

## Bilateral Propagation Network for Depth Completion

Jie Tang<sup>1</sup>, Fei-Peng Tian<sup>2</sup>, Boshi An<sup>3</sup>, Jian Li<sup>1\*</sup>, Ping Tan<sup>4,2</sup>  
<sup>1</sup>National University of Defense Technology,  
<sup>2</sup>LightIllusions, <sup>3</sup>Peking University,  
<sup>4</sup>The Hong Kong University of Science and Technology

### Abstract

Depth completion aims to derive a dense depth map from sparse depth measurements with a synchronized color image. Current state-of-the-art (SOTA) methods are predominantly propagation-based, which work as an iterative refinement on the initial estimated dense depth. However, the initial depth estimations mostly result from direct applications of convolutional layers on the sparse depth map. In this paper, we present a Bilateral Propagation Network (BP-Net), that propagates depth at the earliest stage to avoid directly convolving on sparse data. Specifically, our approach propagates the target depth from nearby depth measurements via a non-linear model, whose coefficients are generated through a multi-layer perceptron conditioned on both radiometric difference and spatial distance. By integrating bilateral propagation with multi-modal fusion and depth refinement in a multi-scale framework, our BP-Net demonstrates outstanding performance on both indoor and outdoor scenes. It achieves SOTA on the NYUv2 dataset and ranks 1st on the KITTI depth completion benchmark at the time of submission. Experimental results not only show the effectiveness of bilateral propagation but also emphasize the significance of early-stage propagation in contrast to the refinement stage. Our code and trained models will be available on the project page.

### 1. Introduction

Dense depth perception, *i.e.* estimating per-pixel distance to the camera, is crucial for 3D tasks, with a rapid growth of applications, such as augmented reality, autonomous driving and robotics. However, due to current hardware limitations, it is still difficult to directly measure dense depth maps. Depth measured from LiDAR [9] or SfM [38] is typically sparse, which is insufficient for real applications like scene reconstruction [47] and robot grasping [29]. Depth completion<sup>1</sup>, *i.e.* estimating dense depth map from sparse

\*indicates the corresponding author.

<sup>1</sup>The exact name should be *image guided depth completion*, if considering some early attempts don't utilize color images.

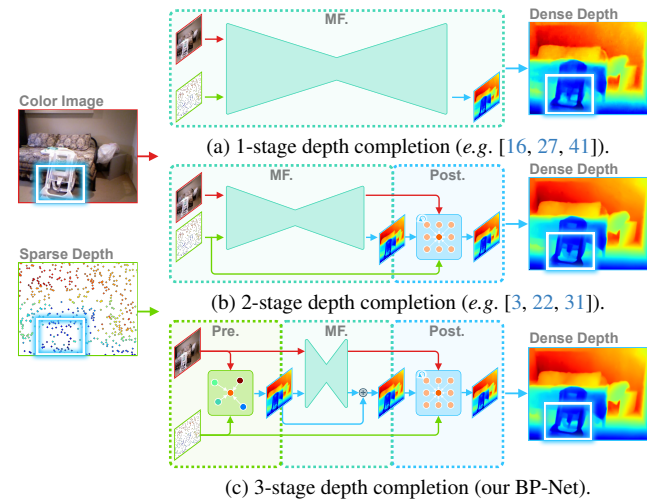


Figure 1. Depth completion is to produce a dense depth map from sparse depth measurements with a synchronized color image. (a) Earlier methods take a 1-stage approach to solve this problem by multi-modal fusion network (MF). (b) Some recent works include a post-processing stage (Post.) to refine the results. (c) Our BP-Net further introduces a pre-processing stage (Pre.) to make the multi-modal fusion (MF) network more effective for better results.

depth measurements with a synchronized color image, is a promising solution with the advantage of exploiting complementary information from data in different modalities. A key challenge in this task is to effectively process irregularly sampled sparse depth points.

In the context of depth completion, deep learning-based methods have achieved great success. As shown in Fig. 1a, earlier methods [16, 27, 41] often take a 1-stage depth completion strategy by a multi-modal fusion (MF) network to recover dense depth maps. These methods mostly suffer from excessive smoothing at depth edges which leads to loss of details in the results. More recent works [3, 22, 31] have started to adopt a 2-stage approach as shown in Fig. 1b, with an additional post-processing stage via propagation-based networks. By revisiting sparse depth measurements to iteratively refine the regressed depth, such propagation-

based post-processing can alleviate the over-smoothing problems. Thus, they have better performance and dominate the current state-of-the-art (SOTA).

Although deep learning-based methods generate superior results than conventional methods [19], these 1-stage and 2-stage methods generally suffer from two common issues. Firstly, they typically use 0 value to indicate an unknown pixel in the input sparse depth map, inducing ambiguity to distinguish between valid and invalid depth measurements as discussed in [23]. Secondly, the spatially invariant convolution is not an ideal processor for irregularly sampled sparse depth points [44]. These issues downgrade the performance of the multi-modal fusion network, and the post-processing stage may not be able to solve these issues effectively.

