieved successful performance but suffered from the limitation of being trainable and renderable only in a single scene, which restricts their practical applications [13, 60, 61]. One solution to this issue is conditioning on CNN features from the known view images, which align with the input coordinates of NeRF. PixelNeRF [63] encodes input images into pixel-aligned feature grids, combining image features with corresponding spatial positions and view directions in a shared MLP to output colors and densities. MVSNeRF [4] utilizes a cost volume to model the scene, with interpolated features on volume conditioned. Our approach also employs CNN features from known views, and we input the processed, pixel-aligned features into NeRF’s MLP network to predict colors and densities. However, unlike PixelNeRF and similar methods [4, 7, 63], which use average pooling for handling multiple views, our approach learns multi-view information and assigns weights to each view based on its relevance.

Generalizable NeRF with Transformers. More recently, generalizable novel view synthesis methods [10, 11, 19, 22,

[23, 40, 43, 51, 52] have incorporated transformer-based networks to enhance visual features from known views. These approaches employ self-attention or cross-attention mechanisms along various dimensions such as depth, view, or epipolar, enabling high-quality feature interaction and aggregation. GeoNeRF [21] concatenates view-independent tokens with view-dependent tokens and feeds them into a cross-view aggregation transformer network to enhance the features of the cost volume. GPNR [43] employs a 3-stage transformer-based aggregation network that sequentially interacts with view, epipolar, and view information. GNT [49] and its subsequent work, GNT-MOVE [11], utilize a 2-stage transformer-based aggregation network, first performing cross-view aggregation and then engaging depth information interaction. ContraNeRF [59] initially employs a two-stage transformer-based network for geometry-aware feature extraction, followed by the computation of positive and negative sample contrastive loss based on ground-truth depth values.

Inspired by these developments, we have analyzed the limitations of single-dimensional aggregation transformer networks and introduced EVE-NeRF that achieves efficient interaction between the complementary appearance and geometry information across different dimensions. Moreover, our method does not require any ground-truth depth information for model training.

3. Problem Formulation

Our objective is to train a generalizable NeRF capable of comprehending 3D information from scenes the model has never encountered before and rendering new perspective images. Specifically, given M source images for a particular scene $I = \{I_i\}_{i=1}^M$ and their corresponding camera intrinsic and extrinsic parameters $K = \{K_i\}_{i=1}^M$, $P = \{P_i = [R_i, t_i]\}_{i=1}^M$, most generalizable NeRF methods [4, 52, 63] can be formulated with a generalizable feature extraction network \mathcal{F}_θ and a rendering network \mathcal{G}_ϕ :

$$\mathcal{F}_\theta : (I, K, P) \rightarrow z, \quad \mathcal{G}_\phi : (x, d, z) \rightarrow (c, \sigma), \quad (1)$$

where x and d represent the 3D point position and the direction of the target ray’s sampling points, while c and σ are the predicted color and density, respectively. Similar to vanilla NeRF, c and σ are utilized to compute the final color value of the target ray through volume rendering. The variable z represents generalizable 3D representation of the scene provided by the feature extraction network. Θ and ϕ denote the learnable parameters of the networks.

4. Methodology

Overview. Figure 2 provides an overview of EVE-NeRF, which includes a lightweight CNN-based image feature extractor, two dual-branch transformer-based modules named

View-Epipolar Interaction (VEI) and Epipolar-View Interaction (EVI), respectively, and a conditioned NeRF decoder. Source images are first forwarded to a lightweight CNN and are transferred to feature maps. In the following, VEI and EVI are alternatively organized to aggregate the view-epipolar features in an entangled manner. The inter-branch information interaction mechanism within VEI and EVI capitalizes on the scene-invariant geometry and appearance priors to further calibrate the aggregated features. The output of the Entangle View-Epipolar Information Aggregation is a generalizable 3D representation z . Finally, a conditioned NeRF decoder is employed for predicting the color and density values of the target ray based on z for volume rendering.

4.1. Lightweight CNN

For M source views input $\{I_i\}_{i=1}^M$, we first extract convolutional features $\{F_i^c\}_{i=1}^M$ for each view independently using a lightweight CNN with sharing weights (see Appendix A.4). Unlike previous generalizable NeRF methods [4, 49] that employ deep convolutional networks like U-Net [41], we use this approach since convolutional features with large receptive fields may not be advantageous for extracting scene-generalizable features [43]. Additionally, features derived from the re-projected sampling points guided by epipolar geometry are more focused on local information [19].

4.2. View-Epipolar Interaction

The View-Epipolar Interaction Module (VEI) is designed as a dual-branch structure, with one branch comprising the View Transformer and the other the Along-Epipolar Perception. The VEI input $X \in \mathbb{R}^{N \times M \times C}$ comes from the CNN feature map interpolated features or from the output of the previous layer of EVI, and the VEI output Y_{VEI} is used as the input to the current layer of EVI.

View Transformer. The View Transformer is responsible for aggregating features across view dimensions. The view transformer takes the input X , allowing it to perform self-attention operations in the view dimension (M). To be more specific, the query, key, and value matrices are computed using linear mappings:

$$Q = XW_Q, K = XW_K, V = XW_V, \quad (2)$$

where $W_Q, W_K, W_V \in \mathbb{R}^{C \times C}$ are the linear mappings without biases. These matrices are then split into h heads $Q = [Q^1, \dots, Q^h]$, $K = [K^1, \dots, K^h]$, and $V = [V^1, \dots, V^h]$, each with $d = C/h$ channels. To enable the model to learn the relative spatial relationships between the target view and the source views, we integrate the differences Δd^s (see Appendix A.2) between the target view and the source views into the self-attention mechanism:

$$\tilde{X}^i = \text{softmax} \left((Q^i + \Delta d^s) (K^i)^\top \right) V^i. \quad (3)$$

