

# Depth Field Networks for Generalizable Multi-view Scene Representation

Vitor Guizilini<sup>1\*</sup>, Igor Vasiljevic<sup>1\*</sup>, Jiading Fang<sup>2\*</sup>, Rares Ambrus<sup>1</sup>, Greg Shakhnarovich<sup>2</sup>, Matthew R. Walter<sup>2</sup>, and Adrien Gaidon<sup>1</sup>

<sup>1</sup> Toyota Research Institute, Los Altos, CA

<sup>2</sup> Toyota Technological Institute at Chicago, Chicago, IL




Fig. 1: Our **Depth Field Networks (DeFiNe)** achieve state of the art in multi-view depth estimation, while also enabling predictions from arbitrary viewpoints.

**Abstract.** Modern 3D computer vision leverages learning to boost geometric reasoning, mapping image data to classical structures such as cost volumes or epipolar constraints to improve matching. These architectures are specialized according to the particular problem, and thus require significant task-specific tuning, often leading to poor domain generalization performance. Recently, generalist Transformer architectures have achieved impressive results in tasks such as optical flow and depth estimation by encoding geometric priors as inputs rather than as enforced constraints. In this paper, we extend this idea and propose to learn an implicit, multi-view consistent scene representation, introducing a series of 3D data augmentation techniques as a geometric inductive prior to increase view diversity. We also show that introducing view synthesis as an auxiliary task further improves depth estimation. Our Depth Field Networks (DeFiNe) achieve state-of-the-art results in stereo and video depth estimation without explicit geometric constraints, and improve on zero-shot domain generalization by a wide margin. Project page: <https://sites.google.com/view/tri-define>.

**Keywords:** Multi-view Depth estimation; Representation learning

## 1 Introduction

Estimating 3D structure from a pair of images is a cornerstone problem of computer vision. Traditionally, this is treated as a correspondence problem, whereby one applies a homography to stereo-rectify the images and then matches pixels

\* Denotes equal contribution.

(or patches) along epipolar lines to obtain disparity estimates. Contemporary approaches to stereo are specialized variants of classical methods, relying on correspondences to compute cost volumes, epipolar losses, bundle adjustment objectives, or projective multi-view constraints, among others. These are either baked directly into the model architecture or enforced as part of the loss function.

Applying the principles of classical vision in this way has given rise to architectures that achieve state-of-the-art results on tasks such as stereo depth estimation [18,24], optical flow [46], and multi-view depth [45]. However, this success comes at a cost: each architecture is specialized and purpose-built for a single task. Great strides have been made to alleviate the dependence on strong geometric assumptions [12,50], and two recent trends allow us to *decouple* the task from the architecture: (i) implicit representations and (ii) generalist networks. Our work draws upon both of these directions.

Implicit representations of geometry and coordinate-based networks have recently achieved incredible popularity in the vision community. This direction is pioneered by work on neural radiance fields (NeRF) [29,54], where a point- and ray-based parameterization along with a volume rendering objective allow simple MLP-based networks to achieve state-of-the-art view synthesis results. Follow-up works extend this coordinate-based representation to the pixel domain [57], allowing predicted views to be conditioned on image features. The second trend in computer vision has been the use of generalist architectures, pioneered by Vision Transformers [5]. Emerging as an attention-based architecture for NLP, Transformers have been used for a diverse set of tasks, including depth estimation [23,32], optical flow [16], and image generation [8]. Transformers have also been applied to geometry-free view synthesis [35], demonstrating that attention can learn long-range correspondence between views for 2D-3D tasks. Representation Transformers (SRT) [36] use the Transformer encoder-decoder model to learn scene representations for view synthesis from sparse, high-baseline data with no geometric constraints. However, owing to the quadratic scaling of the self-attention module, experiments are limited to low-resolution images and require very long training periods.

To alleviate the quadratic complexity of self-attention, the Perceiver architecture [17] disentangles the dimensionality of the latent representation from that of the inputs by fixing the size of the latent representation. Perceiver IO [16] extends this architecture to allow for arbitrary outputs, with results on optical flow estimation that outperform traditional cost-volume based methods. Similarly, the recent Input-level Inductive Bias (IIB) architecture [56] uses image features and camera information as input to Perceiver IO to directly regress stereo depth, outperforming baselines that use explicit geometric constraints. Building upon these works, we propose to learn a *geometric scene representation* for depth synthesis from novel viewpoints, including estimation, interpolation, and extrapolation. We expand the IIB framework to the scene representation setting, taking sequences of images and predicting a consistent multi-view latent representation suitable for different downstream tasks. Taking advantage of the query-based nature of the Perceiver IO architecture, we propose a series of 3D augmentations

that increase viewpoint density and diversity during training, thus encouraging (rather than enforcing) multi-view consistency. Furthermore, we show that the introduction of view synthesis as an auxiliary task, decoded from the same latent representation, improves depth estimation without additional ground truth.

We test our model on the popular ScanNet benchmark [3], achieving state-of-the-art real-time results for stereo depth estimation and competitive results for video depth estimation, without relying on memory- or compute-intensive operations such as cost volume aggregation and test-time optimization. We show that our 3D augmentations lead to significant improvements over baselines that are limited to the viewpoint diversity of training data. Furthermore, our zero-shot transfer results from ScanNet to 7-Scenes [40] improve the state-of-the-art by a large margin, demonstrating that our method generalizes better than specialized architectures, which suffer from poor performance on out-of-domain data. Our contributions are summarized as follows:

- We use a generalist Transformer-based architecture to learn a depth estimator from an arbitrary number of posed images. In this setting, we (i) **propose a series of 3D augmentations** that improve the geometric consistency of our learned latent representation; and (ii) show that **jointly learning view synthesis as an auxiliary task improves depth estimation**.
- Our Depth Field Networks (DeFiNe) not only achieve **state-of-the-art stereo depth estimation results** on the widely used ScanNet dataset, but also exhibit superior generalization properties with **state-of-the-art results on zero-shot transfer to 7-Scenes**.
- DeFiNe also **enables depth estimation from arbitrary viewpoints**. We evaluate this novel generalization capability in the context of *interpolation* (between timesteps), and *extrapolation* (future timesteps).