In this paper, we present a bilateral propagation network (BP-Net) to address the above issues. Our BP-Net is a 3-stage method as illustrated in Fig. 1c, with a pre-processing stage to compute an initial dense depth map for the following multi-modal fusion network. The initial dense depth is explicitly propagated from the valid sparse depth measurements, which avoids the above issues of ambiguity and invariant convolution. Specifically, we introduce a non-linear propagation model, where the output depth is a combination of nearby valid depth measurements weighted by learned coefficients. Inspired by the well-known bilateral filtering [42], we propose a novel multi-layer perceptron to produce these combination coefficients conditioned on both *radiometric difference* and *spatial distance*. In this way, depth can be propagated with the preference of nearest values in both domain and range. This bilateral propagation is fully differentiable, that can be optimized in an end-to-end fashion together with the multi-modal fusion and depth refinement networks. Moreover, to deal with long-range depth propagation (*e.g.* a large hole without valid depth measurements nearby), we adopt a multi-scale architecture in which depth is estimated from coarse to fine and the low-resolution result is used as prior for high-resolution estimation. Furthermore, a multi-scale loss is designed to better supervise the multi-scale network.

We perform experiments to validate our BP-Net with both indoor and outdoor scenes, under standard evaluation criteria. Our method achieves SOTA performance on the NYUv2 dataset and ranks 1st on the KITTI benchmark at the time of paper submission. We conduct ablation studies to demonstrate the efficacy of each component and also analyze performance under various sparsity levels. Experimental results not only show the effectiveness of bilateral propagation but also emphasize the significance of early-stage propagation in contrast to the refinement stage. Code and trained models are available at <https://github.com/kakaxi314/BP-Net>.

## 2. Related Work

To estimate dense depth map, many techniques work on data from RGB-D camera, like Kinect, whose goal is depth super-resolution or depth inpainting. To this end, filter-based [10, 15, 18], optimization-based [6, 8, 50], and learning-based [21, 43, 52] methods are proposed to upsample resolution or fill holes in depth map. On the other hand, depth completion from irregularly distributed extremely sparse data measured by LiDAR or SfM attracts increasing attention with a large number of approaches emerging in recent years. In general, previous depth completion approaches mainly focused on the following three problems: how to process sparse data, how to fuse multi-modal data, and how to refine results. We briefly review previous techniques from these three aspects.

**Sparse Data Processing.** As traditional methods, Ku *et al.* [19] utilize a series of classical image processing operators like dilation, hole filling *etc.* to densify sparse depth map. Zhao *et al.* [55] calculate surface normal in spherical coordinate system, and then estimate depth with local surface sharing assumption. In deep learning era, some early attempts focus on sparsity invariant operations to avoid simply convolving on sparse depth map. Uhrig *et al.* [44] replace the conventional convolution operation with sparsity-invariant convolution that keeps track of validation masks at each layer. Huang *et al.* [14] employ the sparsity-invariant layer in an encoder-decoder network structure. Eldesokey *et al.* [7] further extend the validation masks in sparsity invariant layer to a continuous confidence field. There are also works embedding sparse depth by differentiable solver in an end-to-end trained network. Qu *et al.* [34, 35] process sparse depth by solving linear regression in deep network. Conti *et al.* [5] optimize weighted linear regression in multi-scale to inject sparse depth into network. Some works propose to first densify the sparse depth by classical approaches, and then enhance the densified depth by learning-based method. Chen *et al.* [2] adopt nearest interpolation on sparse depth map before feed it to deep network. Wang *et al.* [46] utilize [19] to estimate the initial depth map. Liu *et al.* [23] further replace the hand-crafted interpolation with a differentiable kernel regression layer. In comparison, our BP-Net propagates an initial dense depth map for the following operations, but depending on both image content and spatial distance, while spatial distance from valid depth to the target pixel is not considered in previous methods.

**Multi-modal Fusion.** To fuse multi-modal data of depth map and color image, some methods [16, 27, 28] adopt the early-fusion scheme, *e.g.* Ma *et al.* [27] concatenate the depth map and the RGB image as input for deep network to process. While some [41, 49, 51] adopt the late-fusion scheme, *e.g.* Tang *et al.* [41] propose a guided convolutional module to fuse image feature with depth feature in multi-stage. Some works [33, 48] fuse depth map with es-

timated surface normal, *e.g.* Qiu *et al.* [33] utilize surface normal as intermediate representation for fusion with sparse depth. And some [1, 54] explore the geometry information in 3D points, *e.g.* Chen *et al.* [1] fuse 2D image feature with 3D point feature for depth estimation. Recent, Transformer architecture [45] is also introduced to establish attention mechanism for multi-modal fusion [36, 53]. Our BP-Net mainly focused on sparse data processing stage, with only a simple U-Net structure for multi-modal fusion in early-fusion scheme, and is not conflict to other multi-modal fusion designs in theory.