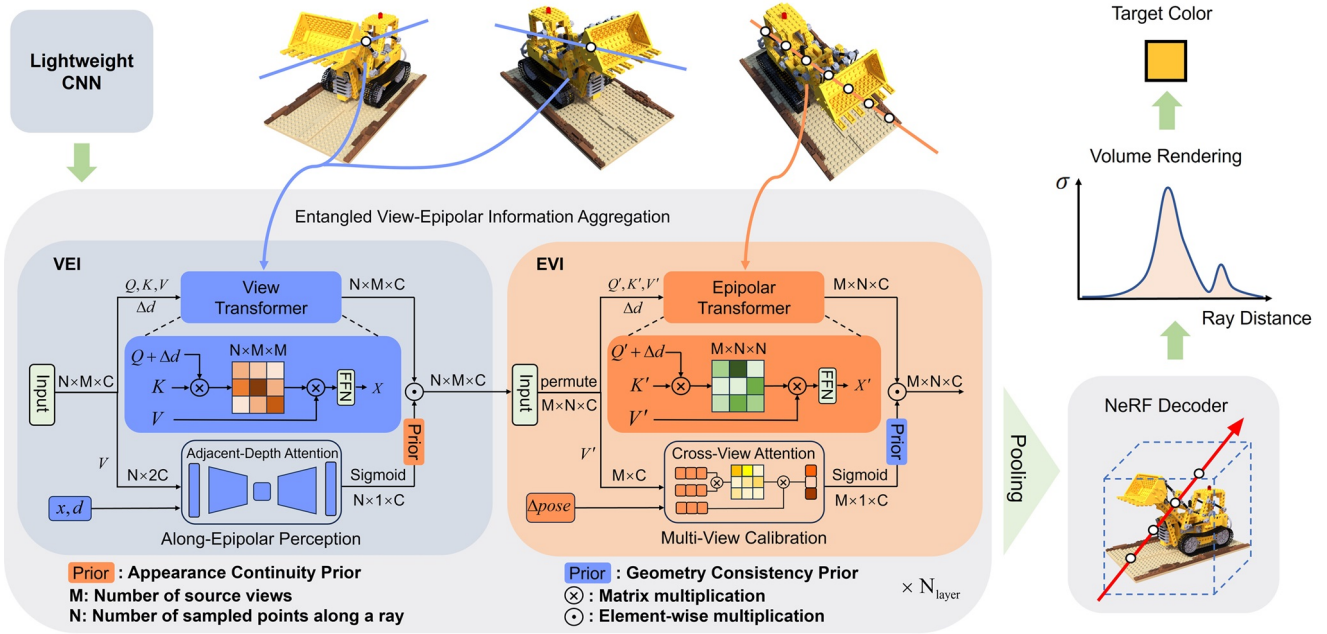


Figure 2. Pipeline of EVE-NeRF. 1) We first employ a lightweight CNN to extract features of the epipolar sampling points from source views. 2) Through the Entangled View-Epipolar Information Aggregation, we complementarily enable information interaction in both the view and epipolar dimensions to produce generalizable multi-view epipolar features. 3) We use the NeRF Decoder to obtain color and density for the sampling points and predict the target color based on volume rendering.

Ultimately, we obtain $\tilde{\mathbf{X}} = [\tilde{\mathbf{X}}^i, \dots, \tilde{\mathbf{X}}^h]$, and we employ a conventional Feed-Forward Network (FFN) to perform point-wise feature transformation:

$$\mathbf{Y} = \text{FFN}(\tilde{\mathbf{X}}) + \tilde{\mathbf{X}}. \quad (4)$$

Along-Epipolar Perception. The Along-Epipolar Perception, serving as the second branch of VEI, aims to extract view-independent depth information to provide appearance continuity prior to the 3D representation. We compute the mean and variance of $\mathbf{V} \in \mathbb{R}^{N \times M \times C}$ in the view dimension (M) within the view transformer to obtain the global view-independent feature $\mathbf{f}^0 \in \mathbb{R}^{N \times 2C}$. We proceed to perceive the depth information along the entire ray through an adjacent-depth attention (1D Convolution AE) in the ray dimension (N). Since the information along an epipolar line is inherently continuous, a convolution operation that is seen as a kind of adjacent attention can learn the appearance continuity prior, which predicts the importance weights w_i^v for the sampling points:

$$\mathbf{f}^1 = \text{concat}(\mathbf{f}^0, \mathbf{x}, \mathbf{d}),$$

$$\{w_i^v\}_{i=1}^N = \text{sigmoid}(\text{AE}(\{\mathbf{f}_i^1\}_{i=1}^N)), \quad (5)$$

where \mathbf{x} and \mathbf{d} refer to the 3D point position and the direction of the target ray's sampling point. Particularly, \mathbf{d} is copied to the same dimension as \mathbf{x} . GeoNeRF [21] also

employs an AE network to predict coherent volume densities. However, our approach is more similar to an adjacent attention mechanism predicting depth importance weights and learning appearance continuity prior based on global epipolar features.

Combining the output of the View Transformer and Along-Epipolar Perception, the final output of VEI is calculated as follows:

$$\mathbf{Y}_{VEI} = \mathbf{w}^v \cdot \mathbf{Y}, \quad (6)$$

where $\mathbf{w}^v = [w_1^v, \dots, w_N^v]$, \mathbf{Y}_{VEI} denotes the VEI's output, and \cdot denotes element-wise multiplication.

4.3. Epipolar-View Interaction

Similar to VEI, The Epipolar-View Interaction Module (EVI) consists of two branches, the Epipolar Transformer and the Multi-View Calibration. The EVI input $\mathbf{X}' \in \mathbb{R}^{M \times N \times C}$ comes from the output of the current layer of VEI, and the EVI output \mathbf{Y}_{EVI} is used as the input to the next layer of EVIs or as the total output of the aggregation network.