## 2 Related Work

**Monocular Depth Estimation.** Supervised depth estimation—the task of estimating per-pixel depth given an RGB image and a corresponding ground-truth depth map—dates back to the pioneering work of Saxena et al. [37]. Since then, deep learning-based architectures designed for supervised monocular depth estimation have become increasingly sophisticated [6,7,9,21,22], generally offering improvements over the standard encoder-decoder convolutional architecture. Self-supervised methods provide an alternative to those that rely on ground-truth depth maps at training time, and are able to take advantage of the new availability of large-scale video datasets. Early self-supervised methods relied on stereo data [10], and then progressed to fully monocular video sequences [59], with increasingly sophisticated losses [41] and architectures [11,13,52].

**Multi-view Stereo.** Traditional multi-view stereo approaches have dominated even in the deep learning era. COLMAP [38] remains the standard framework for structure-from-motion, incorporating sophisticated bundle adjustment and keypoint refinement procedures, at the cost of speed. With the goal of producing closer to real-time estimates, multi-view stereo learning approaches adapt

traditional cost volume-based approaches to stereo [18,2] and multi-view [55,15] depth estimation, often relying on known extrinsics to warp views into the frame of the reference camera. Recently, iterative refinement approaches that employ recurrent neural networks have made impressive strides in optical flow estimation [46]. Follow-on work applies this general recurrent correspondence architecture to stereo depth [24], scene-flow [48], and even SLAM [47]. While their results are impressive, recurrent neural networks can be difficult to train, and test-time optimization increases inference time over a single forward pass.


Recently, Transformer-based architectures [51] have replaced CNNs in many geometric estimation tasks. The Stereo Transformer [23] architecture replaces cost volumes with a correspondence approach inspired by sequence-to-sequence modeling. The Perceiver IO [16] architecture constitutes a large departure from cost volumes and geometric losses. For the task of optical flow, Perceiver IO feeds positionally encoded images through a Transformer [17], rather than using a cost volume for processing. IIB [56] adapts the Perceiver IO architecture to generalized stereo estimation, proposing a novel epipolar parameterization as an additional input-level inductive bias. Building upon this baseline, we propose a series of geometry-preserving 3D data augmentation techniques designed to promote the learning of a *geometrically-consistent latent scene representation*. We also introduce novel view synthesis as an auxiliary task to depth estimation, decoded from the same latent space. Our video-based representation (aided by our 3D augmentations) allows us to generalize to novel viewpoints, rather than be restricted to the stereo setting.

**Video Depth Estimation.** Video and stereo depth estimation methods generally produce monocular depth estimates at test time. ManyDepth [52] combines a monocular depth framework with multi-view stereo, aggregating predictions in a cost volume and thus enabling multi-frame inference at test-time. Recent methods accumulate predictions at train and test time, either with generalized stereo [49] or with sequence data [58]. DeepV2D [45] incorporates a cost-volume based multi-view stereo approach with an incremental pose estimator to iteratively improve depth and pose estimates at train and test time.


Another line of work draws on the availability of monocular depth networks that perform accurate but *multi-view inconsistent* estimates at test time [28]. In this setting, additional geometric constraints are enforced to finetune the network and improve multi-view consistency through epipolar constraints. Consistent Video Depth Estimation [28] refines COLMAP [38] results with a monocular depth network constrained to be multi-view consistent. Subsequent work jointly optimizes depth and pose for added robustness to challenging scenarios with poor calibration [19]. A recent framework incorporates many architectural elements of prior work into a Transformer-based architecture that takes video data as input for multi-view depth [26]. NeuralRecon [43] moves beyond depth-based architectures to learn Truncated Signed Distance Field (TSDF) volumes as a way to improve surface consistency.

**Novel View Synthesis.** Since the emergence of neural radiance fields [29], implicit representations and volume rendering have emerged as the *de facto*






(a) Architecture overview.



(b) Image embeddings.



(c) Camera embeddings.

Fig. 2: **Overview of our proposed DeFiNe architecture**, and the embeddings used to encode and decode information for depth and view synthesis.

standard for view synthesis. They parameterize viewing rays and points with positional encoding, and need to be re-trained on a scene-by-scene basis. Many recent improvements leverage depth supervision to improve view synthesis in a volume rendering framework [1, 4, 34, 53, 60]. An alternative approach replaces volume rendering with a directly learned light field network [42], predicting color values directly from viewing rays. This is the approach we take when estimating the auxiliary view synthesis loss, due to its computational simplicity. Other works attempt to extend the NeRF approach beyond single scene models by incorporating learned *features*, enabling few-shot volume rendering [57]. Feature-based methods have also treated view synthesis as a sequence-learning task, such as the Scene Representation Transformer (SRT) architecture [36].

### 3 The DeFiNe Architecture

Our proposed DeFiNe architecture (Figure 2a) is designed with flexibility in mind, so data from different sources can be used as input and different output tasks can be estimated from the same latent space. Similar to Yifan et al. [56], we use Perceiver IO [17] as our general-purpose Transformer backbone. During the encoding stage, our model takes RGB images from calibrated cameras, with known intrinsics and relative poses. The architecture processes this information according to the modality into different pixel-wise embeddings that serve as input to our Perceiver IO backbone. This encoded information can be queried using only camera embeddings, producing estimates from arbitrary viewpoints.

#### 3.1 Perceiver IO

Perceiver IO [16] is a recent extension of the Perceiver [17] architecture. The Perceiver architecture alleviates one of the main weaknesses of Transformer-based

methods, namely the quadratic scaling of self-attention with input size. This is achieved by using a fixed-size  $N_l \times C_l$  latent representation  $\mathcal{R}$ , and learning to project high-dimensional  $N_e \times C_e$  encoding embeddings onto this latent space using cross-attention layers. Self-attention is performed in the lower-dimensional latent space, producing a *conditioned latent representation*  $\mathcal{R}_c$  that can be queried using  $N_d \times C_d$  decoding embeddings to generate estimates, again using cross-attention layers.