**Depth Refinement.** Directly regressed depth map may suffer from blur effect on object boundaries. Depth refinement mostly follows the spatial propagation mechanism [24], which iteratively refines the regressed depth by a local linear model with learned affinity. Cheng *et al.* [3] adopt convolution operation to update all pixels simultaneously with a fixed kernel size, and further extend it to CSPN++ [4] assembling results from multiple kernel sizes and multiple iteration steps. Park *et al.* [31] provide a non-local neighborhood for propagation by learning the offsets to the regular grid. Lin *et al.* [22] extend the propagation process with affinity matrix adaptive to neighbors of different distances. Liu *et al.* [25] construct graph neural network integrating 3D information to estimate neighbors in each update iteration. Wang *et al.* [46] adjust the kernel scope from large to small during the update process. These propagation-based methods dominate current SOTA, but still require an initial dense depth which is mostly yielded by directly applying convolutional layers on the sparse depth map. In contrast, our BP-Net propagates depth at the earliest stage to avoid convolving on sparse depth map. Though we implement a simple CSPN++ [4] for depth refinement, our method should be also compatible to other depth refinement approaches.

### 3. The Proposed Method

#### 3.1. Overview

Let  $I$  denote a color image, and  $S$  represent the synchronized sparse depth map typically obtained by projecting measured 3D points onto the image plane using calibration parameters. The color image  $I$  provides rich scene context with well-defined boundaries and precise semantics. In contrast, the depth map  $S$ , while sharing the same resolution as  $I$ , consists of sparse and irregular yet fairly accurate depth measurements with definite geometry. The goal of depth completion is to produce a dense depth map  $D$  by exploring the complementary information embedded in  $I$  and  $S$ .

As shown in Fig. 2, the proposed method is a multi-scale network with 6 scales marked from 0 to 5. The dense depth  $D$  is estimated from the lowest resolution in scale 5 to the highest resolution in scale 0. For any scale  $s$  in this network,

the dense depth map is produced in three sequential stages as illustrated in Fig. 1c: a pre-processing stage to generate an initial dense depth map from sparse depth measurements via the proposed bilateral propagation module (in Sec. 3.2), a multi-modal fusion stage to generate a residual depth map by a simple U-Net fusing both image features and depth features in an early fusion scheme (in Sec. 3.3), and a refinement stage to update the dense depth map with sparse depth via convolutional propagation module in an iterative manner (in Sec. 3.4). Finally, we also introduce the implementation details in Sec. 3.5.

#### 3.2. Bilateral Propagation Module

##### 3.2.1 Depth Parameterization

We aim to propagate a dense depth map  $D'$  from the sparse depth map  $S$  to mitigate the issues arising from the sparsity problem in the subsequent multi-modal fusion. Instead of naively applying convolution operations on the irregular sparse depth map, we explicitly model the depth propagation process as:

$$D'_i = \sum_{j \in \mathcal{N}(i)} \omega_{ij} D'_{ij}{}^c = \sum_{j \in \mathcal{N}(i)} \omega_{ij} (\alpha_{ij} S_j + \beta_{ij}). \quad (1)$$

Here, the target depth  $D'_i$  at pixel  $i$  is a local combination of  $N$  nearest valid sparse depth  $S_j$  with coefficients  $\omega_{ij}$ ,  $\alpha_{ij}$  and  $\beta_{ij}$ , where  $j \in \mathcal{N}(i)$  is a nearby pixel. Specifically, our approach first generate depth candidates  $D'_{ij}{}^c$  by an affine transformation of the sparse depth  $S_j$  with coefficients  $\alpha_{ij}$  and  $\beta_{ij}$ , then linearly combine depth candidates  $D'_{ij}{}^c$  with coefficient  $\omega_{ij}$  to produce the target depth  $D'_i$ . We adopt Euclidean distance on the image plane to find  $\mathcal{N}(i)$  containing  $N$  nearest neighboring valid pixels for any pixel  $i$ . Notably, we observe that a relatively small number is sufficient for  $N$ . Based on empirical findings, we set  $N = 4$ . Further discussion about the choice of  $N$  can be found in Sec. 4.3.1.