Epipolar Transformer. The Epipolar Transformer takes the input \mathbf{X}' , enabling self-attention operations in the epipolar dimension (N). In particular, the epipolar transformer shares the same network structure as the view transformer above:

$$\begin{aligned}
Q' &= X'W'_Q, K' = X'W'_K, V' = X'W'_V, \\
\tilde{X}'^i &= \text{softmax} \left(\left(Q'^i + \Delta d'^s \right) \left(K'^i \right)^\top \right) V'^i, \\
Y' &= \text{FFN}(\tilde{X}') + \tilde{X}', \tag{7}
\end{aligned}$$

where $X'[i, j, k] = X[j, i, k]$, $d'^s[i, j, k] = d^s[j, i, k]$, i, j, k denote the 1st (M), 2nd (N), and 3rd dimensions (C) respectively.

Multi-View Calibration. The Multi-View Calibration, serving as the second branch of the EVI module, is employed to aggregate cross-view features and provide geometry consistency prior, aiming at calibrating the epipolar features. We calculate the weight values w_j^e for the target rays in each source view using the cross-view attention mechanism. In this process, we utilize $V' \in \mathbb{R}^{M \times N \times C}$ from the epipolar transformer as the input:

$$\begin{aligned}
q &= \max(V') + \text{linear}(\Delta \text{pose}), \\
\{w_j^e\}_{j=1}^M &= \text{sigmoid}(\text{Self-Attn}(q, q, q)), \tag{8}
\end{aligned}$$

where Δpose (see Appendix A.3) refers to the difference between the source view camera pose and the target view camera pose, and linear denotes the linear layer. Ultimately, incorporating the regression results of multi-view calibration, the output of the EVI is calculated as follows:

$$Y_{EVI} = w^e \cdot Y', \tag{9}$$

where $w^e = [w_1^e, \dots, w_M^e]$, Y_{EVI} denotes the EVI's output, and \cdot denotes element-wise multiplication.

4.4. Conditioned NeRF Decoder

We follow the established techniques of previous works [63] to construct an MLP-based rendering network. We also condition 3D points on a ray using the generalizable 3D representation z based on Eq. 1. Nevertheless, we diverge from the traditional MLP decoder [35], which processes each point on a ray independently. Instead, we take a more advanced approach by introducing cross-point interactions. For this purpose, we employ the ray Transformer from IBRNet [52] in our implementation. After the rendering network predicts the emitted color c and volume density σ , we can generate target pixel color using volume rendering [35]:

$$C = \sum_{i=1}^N T_i (1 - \exp(-\sigma_i \delta_i)) c_i, \quad T_i = \exp\left(-\sum_{j=1}^{i-1} \sigma_j \delta_j\right), \tag{10}$$

where c_i, σ_i which are calculated based on Eq. 1, refer to the color and density of the i -th sampling point on the ray.

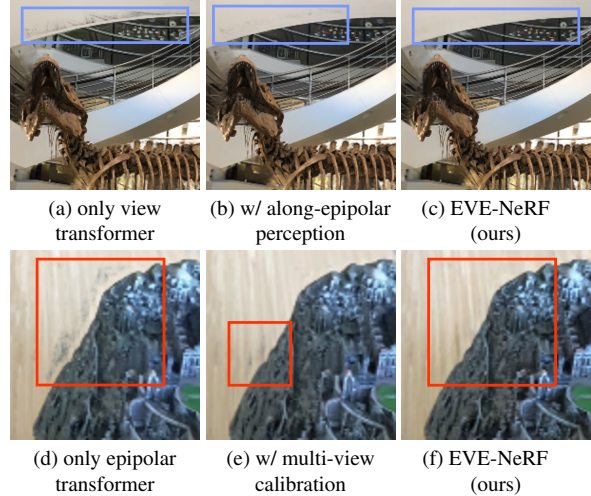


Figure 3. The along-epipolar perception provides appearance continuity prior through adjacent-depth attention along the ray, while the multi-view calibration offers geometry consistency prior via cross-view attention. Our proposed method significantly reduces artifacts in rendering new views compared to single-dimension transformers.

4.5. Training Objectives

EVE-NeRF is trained solely using a photometric loss function, without the need for additional ground-truth 3D data. Specifically, our training loss function is as follows:

$$\mathcal{L} = \sum_{p \in \mathcal{P}} \|C_{pred} - C_{gt}\|_2^2, \tag{11}$$

where \mathcal{P} represents a set of pixel points in a training batch, C_{pred}, C_{gt} respectively represent the rendering color for pixel p and the ground-truth color.

5. Experiments

5.1. Implementation Details

We randomly sample 2,048 rays per batch, each with $N = 88$ sampling points along the rays. Our lightweight CNN and EVE-NeRF models are trained for 250,000 iterations using an Adam optimizer with initial learning rates of $1e^{-3}$ and $5e^{-4}$, respectively, and an exponential learning rate decay. The training is performed end-to-end on 4 V100-32G GPUs for 3 days. To evaluate our model, we use common metrics such as PSNR, SSIM, and LPIPS and compare the results qualitatively and quantitatively with other generalizable neural rendering approaches. More details such as network hyperparameters are provided in Appendix A.4.

5.2. Comparative Studies

To provide a fair comparison with prior works [4, 49, 52], we conducted experiments under 2 different settings: Generaliz-




Figure 4. Qualitative comparison of EVE-NeRF with IBRNet[52] and GNT[49] in setting 1. The first, second, and third rows correspond to the Fern scene from LLFF, the Mic scene from Blender, and the Crest scene from Shiny, respectively. Our method, EVE-NeRF, demonstrates superior capability compared to the baselines in accurately reconstructing the geometry, appearance, and complex texture regions. In particular, our method successfully reconstructs the leaves and the surrounding area in the Fern scene.