### 3.2 Input-Level Embeddings

**Image Embeddings (Figure 2b).** Input  $3 \times H \times W$  images are processed using a ResNet18 [14] encoder, producing a list of feature maps at increasingly lower resolutions and higher dimensionality. Feature maps at 1/4 the original resolution are concatenated with lower-resolution feature maps, after upsampling using bilinear interpolation. The resulting image embeddings are of shape  $H/4 \times W/4 \times 960$ , and are used in combination with the camera embeddings from each corresponding pixel (see below) to encode visual information.

**Camera Embeddings (Figure 2c).** These embeddings capture multi-view scene geometry (e.g., camera intrinsics and extrinsics) as additional inputs during the learning process. Let  $\mathbf{x}_{ij} = (u, v)$  be an image coordinate corresponding to pixel  $i$  in camera  $j$ , with assumed known pinhole  $3 \times 3$  intrinsics  $\mathbf{K}_j$  and  $4 \times 4$  transformation matrix  $T_j = \begin{bmatrix} \mathbf{R}_j & \mathbf{t}_j \\ \mathbf{0} & 1 \end{bmatrix}$  relative to a canonical camera  $T_0$ . Its origin  $\mathbf{o}_j$  and direction  $\mathbf{r}_{ij}$  are given by:

$$\mathbf{o}_j = -\mathbf{R}_j \mathbf{t}_j \quad , \quad \mathbf{r}_{ij} = (\mathbf{K}_j \mathbf{R}_j)^{-1} [u_{ij}, v_{ij}, 1]^T + \mathbf{t}_j \quad (1)$$

Note that this formulation differs from the standard convention [29], which does not consider the camera translation  $\mathbf{t}_j$  when generating viewing rays  $\mathbf{r}_{ij}$ . We ablate this variation in Table 1, showing that it leads to better performance for the task of depth estimation. These two vectors are then Fourier-encoded dimension-wise to produce higher-dimensional vectors, with a mapping of:

$$x \mapsto [x, \sin(f_1 \pi x), \cos(f_1 \pi x), \dots, \sin(f_K \pi x), \cos(f_K \pi x)]^T \quad (2)$$

where  $K$  is the number of Fourier frequencies used ( $K_o$  for the origin and  $K_r$  for the ray directions), equally spaced in the interval  $[1, \frac{\mu}{2}]$ . The resulting camera embedding is of dimensionality  $2(3(K_o+1)+3(K_r+1)) = 6(K_o + K_r + 2)$ . During the encoding stage, camera embeddings are produced per-pixel assuming a camera with 1/4 the original input resolution, resulting in a total of  $\frac{HW}{16}$  vectors. During the decoding stage, embeddings from cameras with arbitrary calibration (i.e., intrinsics and extrinsics) can be queried to produce virtual estimates.

### 3.3 Geometric 3D Augmentations

Data augmentation is a core component of deep learning pipelines [39] that improves model robustness by applying transformations to the training data




Fig. 3: **Geometric augmentations.** (a) Information from camera  $i$  is projected onto a virtual camera at  $T_v$ , creating additional supervision from other viewpoints. (b) Noise  $T'_0$  is introduced to the canonical camera at  $T_0$ , and then propagated to other cameras to preserve relative scene geometry.

consistent with the data distribution in order to introduce desired equivariant properties. In computer vision and depth estimation in particular, standard data augmentation techniques are usually constrained to the 2D space and include color jittering, flipping, rotation, cropping, and resizing [11,56]. Recent works have started looking into 3D augmentations [36] to improve robustness to errors in scene geometry in terms of camera localization (i.e., extrinsics) and parameters (i.e., intrinsics). Conversely, we are interested in *encoding* scene geometry at the input-level, so our architecture can learn a multi-view-consistent geometric latent scene representation. Therefore, in this section we propose a series of 3D augmentations to increase the number and diversity of training views while maintaining the spatial relationship between cameras, thus enforcing desired equivariant properties within this setting.

**Virtual Camera Projection.** One of the key properties of our architecture is that it enables querying from arbitrary viewpoints, since only camera information (viewing rays) is required at the decoding stage. When generating predictions from these novel viewpoints, the network creates *virtual* information consistent with the implicit structure of the learned latent scene representation, conditioned on information from the encoded views. We evaluate this capability in Section 4.5, showing superior performance relative to the explicit projection of information from encoded views. Here, we propose to leverage this property at training time as well, generating additional supervision in the form of *virtual cameras* with corresponding ground-truth RGB images and depth maps obtained by projecting available information onto these new viewpoints (Figure 3a). This novel augmentation technique forces the learned latent scene representation to be viewpoint-independent. Experiments show that this approach provides benefits in both the (a) stereo setting, with only two viewpoints; and (b) video setting, with similar viewpoints and a dominant camera direction.

From a practical perspective, virtual cameras are generated by adding translation noise  $\epsilon_v = [\epsilon_x, \epsilon_y, \epsilon_z]_v \sim \mathcal{N}(0, \sigma_v)$  to the pose of a camera  $i$ . The viewing angle is set to point towards the center  $\mathbf{c}_i$  of the pointcloud  $P_i$  generated by unprojecting information from the selected camera, which is also perturbed

by  $\epsilon_c = [\epsilon_x, \epsilon_y, \epsilon_z]_c \sim \mathcal{N}(0, \sigma_v)$ . When generating ground-truth information, we project the combined pointcloud from all available cameras onto these new viewpoints as a way to preserve full scene geometry. Furthermore, because the resulting RGB image and depth map will be sparse, we can improve efficiency by only querying at these specific locations.

**Canonical Jittering.** When operating in a multi-camera setting, it is standard practice to select one camera to be the *reference* camera, and position all other cameras relative to it [15]. One drawback of this convention is that one camera will always be at the same location (the origin of its own coordinate system) and will therefore produce the same camera embeddings, leading to overfitting. Intuitively, scene geometry should be invariant to the translation and rotation of the entire sensor suite. To enforce this property on our learned latent scene representation, we propose to inject some amount of noise to the canonical pose itself, so it is not located at the origin of the coordinate system. Note that this is different from methods that inject per-camera noise [31] with the goal of increasing robustness to localization errors. We only inject noise *once*, on the canonical camera, and propagate it to other cameras, so relative scene geometry is preserved within a translation and rotation offset (Figure 3b). However, this offset is reflected on the input-level embeddings produced by each camera, and thus forces the latent representation to be invariant to these transformations.