Our depth parameterization is a general formulation for dense depth estimation from sparse valid depth. Many conventional interpolation methods can be viewed as specific cases of our formulation. For instance, nearest interpolation corresponds to the case with  $N = 1$ ,  $\omega_{ij} = 1$ , and bilateral interpolation aligns with a specific case where  $\omega_{ij}$  is determined by a handcrafted kernel. Moreover, in our formulation,  $\alpha$ ,  $\beta$ , and  $\omega$  are dynamically generated depending on both image content and spatial distance, providing adaptability to different sparse and irregular depth distributions. For cases where the depth at the target pixel is far away from those at the neighboring valid pixels, our formulation might learn an  $\alpha$  close to 0, disregarding the neighboring valid depth and directly regressing the target depth through the parameter  $\beta$ . This scenario is challenging for manually designed local filters to address. Consequently, our depth parameterization is much more powerful than a mere local




Figure 2. **Overview of the proposed approach.** Dense depth is estimated in a multi-scale scheme, from low-resolution to high-resolution. In each scale, dense depth map is produced from sparse depth map and image feature by 3-stage processing depicted in Fig. 1c.

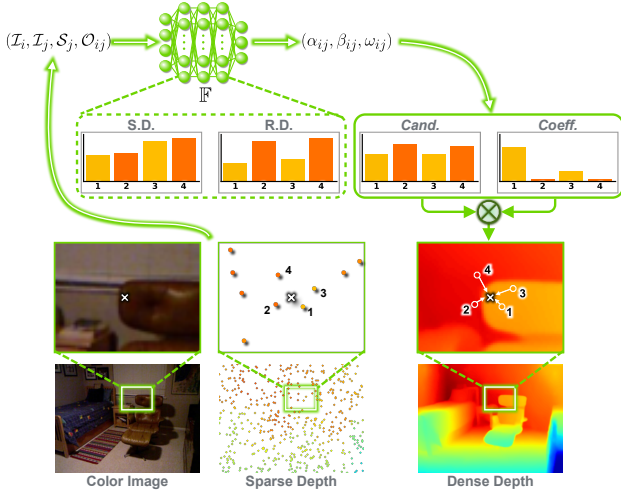


Figure 3. **Bilateral Propagation Module.** The target depth is a convex combination of depth candidates (*Cand.*) weighted by coefficients (*Coeff.*), which are yielded by a MLP  $\mathbb{F}$  conditioned on both spatial distance (*S.D.*) and radiometric difference (*R.D.*) from the nearby depth measurements.

filter on sparse valid depth and thus can help avoid the sparsity issue in the following network operations.

### 3.2.2 Parameter Generation

To obtain the above coefficient parameters  $\alpha$ ,  $\beta$ , and  $\omega$  for depth propagation, we draw inspiration from the well-established bilateral filtering [42]. Bilateral filtering [42] is known for producing edge-preserving results by combining both **radiometric difference** and **spatial distance** in generating filtering weights. As illustrated in Fig. 3, we adopt a MLP as the hyper-network  $\mathbb{F}$ , that is

$$\alpha_{ij}, \beta_{ij}, \omega_{ij} = \mathbb{F}(\mathcal{I}_i, \mathcal{I}_j, \mathcal{S}_j, \mathcal{O}_{ij}). \quad (2)$$

The  $\mathbb{F}$  is conditioned on prior encodings including the image encoding  $\mathcal{I}$  at target pixel  $i$  and source pixel  $j$ , the depth encoding  $\mathcal{S}$  at source pixel  $j$ , and the spatial offset encoding

$\mathcal{O}$  for pixel coordinates from  $i$  to  $j$ . The spatial offset encoding  $\mathcal{O}$  serves as the **spatial distance** term. Instead of designing an explicit radiometric difference term, we use MLP  $\mathbb{F}$ , to **implicitly** consider the **radiometric difference** with image encoding  $\mathcal{I}_i$  and  $\mathcal{I}_j$  as input. We believe our formulation can further explore additional information like scene context, which goes beyond the traditional bilateral filter that utilizes an explicitly designed radiometric difference term, although our initial inspiration comes from it.

The adopted MLP consists of four densely connected layers, each followed by a Batch Normalization [17] and a GeLU [12] layer, with a skip connecting layer adding the outputs from the second and the last layer. In practice,  $\alpha$ ,  $\beta$  are directly regressed. For any target pixel  $i$ ,  $\omega$  is obtained from an extra added Softmax layer. The layer operates among pixel  $i$ 's  $N$  neighboring valid pixels to ensure that  $\sum_{j \in \mathcal{N}(i)} \omega_{ij} = 1$ . This MLP is shared among any  $(i, j)$  pixel pair but takes different encodings as input to generate *spatial-variant* and *content-dependent* parameters.

The paradigm of *content-dependent* parameter generation has been used in previous depth completion methods and gained robust results. GuideNet [41] and RigNet [49] leverage *content-dependent* weights to establish a sophisticated multi-modal feature fusion. CSPN [3] and CSPN++ [4] dynamically generate *content-dependent* affinity matrix to iteratively enhance the initial depth. There are two notable differences between these existing methods and ours. Firstly, the *content-dependent* parameters are only learned from image content in these methods, typically neglecting the spatial distance between the target pixel and nearby depth measurements. In contrast, our method takes both image content and spatial distance into account. This allows the propagation of depth with the preference for nearest values in both content and spatial domains. Secondly, the *content-dependent* parameters are typically employed in the multi-modal fusion or refinement stage in existing methods while we extend this paradigm to the earliest pre-processing stage, enabling subsequent stages to overcome the sparsity problem more effectively.