Method	LLFF			Blender			Shiny		
	PSNR \uparrow	SSIM \uparrow	LPIPS \downarrow	PSNR \uparrow	SSIM \uparrow	LPIPS \downarrow	PSNR \uparrow	SSIM \uparrow	LPIPS \downarrow
PixelNeRF [63]	18.66	0.588	0.463	22.65	0.808	0.202	-	-	-
IBRNet [52]	25.17	0.813	0.200	26.73	0.908	0.101	23.60	0.785	0.180
GPNR [43]	25.35	0.818	0.198	28.29	0.927	0.080	25.72	0.880	0.175
NeuRay [25]	25.72	0.880	0.175	26.48	0.944	0.091	24.12	0.860	0.170
GNT [49]	25.53	0.836	0.178	26.01	0.925	0.088	27.19	0.912	0.083
EVE-NeRF	27.16	0.912	0.134	27.03	0.952	0.072	28.01	0.935	0.083

Table 1. Results for setting 1. Our proposed method, EVE-NeRF, outperforms most of the baselines on the majority of the metrics. With the exception of PixelNeRF [63], all baseline methods [25, 43, 52] employ sequential or independent transformer-based single-dimensional ray aggregation. In contrast, our approach is based on a dual-branch structure, enabling multi-dimensional interactions for both view and epipolar information. The results confirm that our method’s multi-dimensional ray feature aggregation is superior to the single-dimensional aggregation used in the baselines.

able NVS and Few-Shot NVS, as was done in GPNR [43].

Setting 1: Generalizable NVS. Following IBRnet [52], we set up the reference view pool comprising $k \times M$ proximate views. M views are chosen at random from this pool to serve as source views. Throughout the training process, the parameters k and M are subject to uniform random sampling, with k drawn from (1, 3) and M from (8, 12). During evaluation, we fix the number of source views $M = 10$.

For the training dataset, we adopt object renderings of 1,030 models from Google Scanned Object [12], RealEstate10K [66], 100 scenes from the Spaces dataset [14] and 95 real scenes from handheld cellphone captures [34, 52].

For the evaluation dataset, we use Real Forward-Facing [34], Blender [35] and Shiny [56].

Table 1 presents the quantitative results, while Figure 4 showcases the qualitative results. As shown in Table 1, comparison to methods [25, 43, 52] using transformer networks for feature aggregation, our approach outperforms them under most metrics. Our method outperforms SOTA method GNT [49] by 4.43% \uparrow PSNR, 4.83% \uparrow SSIM, 14.3% \downarrow LPIPS in 3 evaluating dataset evenly. Such results verify the effectiveness of introducing the complementary appearance continuity and geometry consistency priors to the feature aggregation. Additionally, as shown in the second row of Figure 4, our method successfully reconstructs intricate textural

Method	DTU			Blender		
	PSNR \uparrow	SSIM \uparrow	LPIPS \downarrow	PSNR \uparrow	SSIM \uparrow	LPIPS \downarrow
PixelNeRF [63]	19.31	0.789	0.671	7.39	0.658	0.411
IBRNet [52]	26.04	0.917	0.190	22.44	0.874	0.195
MVSNeRF [4]	26.63	0.931	0.168	23.62	0.897	0.176
MatchNeRF [7]	26.91	0.934	0.159	23.20	0.897	0.164
EVE-NeRF	27.80	0.937	0.149	23.45	0.903	0.132

Table 2. Results for setting 2. Our method (EVE-NeRF) is trained on DTU and the Google Scanned Object dataset with 3 reference views. Our method outperforms on multiple metrics with other few-shot generalizable neural rendering methods.

details.

Setting 2: Few-Shot NVS. To compare with few-shot generalizable neural rendering methods [4, 7], we conducted novel view synthesis experiments with $M = 3$ input views in both training and evaluating, following the MVSNeRF [4] setting. We split the DTU [20] dataset into 88 scenes for training and 16 scenes for testing, following the methodology of prior works. Additionally, we also conducted training on the Google Scanned Object dataset. As shown in Table 2, we performed a quantitative comparison with 4 few-shot generalizable neural rendering methods [4, 7, 52, 63] on the DTU testing dataset and Blender. With only 3 source views input for setting, our model still achieves good performance. Our method outperforms SOTA methods MatchNeRF [7] by 2.19% \uparrow PSNR, 13.0% \downarrow LPIPS in 2 evaluating dataset evenly. Please refer to Appendix D.1 for the qualitative comparison for setting 2.

Efficiency Comparison. As shown in the Tab 3, our method is compared with GNT [49] and GPNR [43] in terms of efficiency. We perform the testing of setting 1 in LLFF dataset. The result illustrates that our method not only requires less memory and faster rendering time per-image, but also has a higher PSNR for novel view synthesis.

5.3. Ablation Studies

To evaluate the significance of our contributions, we conducted extensive ablation experiments. We trained on setting 1 and tested on the Real Forward-Facing dataset. For efficiency, in both the training and testing datasets we set the resolutions of images reduced by half in both the training and testing datasets, resulting in the resolution of 504×378 for the Real Forward-Facing dataset.

Only view/epipolar transformer. In this ablation, we maintain only the view/epipolar transformers and NeRF decoder. As can be observed in the first row of Table 4, using only view/epipolar transformer reduces PSNR by 6.21% compared to EVE-NeRF due to the limitations of view/epipolar aggregation in only one dimension.

Model	Storage \downarrow	Time \downarrow	PSNR \uparrow
GNT [49]	112MB	208s	25.35
GPNR [43]	110MB	12min	25.53
EVE-NeRF	53.8MB	194s	27.16

Table 3. Efficiency comparison results in LLFF dataset on the same RTX4090. Our method requires less storage, shorter rendering time per new view synthesis, and higher quality reconstruction compared to GNT [49] and GPNR [43].

Model	PSNR \uparrow	SSIM \uparrow	LPIPS \downarrow
only view transformer	25.03	0.886	0.132
+ along-epipolar perception	25.48	0.892	0.128
only epipolar transformer	25.02	0.879	0.147
+ multi-view calibration	25.17	0.883	0.141
naïve dual transformer	25.66	0.890	0.128
+ cross-attention interaction	25.85	0.896	0.120
EVE-NeRF	26.69	0.913	0.102

Table 4. Ablations. The ablation study was conducted by training in a low-resolution setting 1 and testing on LLFF dataset with the resolution of 504×378 .