In order to perform canonical jittering, we randomly sample translation  $\epsilon_t = [\epsilon_x, \epsilon_y, \epsilon_z]^\top \sim \mathcal{N}(0, \sigma_t)$  and rotation  $\epsilon_r = [\epsilon_\phi, \epsilon_\theta, \epsilon_\psi]^\top \sim \mathcal{N}(0, \sigma_r)$  errors from zero-mean normal distributions with pre-determined standard deviations. Represented as Euler angles, we convert each set of rotation errors to a  $3 \times 3$  rotation matrix  $\mathbf{R}_r$ . We then use the rotation matrix and translation error to create a jittered canonical transformation matrix  $T'_0 = \begin{bmatrix} \mathbf{R}_r & \epsilon_t \\ \mathbf{0} & 1 \end{bmatrix}$  that is then propagated to all other  $N$  cameras, such that  $T'_i = T'_0 \cdot T_i, \forall i \in \{1, \dots, N-1\}$ .

**Canonical Randomization.** As an extension to canonical jittering, we also introduce canonical randomization to encourage generalization to different relative camera configurations, while still preserving scene geometry. Assuming  $N$  cameras, we randomly select  $o \in \{0, \dots, N-1\}$  as the canonical index. We then compute the relative transformation matrix  $T'_i$  given world-frame transformation matrix  $T_i$  as  $T'_i = T_i \cdot T_o^{-1} \forall i \in \{0, \dots, N-1\}$ .

### 3.4 Decoders

We use task-specific decoders, each consisting of one cross-attention layer between the  $N_d \times C_d$  queries and the  $N_l \times C_l$  conditioned latent representation  $\mathcal{R}_c$ , followed by a linear layer that creates an output of size  $N_d \times C_o$ , and a sigmoid activation function  $\sigma(x) = \frac{1}{1+e^{-x}}$  to produce values in the interval  $[0, 1]$ . We set  $C_o^d = 1$  for depth estimation and  $C_o^s = 3$  for view synthesis. Depth estimates are scaled to lie within the range  $[d_{\min}, d_{\max}]$ . Note that other decoders can be incorporated with DeFiNe without any modification to the underlying architecture, enabling the generation of multi-task estimates from arbitrary viewpoints.

### 3.5 Losses

We use an L1-log loss  $\mathcal{L}_d = \|\log(d_{ij}) - \log(\hat{d}_{ij})\|_1$  to supervise depth estimation, where  $\hat{d}_{ij}$  and  $d_{ij}$  are depth estimates and ground truth, respectively, for pixel  $j$  at camera  $i$ . For view synthesis, we use an L2 loss  $\mathcal{L}_s = \|\mathbf{p}_{ij} - \hat{\mathbf{p}}_{ij}\|^2$ , where  $\hat{\mathbf{p}}_{ij}$  and  $\mathbf{p}_{ij}$  are RGB estimates and ground truth, respectively, for pixel  $j$  at camera  $i$ . We use a weight coefficient  $\lambda_s$  to balance these two losses, and another  $\lambda_v$  to balance losses from available and virtual cameras. The final loss is of the form:

$$\mathcal{L} = \mathcal{L}_d + \lambda_s \mathcal{L}_s + \lambda_v (\mathcal{L}_{d,v} + \lambda_s (\mathcal{L}_{s,v})) \quad (3)$$

Note that because our architecture enables querying at specific image coordinates, we can improve efficiency at training time by not computing estimates for pixels without ground truth (e.g., sparse depth maps or virtual cameras).

## 4 Experiments

### 4.1 Datasets

**ScanNet** [3] We evaluate our DeFiNe for both *stereo* and *video* depth estimation using ScanNet, an RGB-D video dataset that contains 2.5 million views from around 1500 scenes. For the stereo experiments, we follow the same setting as Kusupati et al. [20]: we downsample scenes by a factor of 20 and use a custom split to create stereo pairs, resulting in 94212 training and 7517 test samples. For the video experiments, we follow the evaluation protocol of Teed et al. [45], with a total of 1405 scenes for training. For the test set, we use a custom split to select 2000 samples from 90 scenes not covered in the training set. Each training sample includes a target frame and a context of  $[-3, 3]$  frames with stride 3. Each test sample includes a pair of frames, with a context of  $[-3, 3]$  relative to the first frame of the pair with stride 3.


**7-Scenes** [40] We also evaluate on the test split of 7-Scenes to measure zero-shot cross-dataset performance. Collected using KinectFusion [30], the dataset consists of  $640 \times 480$  images in 7 settings, with a variable number of scenes in each setting. There are 500–1000 images in each scene. We follow the evaluation protocol of Sun et al. [43], median-scaling predictions using ground-truth information before evaluation.

### 4.2 Stereo Depth Estimation

To test the benefits our proposed geometric 3D augmentation procedures over the IIB [56] baseline, we first evaluate DeFiNe on the task of stereo depth estimation. Here, because each sample provides minimal information about the scene (i.e., only two frames), the introduction of additional virtual supervision should have the largest effect. We report our results in Figure 4a and visualize examples of reconstructed pointclouds in Figure 5. DeFiNe significantly outperforms other methods on this dataset, including IIB. Our virtual view augmentations lead to a large (20%) relative improvement, showing that DeFiNe benefits from a scene representation that encourages multi-view consistency.

Method	Abs.Rel↓	RMSE↓	$\delta_{1.25}$ ↑
DPSNet [15]	0.126	0.314	—
NAS [20]	0.107	0.281	—
IIB [56]	0.116	0.281	0.908
<b>DeFiNe (128 × 192)</b>	0.093	0.246	0.911
<b>DeFiNe (240 × 320)</b>	<b>0.089</b>	<b>0.232</b>	<b>0.915</b>

(a) Depth estimation results.