### 3.2.3 Prior Encoding

The proposed bilateral propagation module is arranged in a multi-scale scheme, with prior encodings from the corresponding resolution. For the lowest resolution with  $s = 5$ , we directly adopt the image feature  $\mathbf{I}^s$  as image encoding. Otherwise, to make image encoding representative, we concatenate the multi-modal fused feature  $\mathbf{F}^{s+1}$  (introduced in Sec. 3.3) with depth feature  $\mathbf{D}^{s+1}$  in scale  $s + 1$ . The depth feature  $\mathbf{D}^{s+1}$  is achieved by inverse projecting estimated depth  $D^{s+1}$  to camera space. Then we utilize deconvolution operation to upsample the concatenated feature map to scale  $s$ . Finally, we concatenate the upsampled feature with  $\mathbf{I}^s$ , and adopt an extra convolution operation to produce the image encoding  $\mathcal{I}^s$ .

We inverse project sparse depth map  $S^s$  to camera space as depth encoding  $\mathcal{S}^s$ . Except  $S^0 = S$ , the sparse depth map  $S^s$  is downsampled from  $S$  via a weighted pooling, whose weights are generated from the image encoding  $\mathcal{I}^s$ . Specially, we adopt a periodic shuffling operator to rearrange the image feature map from shape  $\frac{H}{2^s} \times \frac{W}{2^s} \times 4^s$  to shape  $H \times W$  and apply the exponential transformation, that guarantees the generated weight map is positive and has the same resolution as  $S$ . More implementation details are introduced in the supplementary material. While the points in low-resolution depth maps are denser, we only use the valid ones for depth propagation shown in Eq. (1) and Eq. (2).

We adopt the directed distance from the target pixel to the valid source pixel, *i.e.* a 2D vector representing offsets along two axes of the image plane, as the spatial distance encoding  $\mathcal{O}$ . Position encoding, which transforms the position to a high dimensional space via Fourier transform or learned embedding, has been used in some recent popular network structures, *e.g.* Transformer [45] and NeRF [30]. We find high dimensional position encoding can induce little performance improvement in our formulation, despite the increased computational cost. We attribute this to the simplicity of a spatial distance term being adequate for co-efficient generation in our local combination model, akin to bilateral filtering.

### 3.3. Multi-Modal Fusion

For each scale, we can get a dense depth map  $D'^s$  from the bilateral propagation module which is suitable for CNN to process. We adopt image encoding  $\mathcal{I}^s$  as the image feature, and inverse project  $D'^s$  to camera space as the depth feature. Then, we follow the early-fusion scheme that simply concatenates the image feature and depth feature, and feeds the concatenated feature as input to a U-Net structure [37] for multi-modal fusion. The adopted U-Net is an encoder-decoder CNN, aggregating local and global features in multi-scale. The encoder part consists of 2 ResNet blocks [11] for feature extraction in each scale and convolution layer with stride 2 to reduce the resolution of the feature

map. The decoder part consists of a deconvolution layer with stride 2 to upsample the feature map and skip connection to fuse the upsampled feature and encoded feature with the same resolution by concatenation operation. The U-Net finally yields a fused feature map  $\mathbf{F}^s$ . And a convolution layer is applied on  $\mathbf{F}^s$  to estimate a residual depth map, which is then added to the previous depth map  $D'^s$ , resulting in a processed depth map  $D''^s$ .

### 3.4. Depth Refinement

We implement a simple convolutional spatial propagation module similar to CSPN++ [4] as the depth refinement for each scale. This propagation process updates the processed depth map  $D''$  aided by sparse valid depth map  $S$ . The updating equation at step  $t$  for a preset propagation kernel size  $k$  can be written as:

$$\begin{aligned} \hat{D}_{i,k,t} &= \kappa_{i,k} \hat{D}_{i,k,t-1} + \sum_{j \in \mathcal{N}_k(i) \setminus i} \kappa_{j,k} \hat{D}_{j,k,t-1}, \\ \kappa_{i,k} &= 1 - \sum_{j \in \mathcal{N}_k(i) \setminus i} \kappa_{j,k}, \\ \kappa_{j,k} &= \frac{\hat{\kappa}_{j,k}}{\sum_{j \in \mathcal{N}_k(i) \setminus i} |\hat{\kappa}_{j,k}|}, \end{aligned} \quad (3)$$

where  $\kappa$  is the content-dependent affinity map and  $\hat{\kappa}$  is generated by convolutional layers depending on the fused feature  $\mathbf{F}$ .  $\hat{D}_t$  is initialized with  $D''$  for  $t = 0$ . The  $l^1$ -norm constraint for  $\hat{\kappa}$  guarantees the stability of propagation process [3]. Note that  $\mathcal{N}_k$  used here is different from  $\mathcal{N}$  used in our bilateral propagation module.  $\mathcal{N}_k$  is the set of neighboring pixels in a  $k \times k$  local window, irrelevant to the validation of sparse depth measurement.