Along-epipolar perception. Compared with only view transformer, we retained view transformer with along-epipolar perception in this ablation. As shown in Table 4 and Figure 3, using along-epipolar perception would increase PSNR by 1.80%. The appearance continuity prior provided by along-epipolar perception compensates for the missing epipolar information in the pure view aggregation model.

Multi-view calibration. Similarly, against to only epipolar transformer, we kept the epipolar transformer with multi-view calibration. As can be observed in Table 4 and Figure 3, adopting multi-view calibration would improve the performance of generalizable rendering. It verifies that the multi-view calibration can enhance epipolar aggregating ability via geometry prior.

Naïve dual network architecture. To validate the effectiveness of our proposed entangled aggregation, we compared our EVE-NeRF with two other architectures: the Naïve Dual Transformer and the Dual Transformer with Cross-Attention Interaction (see Appendix A.4). Our method outperforms both naïve dual transformer by 3.85% and 3.14% PSNR. Visualization results are provided in Appendix D.3. It is demonstrated that our proposed EVE-NeRF has a robust ability to fusing view and epipolar pattern.

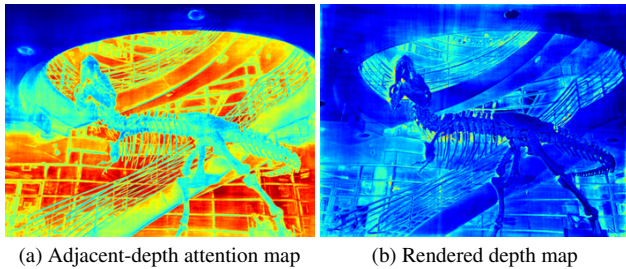


Figure 5. Visualizations of adjacent-depth attention map and rendered depth map. By capitalizing on the appearance continuity prior, adjacent-depth attention boosts the coherence in depth map.

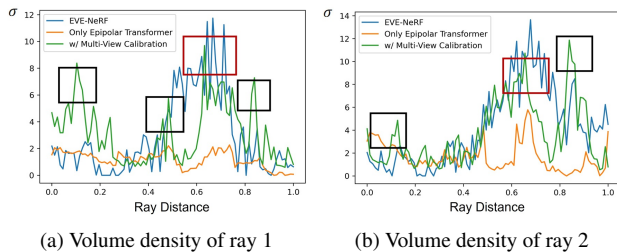


Figure 6. Line charts of the volume density from a novel view of the Chair [35]. Red boxes represent correct peaks and black boxes represent abnormal peaks. Multi-View Calibration learns more complex signals but with more noise. Volume densities predicted by EVE-NeRF are more likely to have a single-peak distribution.

5.4. Visualization on Entangled Information Interaction

To further validate the entangled information interaction module’s ability of providing the *de facto* appearance continuity prior and geometry consistency prior, we visualize and analyze the importance weights predicted by the along-epipolar perception and multi-view calibration.

The along-epipolar perception provides appearance continuity prior and regresses the importance weights for the target ray’s sampled depths. Specifically, we obtain a depth map by multiplying the depth weights with the marching distance along the ray. As shown in Figure 5, the adjacent-depth attention map demonstrates a more coherent character, indicating that the along-epipolar perception provides beneficial appearance consistency prior.

The multi-view calibration provides geometry consistency prior and predicts the importance weights for the source views. As shown in Figure 6, we visualize two line charts of the point density. Multi-view calibration learns more complex light signals, but with a multi-peaked distribution. EVE-NeRF predicts point density distributions with distinct peaks and reduced noise in the light signals. As shown in Figure 7, the view transformer and the multi-view calibration correctly predict the correspondence between the target pixel and the source views, such as the back of the

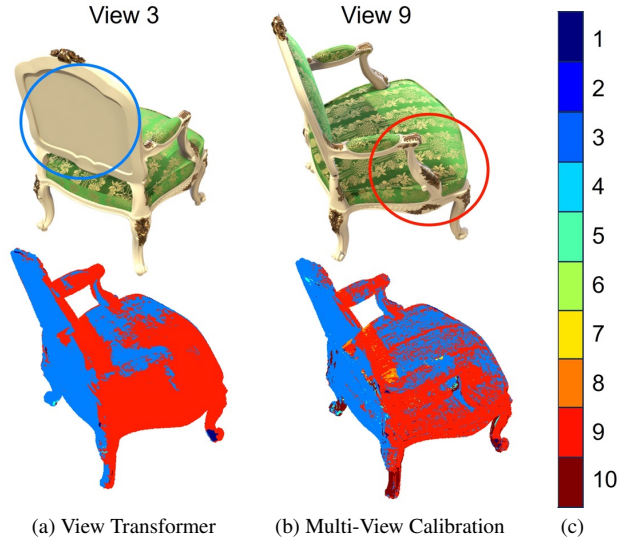


Figure 7. Each color represents the source view ID corresponding to the maximum weight for the target pixel. Both the view transformer and the multi-view calibration have successfully learned the cross-view information from the source views.

chair. Furthermore, both methods predict that the pixels in the upper right part of the chair correspond to source view 3, where the upper right part of the chair is occluded. We believe that EVE-NeRF learns about the awareness of visibility, even when the target pixel is occluded.

6. Conclusion

We propose a new Generalizable NeRF named EVE-NeRF that aggregates cross-view and along-epipolar information in an entangled manner. The core of EVE-NeRF consists of our new proposed View Epipolar Interaction Module (VEI) and Epipolar View Interaction Module (EVI) that are organized alternately. VEI and EVI can project the scene-invariant appearance continuity and geometry consistency priors, which serve to offset information losses that typically arises from single-dimensional aggregation operations. We demonstrate the superiority of our method in both generalizable and few-shot NVS settings compared with the state-of-the-art methods. Additionally, extensive ablation studies confirm that VEI and EVI can enhance information interaction across view and epipolar dimensions to yield better generalizable 3D representation.

Acknowledgement

This work is supported by the National Natural Science Foundation of China (62293554, 62206249, U2336212), the Natural Science Foundation of Zhejiang Province, China (LZ24F020002), Young Elite Scientists Sponsorship Program by CAST (2023QNRC001).