(b) Virtual depth estimation results.

Fig. 4: **Depth estimation results on ScanNet-Stereo.** (a) We outperform contemporary methods by a large margin. (b) Depth estimation results on virtual cameras using different values for  $\sigma_v$  at training and test time.

	Variation	Lower is better ↓			Higher is better ↑		
		Abs. Rel	Sq. Rel	RMSE	$\delta_{1.25}$	$\delta_{1.25^2}$	$\delta_{1.25^3}$
1	Depth-Only	0.098	0.046	0.257	0.902	0.972	0.990
2	w/ Conv. RGB encoder [16]	0.114	0.058	0.294	0.866	0.961	0.982
3	w/ 64-dim R18 RGB encoder	0.104	0.049	0.270	0.883	0.966	0.985
4	w/o camera information	0.229	0.157	0.473	0.661	0.874	0.955
5	w/o global rays encoding	0.097	0.047	0.261	0.897	0.962	0.988
6	w/ equal loss weights	0.095	0.047	0.259	0.908	0.968	0.990
7	w/ epipolar cues [56]	0.094	0.048	0.254	0.905	0.972	0.990
8	w/o Augmentations	0.117	0.060	0.291	0.870	0.959	0.981
9	w/o Virtual Cameras	0.104	0.058	0.268	0.891	0.965	0.986
10	w/o Canonical Jittering	0.099	0.046	0.261	0.897	0.970	0.988
11	w/o Canonical Randomization	0.096	0.044	0.253	0.905	0.971	0.989
	<b>DeFiNe</b>	<b>0.093</b>	<b>0.042</b>	<b>0.246</b>	<b>0.911</b>	<b>0.974</b>	<b>0.991</b>

Table 1: **Ablation study for ScanNet-Stereo**, using different variations.

### 4.3 Ablation Study

We perform a detailed ablation study to evaluate the effectiveness of each component in our proposed architecture, with results shown in Table 1. Firstly, we evaluate performance when (Table 1:1) learning only depth estimation, and see that the joint learning of view synthesis as an auxiliary task leads to significant improvements. The claim that depth estimation improves view synthesis has been noted before [4,53], and is attributed to the well-known fact that multi-view consistency facilitates the generation of images from novel viewpoints. However, our experiments also show the inverse: view synthesis improves depth estimation. Our hypothesis is that appearance is required to learn multi-view consistency since it enables visual correlation across frames. By introducing view synthesis as an additional task, we are also encoding appearance-based information into our latent representation. This leads to improvements in depth estimation even though no explicit feature matching is performed at an architecture or loss level.

We also ablate different variations of our RGB encoder for the generation of image embeddings and show that (Table 1:2) our proposed multi-level feature map concatenation (Figure 2b) leads to the best results relative to the stan-




Fig. 5: **Reconstructed two-view pointclouds**, from ScanNet-Stereo. DeFiNe pointclouds are generated using both depth maps and RGB images queried from our learned latent representation.




Fig. 6: **Depth estimation and view synthesis results** on ScanNet. Although view synthesis is not our primary goal, it can be achieved with minimal modifications, and we show that it improves depth estimation performance.

dard single convolutional layer, or (Table 1:3) using 1/4-resolution ResNet18 64-dimensional feature maps. Similarly, we also ablate some of our design choices, namely (Table 1:4) the use of camera embeddings instead of positional encodings; (Table 1:5) global viewing rays (Section 3.2) instead of traditional relative viewing rays; (Table 1:6) the use of  $\lambda_s = 1$  in the loss calculation (Equation 3) such that both depth and view synthesis tasks have equal weights; and (Table 1:7) the use of epipolar cues as additional geometric embeddings, as proposed by IIB [56]. As expected, camera embeddings are crucial for multi-view consistency and global viewing rays improve over the standard relative formulation. Interestingly, using a smaller  $\lambda_s$  degrades depth estimation performance, providing further evidence that the joint learning of view synthesis is beneficial for multi-view consistency. We did not observe meaningful improvements when incorporating the epipolar cues from IIB [56], indicating that DeFiNe is capable of directly capturing these constraints at an input-level due to the increase in viewpoint diversity. Lastly, we ablate the impact of our various proposed geometric augmentations (Section 3.3), showing that they are key to our reported state-of-the-art performance.

Lastly, we evaluate depth estimation from virtual cameras, using different noise levels  $\sigma_v$  at test time. We also train models using different noise levels and report the results in Figure 4b. From these results, we can see that the optimal virtual noise level, when evaluating at the target location, is  $\sigma_v = 0.25$  m (yellow line), relative to the baseline without virtual noise (blue line). However, models trained with higher virtual noise (e.g., the orange line, with  $\sigma_v = 1$  m) are more robust to larger deviations from the target location.

#### 4.4 Video Depth Estimation

To highlight the flexibility of our proposed architecture, we also experiment using video data from ScanNet following the training protocol of Tang et al. [44]. We evaluate performance on both ScanNet itself, using their evaluation protocol [44], as well as zero-shot transfer (without fine-tuning) to the 7-Scenes dataset. Table 2 reports quantitative results, while Figure 6 provides qualitative examples. On ScanNet, DeFiNe outperforms most published approaches, significantly improving over DeMoN [49], BA-Net [44], and CVD [28] both in terms of performance and speed. We are competitive with DeepV2D [45] in terms of performance, and roughly  $14\times$  faster, owing to the fact that DeFiNe does not require bundle adjustment or any sort of test-time optimization. In fact, our inference time of 49 ms can be split into 44 ms for encoding and only 5 ms for decoding, enabling very efficient generation of depth maps after information has been encoded once. The only method that outperforms DeFiNe in terms of speed is NeuralRecon [43], which uses a sophisticated TSDF integration strategy. Performance-wise, we are also competitive with NeuralRecon, improving over their reported results in one of the three metrics (Sq. Rel).