For each refinement step, we also update the propagated depth  $\hat{D}_t$  by embedding sparse depth measurement after performing Eq. (3),

$$\hat{D}_{i,k,t} \leftarrow (1 - \gamma_{i,k} \mathbb{I}(S_i)) \hat{D}_{i,k,t} + \gamma_{i,k} \mathbb{I}(S_i) S_i, \quad (4)$$

where  $\gamma$  works as confidence generated by convolutional layers followed by a sigmoid layer based on the fused feature  $\mathbf{F}$ . The whole process iterates  $T$  steps, and the final depth is a combination of multi-kernel and multi-step:

$$D_i = \sum_{t \in \mathcal{T}} \sum_{k \in \mathcal{K}} \tau_{i,t} \sigma_{i,k} \hat{D}_{i,k,t}. \quad (5)$$

Here,  $\mathcal{K}$  is a set of  $\{3, 5, 7\}$  for 3 different kernel sizes,  $\mathcal{T}$  is a set of  $\{0, \lfloor T/2 \rfloor, T\}$  indicating different iteration steps.  $\tau$  and  $\sigma$  work as confidence maps generated by convolutional layers followed by a Softmax layer, normalized across different iteration steps and different kernel sizes respectively. Empirically, we set a low iteration number  $T$  for refinement on low-resolution depth map, that  $T$  is from 2 to 12 with an incremental step of 2 for a total of 6 scales from low-resolution to high-resolution.

### 3.5. Implementation Details

#### 3.5.1 Loss Function

Our method is trained in an end-to-end manner, with a multi-scale loss to provide adequate supervision on the depth map estimated in each scale. The loss function is

$$L = \sum_{s=0}^5 \sum_{i \in \mathcal{P}_v} \lambda_s \|D_i^{gt} - \mathbb{U}_s(D^s)_i\|^2, \quad (6)$$

where  $\mathcal{P}_v$  represents the set of valid pixels in the ground truth depth map  $D^{gt}$ .  $\mathbb{U}_s$  is bilinear interpolation operation to upsample predicted depth map in scale  $s$  to the same resolution as  $D^{gt}$ , and  $\lambda_s$  is hyper-parameter to balance loss in each scale, which is set as  $4^{-s}$  empirically.

#### 3.5.2 Training Setting

We implement our method in pytorch [32] and train it on a GPU workstation with 4 Nvidia RTX 3090 GPUs. Our network is mainly stacked with ResNet blocks [11], in which a DropPath [20] is added before residual addition as a regularization for training. We adopt AdamW [26] with 0.05 weight decay as the optimizer, and clip gradient whose  $l^2$ -norm is larger than 0.1. Our method is trained from scratch in roughly 300K iterations with OneCycle learning rate policy [40] gradually reducing the learning rate to 25% of the largest learning rate. We set batch size as 8 and the largest learning rate as 0.001 for KITTI dataset. For NYUv2 dataset, whose image resolution is lower, we increase the batch size to 16 and the largest learning rate to 0.002 correspondingly. The final model is yielded by Exponential Moving Average (EMA) with 0.9999 decay.

## 4. Experiments

### 4.1. Datasets & Evaluation Metrics

We conduct comprehensive experiments to verify our method on both indoor and outdoor scenes.

**Indoor scene.** The NYUv2 dataset [39] is an indoor dataset containing 464 scenes gathered by a Kinect sensor. We take 50K frames sampled from 249 scenes by Ma *et al.* [27] as the training data, and evaluate on the official test set consisting of 654 samples from 215 scenes. Following the common practice [31, 41, 53], images are down-sampled to  $320 \times 240$  and then center-cropped to  $304 \times 228$ , and for each frame, sparse depth is generated by random sampling 500 points from the ground truth depth map. Due to the input resolution for our network being a multiple of 32, we further pad the image to  $320 \times 256$  as input, but evaluate only the valid region of size  $304 \times 228$  to keep a fair comparison with other methods. The standard evaluation