References

- [1] Jonathan T Barron, Ben Mildenhall, Matthew Tancik, Peter Hedman, Ricardo Martin-Brualla, and Pratul P Srinivasan. Mip-nerf: A multiscale representation for anti-aliasing neural radiance fields. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 5855–5864, 2021. [1](#), [2](#)
- [2] Jonathan T Barron, Ben Mildenhall, Dor Verbin, Pratul P Srinivasan, and Peter Hedman. Mip-nerf 360: Unbounded anti-aliased neural radiance fields. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 5470–5479, 2022. [1](#)
- [3] Jonathan T Barron, Ben Mildenhall, Dor Verbin, Pratul P Srinivasan, and Peter Hedman. Zip-nerf: Anti-aliased grid-based neural radiance fields. *arXiv preprint arXiv:2304.06706*, 2023. [2](#)
- [4] Anpei Chen, Zexiang Xu, Fuqiang Zhao, Xiaoshuai Zhang, Fanbo Xiang, Jingyi Yu, and Hao Su. Mvsnerf: Fast generalizable radiance field reconstruction from multi-view stereo. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 14124–14133, 2021. [2](#), [3](#), [5](#), [7](#), [14](#), [15](#)
- [5] Anpei Chen, Zexiang Xu, Andreas Geiger, Jingyi Yu, and Hao Su. Tensorf: Tensorial radiance fields. In *European Conference on Computer Vision*, pages 333–350. Springer, 2022. [2](#)
- [6] Chun-Fu Richard Chen, Quanfu Fan, and Rameswar Panda. Crossvit: Cross-attention multi-scale vision transformer for image classification. In *Proceedings of the IEEE/CVF international conference on computer vision*, pages 357–366, 2021. [13](#), [14](#)
- [7] Yuedong Chen, Haofei Xu, Qianyi Wu, Chuanxia Zheng, Tat-Jen Cham, and Jianfei Cai. Explicit correspondence matching for generalizable neural radiance fields. *arXiv preprint arXiv:2304.12294*, 2023. [2](#), [7](#), [14](#), [15](#)
- [8] Zhiqin Chen, Thomas Funkhouser, Peter Hedman, and Andrea Tagliasacchi. Mobilenerf: Exploiting the polygon rasterization pipeline for efficient neural field rendering on mobile architectures. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 16569–16578, 2023. [2](#), [16](#)
- [9] Zheng Chen, Yulun Zhang, Jinjin Gu, Linghe Kong, Xiaokang Yang, and Fisher Yu. Dual aggregation transformer for image super-resolution. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 12312–12321, 2023. [13](#), [14](#)
- [10] Julian Chibane, Aayush Bansal, Verica Lazova, and Gerard Pons-Moll. Stereo radiance fields (srf): Learning view synthesis for sparse views of novel scenes. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 7911–7920, 2021. [2](#)
- [11] Wenyan Cong, Hanxue Liang, Peihao Wang, Zhiwen Fan, Tianlong Chen, Mukund Varma, Yi Wang, and Zhangyang Wang. Enhancing nerf akin to enhancing llms: Generalizable nerf transformer with mixture-of-view-experts. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 3193–3204, 2023. [2](#), [3](#)
- [12] Laura Downs, Anthony Francis, Nate Koenig, Brandon Kinman, Ryan Hickman, Krista Reymann, Thomas B McHugh, and Vincent Vanhoucke. Google scanned objects: A high-quality dataset of 3d scanned household items. In *2022 International Conference on Robotics and Automation (ICRA)*, pages 2553–2560. IEEE, 2022. [6](#)
- [13] Tuo Feng, Wenguan Wang, Xiaohan Wang, Yi Yang, and Qinghua Zheng. Clustering based point cloud representation learning for 3d analysis. In *CVPR*, pages 8283–8294, 2023. [2](#)
- [14] John Flynn, Michael Broxton, Paul Debevec, Matthew DuVall, Graham Fyffe, Ryan Overbeck, Noah Snavely, and Richard Tucker. Deepview: View synthesis with learned gradient descent. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 2367–2376, 2019. [6](#)
- [15] Chen Gao, Ayush Saraf, Johannes Kopf, and Jia-Bin Huang. Dynamic view synthesis from dynamic monocular video. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 5712–5721, 2021. [2](#)
- [16] Qingji Guan, Yaping Huang, Yawei Luo, Ping Liu, Mingliang Xu, and Yi Yang. Discriminative feature learning for thorax disease classification in chest x-ray images. *IEEE Transactions on Image Processing*, 30:2476–2487, 2021. [1](#)
- [17] Peter Hedman, Pratul P Srinivasan, Ben Mildenhall, Jonathan T Barron, and Paul Debevec. Baking neural radiance fields for real-time view synthesis. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 5875–5884, 2021. [2](#), [16](#)
- [18] Wenbo Hu, Yuling Wang, Lin Ma, Bangbang Yang, Lin Gao, Xiao Liu, and Yuewen Ma. Tri-miprf: Tri-mip representation for efficient anti-aliasing neural radiance fields. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 19774–19783, 2023. [2](#)
- [19] Xin Huang, Qi Zhang, Ying Feng, Xiaoyu Li, Xuan Wang, and Qing Wang. Local implicit ray function for generalizable radiance field representation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 97–107, 2023. [2](#), [3](#)
- [20] Rasmus Jensen, Anders Dahl, George Vogiatzis, Engin Tola, and Henrik Aanæs. Large scale multi-view stereopsis evaluation. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 406–413, 2014. [7](#)
- [21] Mohammad Mahdi Johari, Yann Lepoittevin, and François Fleuret. Geonerf: Generalizing nerf with geometry priors. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 18365–18375, 2022. [2](#), [3](#), [4](#), [12](#)
- [22] Jonáš Kulháněk, Erik Derner, Torsten Sattler, and Robert Babuška. Viewformer: Nerf-free neural rendering from few images using transformers. In *European Conference on Computer Vision*, pages 198–216. Springer, 2022. [2](#)
- [23] Jiaxin Li, Zijian Feng, Qi She, Henghui Ding, Changhu Wang, and Gim Hee Lee. Mine: Towards continuous depth mpi with nerf for novel view synthesis. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 12578–12588, 2021. [3](#)
- [24] Steven Liu, Xiuming Zhang, Zhoulong Zhang, Richard Zhang, Jun-Yan Zhu, and Bryan Russell. Editing conditional radiance fields. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 12599–12609, 2021. [1](#)