Next, we evaluate zero-shot transfer from ScanNet to 7-Scenes, which is a popular test of generalization for video depth estimation. In this setting, DeFiNe significantly improves over all other methods, including DeepV2D (which fails to generalize) and NeuralRecon ( $\sim 40\%$  improvement). We attribute this large gain to the highly intricate and specialized nature of these other architectures. In contrast, our method has no specialized module and instead learns a geometrically-consistent multi-view latent representation.

In summary, we achieve competitive results on ScanNet and significantly improve the state-of-the-art for video depth generalization, as evidenced by the large gap between DeFiNe and the best-performing methods on 7-Scenes.

#### 4.5 Depth from Novel Viewpoints

We previously discussed the strong performance that DeFiNe achieves on traditional depth estimation benchmarks, and showed how it improves out-of-domain generalization by a wide margin. Here, we explore another aspect of generalization that our architecture enables: viewpoint generalization. This is possible because, in addition to traditional depth estimation from RGB images, DeFiNe can also generate depth maps from arbitrary viewpoints since it only requires camera embeddings to decode estimates. We explore this capability in two different ways: *interpolation* and *extrapolation*. When interpolating, we encode frames




Method	Abs. Rel↓	Sq. Rel↓	RMSE↓	Speed (ms)↓
<b>ScanNet test split [44]</b>				
DeMoN [49]	0.231	0.520	0.761	110
MiDas-v2 [33]	0.208	0.318	0.742	-
BA-Net [44]	0.091	0.058	0.223	95
CVD [28]	0.073	0.037	0.217	2400
DeepV2D [45]	0.057	<b>0.010</b>	<u>0.168</u>	690
NeuralRecon [43]	<b>0.047</b>	0.024	<b>0.164</b>	<b>30</b>
<b>DeFiNe (128 × 192)</b>	0.059	0.022	0.184	<u>49</u>
<b>DeFiNe (240 × 320)</b>	<u>0.056</u>	<u>0.019</u>	0.176	78
<b>Zero-shot transfer to 7-Scenes [40]</b>				
DeMoN [49]	0.389	0.420	0.855	110
NeuralRGBD [25]	0.176	0.112	0.441	202
DPSNet [15]	0.199	0.142	0.438	322
DeepV2D [45]	0.437	0.553	0.869	347
CNMNet [27]	0.161	0.083	0.361	80
NeuralRecon [43]	0.155	0.104	0.347	<b>30</b>
EST [56]	<u>0.118</u>	<u>0.052</u>	<u>0.298</u>	71
<b>DeFiNe (128 × 192)</b>	<b>0.100</b>	<b>0.039</b>	<b>0.253</b>	<u>49</u>

Table 2: **Depth estimation results on ScanNet and 7-Scenes.** DeFiNe is competitive with other state-of-the-art methods on ScanNet, and outperforms all published methods in zero-shot transfer to 7-Scenes by a large margin.


at  $\{t-5, t+5\}$ , and decode virtual depth maps at locations  $\{t-4, \dots, t+4\}$ . When extrapolating, we encode frames at  $\{t-5, \dots, t-1\}$ , and decode virtual depth maps at locations  $\{t, \dots, t+8\}$ . We use the same training and test splits as in our stereo experiments, with a downsampling factor of 20 to encourage smaller overlaps between frames. As baselines for comparison, we consider the explicit projection of 3D information from encoded frames onto these new viewpoints. We evaluate both standard depth estimation networks [11,13,22] as well as DeFiNe itself, that can be used to either explicitly project information from encoded frames onto new viewpoints (projection), or query from the latent representation at that same location (query).

Figure 7 reports results in terms of root mean squared error (RMSE) considering only valid projected pixels. Of particular note, our multi-frame depth estimation architecture significantly outperforms other single-frame baselines. However, and most importantly, results obtained by implicit querying consistently outperform those obtained via explicit projection. This indicates that our model is able to improve upon available encoded information via the learned latent representation. Furthermore, we also generate geometrically consistent estimates for areas without valid explicit projections (Figure 7c). As the camera tilts to the right, the floor is smoothly recreated in unobserved areas, as well as the partially observed bench. Interestingly, the chair at the extreme right was not recreated, which could be seen as a failure case. However, because the chair



(a) Depth interpolation results.

(b) Depth extrapolation results.



(c) Depth extrapolation to future timesteps. Images and ground-truth depth maps are displayed only for comparison. DeFiNe can complete unseen portions of the scene in a geometrically-consistent way, generating dense depth maps from novel viewpoints.

Fig. 7: **Depth estimation results** from novel viewpoints.

was never observed in the first place, it is reasonable for the model to assume the area is empty, and recreate it as a continuation of the floor.

## 5 Conclusion

We introduced Depth Field Networks (DeFiNe), a generalist framework for training multi-view consistent depth estimators. Rather than explicitly enforcing geometric constraints at an architecture or loss level, we use geometric embeddings to condition network inputs, alongside visual information. To learn a geometrically-consistent latent representation, we propose a series of 3D augmentations designed to promote viewpoint, rotation, and translation equivariance. We also show that the introduction of view synthesis as an auxiliary task improves depth estimation without requiring additional ground truth. We achieve state-of-the-art results on the popular ScanNet stereo benchmark, and competitive results on the ScanNet video benchmark with no iterative refinement or explicit geometry modeling. We also demonstrate strong generalization properties by achieving state-of-the-art results on zero-shot transfer from ScanNet to 7-Scenes. The general nature of our framework enables many exciting avenues for future work, including additional tasks such as optical flow, extension to dynamic scenes, spatio-temporal representations, and uncertainty estimation.