	KITTI				NYUv2		
	RMSE↓ (mm)	MAE↓ (mm)	iRMSE↓ (1/km)	iMAE↓ (1/km)	RMSE↓ (m)	REL↓	$\delta_{1.25}$ ↑ (%)
S2D [27]	814.73	249.95	2.80	1.21	0.230	0.044	97.1
CSPN [3]	1019.64	279.46	2.93	1.15	0.117	0.016	99.2
DeepLiDAR [33]	758.38	226.50	2.56	1.15	0.115	0.022	99.3
CSPN++ [4]	743.69	209.28	2.07	0.90	0.115	–	–
GuideNet [41]	736.24	218.83	2.25	0.99	0.101	0.015	99.5
FCFR [13]	735.81	217.15	2.20	0.98	0.106	0.015	99.5
ACMNet [54]	744.91	206.09	2.08	0.90	0.105	0.015	99.4
NLSPN [31]	741.68	199.59	1.99	0.84	0.092	0.012	<b>99.6</b>
RigNet [49]	712.66	203.25	2.08	0.90	0.090	0.013	<b>99.6</b>
DySPN [22]	709.12	192.71	1.88	0.82	0.090	0.012	<b>99.6</b>
BEV@DC [56]	697.44	<b>189.44</b>	1.83	0.82	<b>0.089</b>	0.012	<b>99.6</b>
CFormer [53]	708.87	203.45	2.01	0.88	0.090	0.012	–
LRRU [46]	696.51	189.96	1.87	<b>0.81</b>	0.091	<b>0.011</b>	<b>99.6</b>
BP-Net	<b>684.90</b>	194.69	<b>1.82</b>	0.84	<b>0.089</b>	0.012	<b>99.6</b>

Table 1. **Performance on KITTI and NYUv2 datasets.** For the KITTI dataset, results are evaluated by the KITTI testing server and ranked by the RMSE (in *mm*). For the NYUv2 dataset, authors report their performance on the official test set in their papers. The best result under each criterion is in **bold**.

metrics are root mean squared error (RMSE), mean absolute relative error (REL), and  $\delta_\theta$  represents the percentage of pixels whose error is less than a threshold  $\theta$ .

**Outdoor scene.** The KITTI depth completion (DC) dataset [44] is collected by an autonomous driving vehicle, whose ground truth depth is from temporally registered LiDAR scans, further verified by stereo image pairs. This dataset provides 86, 898 and 1,000 frames for training and validation respectively. Another 1,000 frames for testing are evaluated on the remote server with a public leaderboard<sup>2</sup> for ranking. We randomly crop frames to  $1216 \times 256$  for training and directly use the full-resolution frames as input for testing. The standard evaluation metrics are root mean squared error (RMSE), mean absolute error (MAE), root mean squared error of the inverse depth (iRMSE), and mean absolute error of the inverse depth (iMAE), among which RMSE is chosen as the primary metric for ranking.

### 4.2. Comparison with SOTAs

We evaluate our method on the test set of the NYUv2 dataset and KITTI DC dataset separately. Tab. 1 lists the quantitative comparison of our method and other top-ranking published methods. On the KITTI DC leaderboard, our method ranks 1st and exceeds all other methods under the primary RMSE metric at the time of paper submission. It also has comparable performance under other evaluation metrics. On the NYUv2 dataset, our method achieves the best RMSE and the best  $\delta_{1.25}$ . Due to the diverse viewpoints and sparser input depth map in the NYUv2 dataset, it is more challenging to obtain a large performance gain in the NYUv2 dataset compared with the KITTI dataset.

Fig. 4 provides visual comparisons with other SOTA open-source methods on the validation set of KITTI DC and the test set of NYUv2. We adopt the public code and best

<sup>2</sup>[http://www.cvlibs.net/datasets/kitti/eval\\_depth.php?benchmark](http://www.cvlibs.net/datasets/kitti/eval_depth.php?benchmark)