- ance fields. In *Proceedings of the IEEE/CVF international conference on computer vision*, pages 5773–5783, 2021. [2](#)
- [25] Yuan Liu, Sida Peng, Lingjie Liu, Qianqian Wang, Peng Wang, Christian Theobalt, Xiaowei Zhou, and Wenping Wang. Neural rays for occlusion-aware image-based rendering. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 7824–7833, 2022. [6](#)
- [26] Keyang Luo, Tao Guan, Lili Ju, Haipeng Huang, and Yawei Luo. P-mvsnet: Learning patch-wise matching confidence aggregation for multi-view stereo. In *ICCV*, pages 10452–10461, 2019. [1](#)
- [27] Keyang Luo, Tao Guan, Lili Ju, Yuesong Wang, Zhuo Chen, and Yawei Luo. Attention-aware multi-view stereo. In *CVPR*, pages 1590–1599, 2020. [1](#)
- [28] Yawei Luo and Yi Yang. Large language model and domain-specific model collaboration for smart education. *FITEE*, 2024. [1](#)
- [29] Yawei Luo, Ping Liu, Tao Guan, Junqing Yu, and Yi Yang. Significance-aware information bottleneck for domain adaptive semantic segmentation. In *ICCV*, pages 6778–6787, 2019. [1](#)
- [30] Yawei Luo, Liang Zheng, Tao Guan, Junqing Yu, and Yi Yang. Taking a closer look at domain shift: Category-level adversaries for semantics consistent domain adaptation. In *CVPR*, pages 2507–2516, 2019.
- [31] Yawei Luo, Ping Liu, Tao Guan, Junqing Yu, and Yi Yang. Adversarial style mining for one-shot unsupervised domain adaptation. In *NeurIPS*, pages 20612–20623, 2020.
- [32] Yawei Luo, Ping Liu, Liang Zheng, Tao Guan, Junqing Yu, and Yi Yang. Category-level adversarial adaptation for semantic segmentation using purified features. *T-PAMI*, 2021. [1](#)
- [33] Shaojie Ma, Yawei Luo, and Yi Yang. Personas-based student grouping using reinforcement learning and linear programming. *Knowledge-Based Systems*, page 111071, 2023. [1](#)
- [34] Ben Mildenhall, Pratul P Srinivasan, Rodrigo Ortiz-Cayon, Nima Khademi Kalantari, Ravi Ramamoorthi, Ren Ng, and Abhishek Kar. Local light field fusion: Practical view synthesis with prescriptive sampling guidelines. *ACM Transactions on Graphics (TOG)*, 38(4):1–14, 2019. [6](#)
- [35] Ben Mildenhall, Pratul P Srinivasan, Matthew Tancik, Jonathan T Barron, Ravi Ramamoorthi, and Ren Ng. Nerf: Representing scenes as neural radiance fields for view synthesis. *Communications of the ACM*, 65(1):99–106, 2021. [1](#), [2](#), [5](#), [6](#), [8](#), [15](#), [16](#)
- [36] Amirali Molaei, Amirhossein Aminimehr, Armin Tavakoli, Amirhossein Kazerouni, Bobby Azad, Reza Azad, and Dorit Merhof. Implicit neural representation in medical imaging: A comparative survey. In *ICCV*, pages 2381–2391, 2023. [1](#)
- [37] Keunhong Park, Utkarsh Sinha, Jonathan T Barron, Sofien Bouaziz, Dan B Goldman, Steven M Seitz, and Ricardo Martin-Brualla. Nerfies: Deformable neural radiance fields. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 5865–5874, 2021. [2](#)
- [38] Albert Pumarola, Enric Corona, Gerard Pons-Moll, and Francesc Moreno-Noguer. D-nerf: Neural radiance fields for dynamic scenes. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 10318–10327, 2021. [2](#)
- [39] Christian Reiser, Rick Szeliski, Dor Verbin, Pratul Srinivasan, Ben Mildenhall, Andreas Geiger, Jon Barron, and Peter Hedman. Merf: Memory-efficient radiance fields for real-time view synthesis in unbounded scenes. *ACM Transactions on Graphics (TOG)*, 42(4):1–12, 2023. [2](#)
- [40] Jeremy Reizenstein, Roman Shapovalov, Philipp Henzler, Luca Sbordone, Patrick Labatut, and David Novotny. Common objects in 3d: Large-scale learning and evaluation of real-life 3d category reconstruction. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 10901–10911, 2021. [2](#), [3](#)
- [41] Olaf Ronneberger, Philipp Fischer, and Thomas Brox. U-net: Convolutional networks for biomedical image segmentation. In *Medical Image Computing and Computer-Assisted Intervention—MICCAI 2015: 18th International Conference, Munich, Germany, October 5-9, 2015, Proceedings, Part III 18*, pages 234–241. Springer, 2015. [3](#)
- [42] Karen Simonyan and Andrew Zisserman. Two-stream convolutional networks for action recognition in videos. *Advances in neural information processing systems*, 27, 2014. [13](#)
- [43] Mohammed Suhail, Carlos Esteves, Leonid Sigal, and Ameesh Makadia. Generalizable patch-based neural rendering. In *European Conference on Computer Vision*, pages 156–174. Springer, 2022. [1](#), [2](#), [3](#), [6](#), [7](#)
- [44] Mohammed Suhail, Carlos Esteves, Leonid Sigal, and Ameesh Makadia. Light field neural rendering. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 8269–8279, 2022. [2](#), [15](#), [16](#)
- [45] Cheng Sun, Min Sun, and Hwann-Tzong Chen. Direct voxel grid optimization: Super-fast convergence for radiance fields reconstruction. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 5459–5469, 2022. [2](#)
- [46] Jingxiang Sun, Xuan Wang, Yong Zhang, Xiaoyu Li, Qi Zhang, Yebin Liu, and Jue Wang. Fenerf: Face editing in neural radiance fields. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 7672–7682, 2022. [2](#)
- [47] Fengrui Tian, Shaoyi Du, and Yueqi Duan. Mononerf: Learning a generalizable dynamic radiance field from monocular videos. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 17903–17913, 2023. [2](#)
- [48] Edgar Tretschk, Ayush Tewari, Vladislav Golyanik, Michael Zollhöfer, Christoph Lassner, and Christian Theobalt. Non-rigid neural radiance fields: Reconstruction and novel view synthesis of a dynamic scene from monocular video. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 12959–12970, 2021. [2](#)
- [49] Mukund Varma, Peihao Wang, Xuxi Chen, Tianlong Chen, Subhashini Venugopalan, and Zhangyang Wang. Is attention all that nerf needs? In *The Eleventh International Conference on Learning Representations*, 2022. [1](#), [2](#), [3](#), [5](#), [6](#), [7](#), [12](#)
- [50] Can Wang, Menglei Chai, Mingming He, Dongdong Chen, and Jing Liao. Clip-nerf: Text-and-image driven manipulation