## References

1. Azinović, D., Martin-Brualla, R., Goldman, D.B., Nießner, M., Thies, J.: Neural RGB-D surface reconstruction. arXiv preprint arXiv:2104.04532 (2021)
2. Chang, J.R., Chen, Y.S.: Pyramid stereo matching network. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR). pp. 5410–5418 (2018)
3. Dai, A., Chang, A.X., Savva, M., Halber, M., Funkhouser, T., Nießner, M.: ScanNet: Richly-annotated 3D reconstructions of indoor scenes. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR). pp. 5828–5839 (2017)
4. Deng, K., Liu, A., Zhu, J.Y., Ramanan, D.: Depth-supervised NeRF: Fewer views and faster training for free. arXiv preprint arXiv:2107.02791 (2021)
5. Dosovitskiy, A., Beyer, L., Kolesnikov, A., Weissenborn, D., Zhai, X., Unterthiner, T., Dehghani, M., Minderer, M., Heigold, G., Gelly, S., et al.: An image is worth 16x16 words: Transformers for image recognition at scale. arXiv preprint arXiv:2010.11929 (2020)
6. Eigen, D., Fergus, R.: Predicting depth, surface normals and semantic labels with a common multi-scale convolutional architecture. In: Proceedings of the International Conference on Computer Vision (ICCV). pp. 2650–2658 (2015)
7. Eigen, D., Puhrsch, C., Fergus, R.: Depth map prediction from a single image using a multi-scale deep network. In: Advances in Neural Information Processing Systems (NeurIPS) (2014)
8. Esser, P., Rombach, R., Ommer, B.: Taming transformers for high-resolution image synthesis. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 12873–12883 (2021)
9. Fu, H., Gong, M., Wang, C., Batmanghelich, K., Tao, D.: Deep ordinal regression network for monocular depth estimation. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR). pp. 2002–2011 (2018)
10. Godard, C., Mac Aodha, O., Brostow, G.J.: Unsupervised monocular depth estimation with left-right consistency. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR). pp. 270–279 (2017)
11. Godard, C., Mac Aodha, O., Firman, M., Brostow, G.J.: Digging into self-supervised monocular depth prediction. In: Proceedings of the International Conference on Computer Vision (ICCV) (2019)
12. Gordon, A., Li, H., Jonschkowski, R., Angelova, A.: Depth from videos in the wild: Unsupervised monocular depth learning from unknown cameras. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR) (2019)
13. Guizilini, V., Ambrus, R., Pillai, S., Raventos, A., Gaidon, A.: 3D packing for self-supervised monocular depth estimation. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR) (2020)
14. He, K., Zhang, X., Ren, S., Sun, J.: Deep residual learning for image recognition. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR). pp. 770–778 (2016)
15. Im, S., Jeon, H.G., Lin, S., Kweon, I.S.: DPSNet: End-to-end deep plane sweep stereo. In: Proceedings of the International Conference on Learning Representations (ICLR) (2019)
16. Jaegle, A., Borgeaud, S., Alayrac, J.B., Doersch, C., Ionescu, C., Ding, D., Kop-pula, S., Zoran, D., Brock, A., Shelhamer, E., et al.: Perceiver IO: A general architecture for structured inputs & outputs. arXiv preprint arXiv:2107.14795 (2021)

17. Jaegle, A., Gimeno, F., Brock, A., Vinyals, O., Zisserman, A., Carreira, J.: Perceiver: General perception with iterative attention. In: Proceedings of the International Conference on Machine Learning (ICML). pp. 4651–4664 (2021)
18. Kendall, A., Martirosyan, H., Dasgupta, S., Henry, P., Kennedy, R., Bachrach, A., Bry, A.: End-to-end learning of geometry and context for deep stereo regression. In: Proceedings of the International Conference on Computer Vision (ICCV). pp. 66–75 (2017)
19. Kopf, J., Rong, X., Huang, J.B.: Robust consistent video depth estimation. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR). pp. 1611–1621 (2021)
20. Kusupati, U., Cheng, S., Chen, R., Su, H.: Normal assisted stereo depth estimation. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR). pp. 2189–2199 (2020)
21. Laina, I., Ruppel, C., Belagiannis, V., Tombari, F., Navab, N.: Deeper depth prediction with fully convolutional residual networks. In: Proceedings of the International Conference on 3D Vision (3DV). pp. 239–248 (2016)
22. Lee, J.H., Han, M.K., Ko, D.W., Suh, I.H.: From big to small: Multi-scale local planar guidance for monocular depth estimation. arXiv:1907.10326 (2019)
23. Li, Z., Liu, X., Drenkow, N., Ding, A., Creighton, F.X., Taylor, R.H., Unberath, M.: Revisiting stereo depth estimation from a sequence-to-sequence perspective with transformers. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. pp. 6197–6206 (2021)
24. Lipson, L., Teed, Z., Deng, J.: RAFT-Stereo: Multilevel recurrent field transforms for stereo matching. In: Proceedings of the International Conference on 3D Vision (3DV). pp. 218–227 (2021)
25. Liu, C., Gu, J., Kim, K., Narasimhan, S.G., Kautz, J.: Neural RGB  $\rightarrow$  D sensing: Depth and uncertainty from a video camera. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR). pp. 10986–10995 (2019)
26. Long, X., Liu, L., Li, W., Theobalt, C., Wang, W.: Multi-view depth estimation using epipolar spatio-temporal networks. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR). pp. 8258–8267 (2021)
27. Long, X., Liu, L., Theobalt, C., Wang, W.: Occlusion-aware depth estimation with adaptive normal constraints. In: Proceedings of the European Conference on Computer Vision (ECCV). pp. 640–657 (2020)
28. Luo, X., Huang, J.B., Szeliski, R., Matzen, K., Kopf, J.: Consistent video depth estimation. ACM Transactions on Graphics (TOG) **39**(4) (2020)
29. 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: Proceedings of the European Conference on Computer Vision (ECCV). pp. 405–421 (2020)
30. Newcombe, R.A., Izadi, S., Hilliges, O., Molyneaux, D., Kim, D., Davison, A.J., Kohi, P., Shotton, J., Hodges, S., Fitzgibbon, A.: KinectFusion: Real-time dense surface mapping and tracking. In: Proceedings of the IEEE International Symposium on Mixed and Augmented Reality (ISMAR) (2011)
31. Novotny, D., Larlus, D., Vedaldi, A.: Learning 3D object categories by looking around them. In: Proceedings of the International Conference on Computer Vision (ICCV). pp. 5218–5227 (2017)
32. Ranftl, R., Bochkovskiy, A., Koltun, V.: Vision transformers for dense prediction. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. pp. 12179–12188 (2021)