- of neural radiance fields. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 3835–3844, 2022. [2](#)
- [51] Dan Wang, Xinrui Cui, Septimiu Salcudean, and Z Jane Wang. Generalizable neural radiance fields for novel view synthesis with transformer. *arXiv preprint arXiv:2206.05375*, 2022. [2](#), [3](#)
- [52] Qianqian Wang, Zhicheng Wang, Kyle Genova, Pratul P Srinivasan, Howard Zhou, Jonathan T Barron, Ricardo Martin-Brualla, Noah Snavely, and Thomas Funkhouser. Ibrnet: Learning multi-view image-based rendering. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 4690–4699, 2021. [2](#), [3](#), [5](#), [6](#), [7](#), [12](#)
- [53] Rui Wang, Dongdong Chen, Zuxuan Wu, Yinpeng Chen, Xiyang Dai, Mengchen Liu, Yu-Gang Jiang, Luwei Zhou, and Lu Yuan. Bevt: Bert pretraining of video transformers. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 14733–14743, 2022. [13](#)
- [54] Yuesong Wang, Tao Guan, Zhuo Chen, Yawei Luo, Keyang Luo, and Lili Ju. Mesh-guided multi-view stereo with pyramid architecture. In *CVPR*, pages 2039–2048, 2020. [1](#)
- [55] Yuesong Wang, Zhaojie Zeng, Tao Guan, Wei Yang, Zhuo Chen, Wenkai Liu, Luoyuan Xu, and Yawei Luo. Adaptive patch deformation for textureless-resilient multi-view stereo. In *CVPR*, pages 1621–1630, 2023. [1](#)
- [56] Suttisak Wizadwongsa, Pakkapon Phongthawee, Jiraphon Yenphraphai, and Supasorn Suwajanakorn. Nex: Real-time view synthesis with neural basis expansion. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 8534–8543, 2021. [6](#), [15](#), [16](#)
- [57] Fanbo Xiang, Zexiang Xu, Milos Hasan, Yannick Hold-Geoffroy, Kalyan Sunkavalli, and Hao Su. Neutex: Neural texture mapping for volumetric neural rendering. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 7119–7128, 2021. [2](#)
- [58] Tianhan Xu and Tatsuya Harada. Deforming radiance fields with cages. In *European Conference on Computer Vision*, pages 159–175. Springer, 2022. [2](#)
- [59] Hao Yang, Lanqing Hong, Aoxue Li, Tianyang Hu, Zhenguo Li, Gim Hee Lee, and Liwei Wang. Contranerf: Generalizable neural radiance fields for synthetic-to-real novel view synthesis via contrastive learning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 16508–16517, 2023. [3](#)
- [60] Junbo Yin, Jin Fang, Dingfu Zhou, Liangjun Zhang, Cheng-Zhong Xu, Jianbing Shen, and Wenguan Wang. Semi-supervised 3d object detection with proficient teachers. In *ECCV*, pages 727–743, 2022. [2](#)
- [61] Junbo Yin, Dingfu Zhou, Liangjun Zhang, Jin Fang, Cheng-Zhong Xu, Jianbing Shen, and Wenguan Wang. Proposal-contrast: Unsupervised pre-training for lidar-based 3d object detection. In *ECCV*, pages 17–33, 2022. [2](#)
- [62] Alex Yu, Ruilong Li, Matthew Tancik, Hao Li, Ren Ng, and Angjoo Kanazawa. Plenotrees for real-time rendering of neural radiance fields. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 5752–5761, 2021. [1](#), [2](#)
- [63] Alex Yu, Vickie Ye, Matthew Tancik, and Angjoo Kanazawa. pixelnerf: Neural radiance fields from one or few images. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 4578–4587, 2021. [2](#), [3](#), [5](#), [6](#), [7](#)
- [64] Kai Zhang, Gernot Riegler, Noah Snavely, and Vladlen Koltun. Nerf++: Analyzing and improving neural radiance fields. *arXiv preprint arXiv:2010.07492*, 2020. [2](#)
- [65] Zerong Zheng, Han Huang, Tao Yu, Hongwen Zhang, Yandong Guo, and Yebin Liu. Structured local radiance fields for human avatar modeling. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 15893–15903, 2022. [2](#)
- [66] Tinghui Zhou, Richard Tucker, John Flynn, Graham Fyffe, and Noah Snavely. Stereo magnification: Learning view synthesis using multiplane images. *arXiv preprint arXiv:1805.09817*, 2018. [6](#)