33. Ranftl, R., Lasinger, K., Hafner, D., Schindler, K., Koltun, V.: Towards robust monocular depth estimation: Mixing datasets for zero-shot cross-dataset transfer. *IEEE Transactions on Pattern Analysis and Machine Intelligence* (2020)
34. Rematas, K., Liu, A., Srinivasan, P.P., Barron, J.T., Tagliasacchi, A., Funkhouser, T., Ferrari, V.: Urban radiance fields. *arXiv preprint arXiv:2111.14643* (2021)
35. Rombach, R., Esser, P., Ommer, B.: Geometry-free view synthesis: Transformers and no 3D priors. In: *Proceedings of the International Conference on Computer Vision (ICCV)*. pp. 14356–14366 (2021)
36. Sajjadi, M.S., Meyer, H., Pot, E., Bergmann, U., Greff, K., Radwan, N., Vora, S., Lucic, M., Duckworth, D., Dosovitskiy, A., Uszkoreit, J., Funkhouser, T., Tagliasacchi, A.: Scene representation transformer: Geometry-free novel view synthesis through set-latent scene representations. *arXiv preprint arXiv:2111.13152* (2021)
37. Saxena, A., Chung, S., Ng, A.: Learning depth from single monocular images. In: *Advances in Neural Information Processing Systems (NeurIPS)* (2005)
38. Schonberger, J.L., Frahm, J.M.: Structure-from-motion revisited. In: *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*. pp. 4104–4113 (2016)
39. Shorten, C., Khoshgoftaar, T.M.: A survey on image data augmentation for deep learning. *Journal of Big Data* **6**(1), 1–48 (2019)
40. Shotton, J., Glocker, B., Zach, C., Izadi, S., Criminisi, A., Fitzgibbon, A.: Scene coordinate regression forests for camera relocalization in RGB-D images. In: *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*. pp. 2930–2937 (2013)
41. Shu, C., Yu, K., Duan, Z., Yang, K.: Feature-metric loss for self-supervised learning of depth and egomotion. In: *Proceedings of the European Conference on Computer Vision (ECCV)*. pp. 572–588 (2020)
42. Sitzmann, V., Rezhikov, S., Freeman, B., Tenenbaum, J., Durand, F.: Light field networks: Neural scene representations with single-evaluation rendering. In: *Advances in Neural Information Processing Systems (NeurIPS)* (2021)
43. Sun, J., Xie, Y., Chen, L., Zhou, X., Bao, H.: NeuralRecon: Real-time coherent 3D reconstruction from monocular video. In: *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*. pp. 15598–15607 (June 2021)
44. Tang, C., Tan, P.: BA-Net: Dense bundle adjustment network. *arXiv preprint arXiv:1806.04807* (2018)
45. Teed, Z., Deng, J.: DeepV2D: Video to depth with differentiable structure from motion. In: *Proceedings of the International Conference on Learning Representations (ICLR)* (2020)
46. Teed, Z., Deng, J.: RAFT: Recurrent all-pairs field transforms for optical flow. In: *Proceedings of the European Conference on Computer Vision (ECCV)* (2020)
47. Teed, Z., Deng, J.: DROID-SLAM: Deep visual SLAM for monocular, stereo, and RGB-D cameras. In: *Advances in Neural Information Processing Systems (NeurIPS)* (2021)
48. Teed, Z., Deng, J.: Raft-3D: Scene flow using rigid-motion embeddings. In: *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*. pp. 8375–8384 (2021)
49. Ummenhofer, B., Zhou, H., Uhrig, J., Mayer, N., Ilg, E., Dosovitskiy, A., Brox, T.: DeMoN: Depth and motion network for learning monocular stereo. In: *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*. pp. 5038–5047 (2017)

50. Vasiljevic, I., Guizilini, V., Ambrus, R., Pillai, S., Burgard, W., Shakhnarovich, G., Gaidon, A.: Neural ray surfaces for self-supervised learning of depth and ego-motion. In: *Proceedings of the International Conference on 3D Vision (3DV)* (2020)
51. Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A.N., Kaiser, L.u., Polosukhin, I.: Attention is all you need. In: *Advances in Neural Information Processing Systems (NeurIPS)* (2017)
52. Watson, J., Mac Aodha, O., Prisacariu, V., Brostow, G., Firman, M.: The temporal opportunist: Self-supervised multi-frame monocular depth. In: *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*. pp. 1164–1174 (2021)
53. 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: *Proceedings of the International Conference on Computer Vision (ICCV)* (2021)
54. Xie, Y., Takikawa, T., Saito, S., Litany, O., Yan, S., Khan, N., Tombari, F., Tompkin, J., Sitzmann, V., Sridhar, S.: Neural fields in visual computing and beyond. *arXiv preprint arXiv:2111.11426* (2021)
55. Yao, Y., Luo, Z., Li, S., Fang, T., Quan, L.: MVSNet: Depth inference for unstructured multi-view stereo. In: *Proceedings of the European Conference on Computer Vision (ECCV)*. pp. 767–783 (2018)
56. Yifan, W., Doersch, C., Arandjelović, R., Carreira, J., Zisserman, A.: Input-level inductive biases for 3D reconstruction. *arXiv preprint arXiv:2112.03243* (2021)
57. Yu, A., Ye, V., Tancik, M., Kanazawa, A.: pixelNeRF: Neural radiance fields from one or few images. In: *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)* (2021)
58. Zhou, H., Ummenhofer, B., Brox, T.: DeepTAM: Deep tracking and mapping. In: *Proceedings of the European Conference on Computer Vision (ECCV)*. pp. 822–838 (2018)
59. Zhou, T., Brown, M., Snavely, N., Lowe, D.G.: Unsupervised learning of depth and ego-motion from video. In: *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)* (2017)
60. Zhu, Z., Peng, S., Larsson, V., Xu, W., Bao, H., Cui, Z., Oswald, M.R., Pollefeys, M.: NICE-SLAM: Neural implicit scalable encoding for SLAM. *arXiv preprint arXiv:2112.12130* (2021)