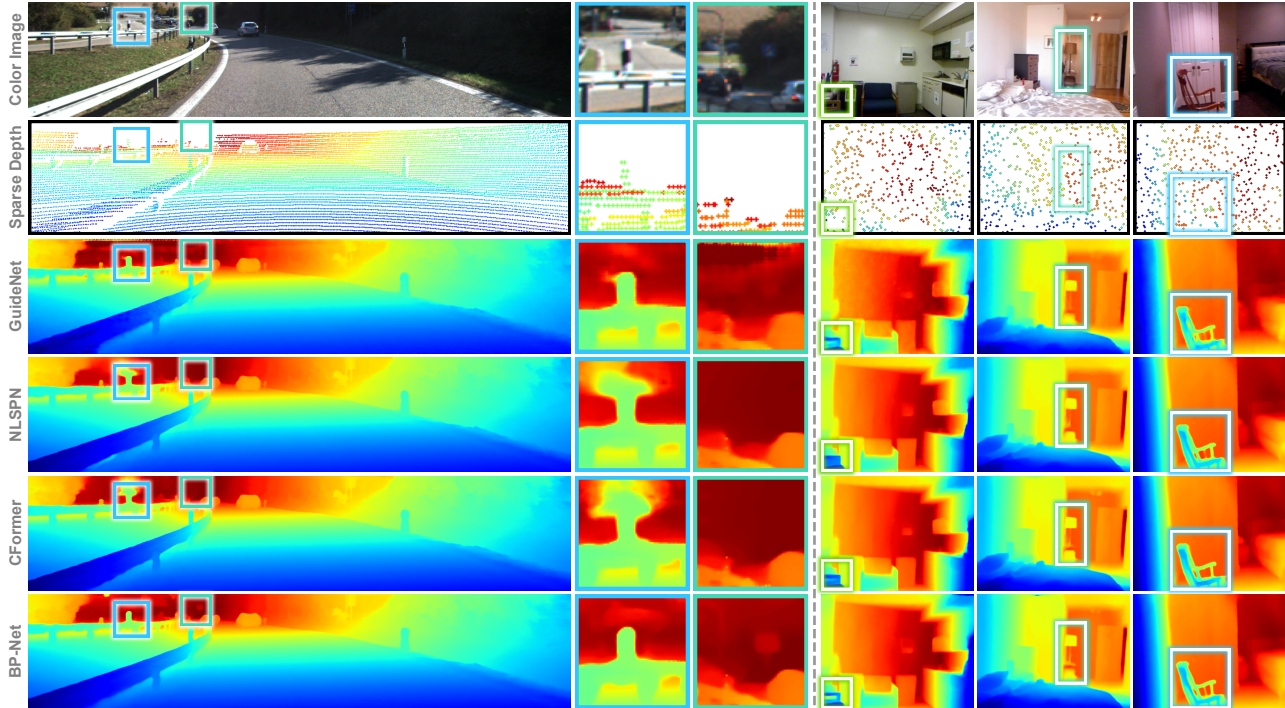


Figure 4. Qualitative comparison with ‘GuideNet’ [41], ‘NLSPN’ [31] and ‘CFormer’ [53] on indoor and outdoor scenes. Sparse depth points are enlarged for better visualization. Our method is shown in the last row, with highlighted rectangles for easy comparison.




Figure 5. The RMSE performance under different numbers of neighboring pixels.

models from their authors to produce results and keep the sparse depth map the same for comparing methods. Our results listed in the last row exhibit clearer object boundaries and richer details, while other methods struggle to estimate accurate depth in these challenging regions.

### 4.3. Ablation Studies

We conduct ablation studies on the NYUv2 dataset to reveal the effects of the different components in our method.

#### 4.3.1 Different Number for Neighboring Pixels

In our bilateral propagation module, the target depth is propagated from  $N$  nearest neighboring valid depth measurements according to our depth parametrization in Eq. (1). We experiment with different values of  $N$  and plot the corresponding RMSEs in Fig. 5. While different choices of  $N$  might slightly change the results at the *Pre.* stage, they can

Content Propagation	Spatial Propagation	Multi-scale Architecture	Multi-scale Loss	RMSE↓ (mm)
✓		✓	✓	90.40
	✓	✓	✓	90.93
✓	✓			91.94
✓	✓	✓		92.46
✓	✓	✓	✓	88.69

Table 2. Ablation studies on NYUv2 dataset.

all improve the *MF.* stage for more effective depth completion. Thus, our method is not very sensitive to this hyperparameter. The performance gain is most noticeable when  $N$  increases from 1 to 2, after which it gradually decreases and trends toward saturation. Finally, we choose  $N = 4$  in our method.

#### 4.3.2 Effect of Bilateral Parameters Generation

Our bilateral propagation module is a non-linear model, whose coefficients are dynamically generated depending on image content and spatial distance as defined in Eq. (2). To verify the contribution of these terms, We train a variant named *content propagation*, where the coefficients depend only on  $I_i, I_j, S_j$ . We further train another variant named *spatial propagation*, where the coefficients depend only on  $O_{ij}$ . Tab. 2 lists the performance of these two variants and reveals that both *content propagation* and *spatial propagation* are useful, and removing any of them decreases the performance.

Pre.	MF.	Post.	RMSE↓ (mm)	Params.↓ (M)	Madds.↓ (G)	GPU RAM.↓ (MB)	Runtime.↓ (ms)
✓	✓		91.68(-0.00)	76.15(+ 0.00)	105.76(+ 0.00)	5206(+ 0)	7.78(+ 0.00)
	✓		89.56(-2.12)	77.26(+ 1.11)	113.15(+ 7.39)	5215(+ 9)	12.95(+ 5.17)
	✓	✓	90.04(-1.68)	88.77(+12.62)	129.73(+23.87)	5302(+ 96)	14.36(+ 6.58)
✓	✓	✓	88.69(-2.99)	89.87(+13.72)	137.12(+31.36)	5312(+106)	19.68(+11.90)

Table 3. Effect of 3-stage estimation.

### 4.3.3 Effect of Multi-Scale Depth Estimation

We also verify the effectiveness of our multi-scale scheme, which estimates dense depth in a multi-scale architecture with a multi-scale loss. We train a variant by replacing the multi-scale loss with a loss only on the original scale and train another variant by further replacing the multi-scale architecture with depth estimation only on the original scale. Based on the results listed in Tab. 2, we can see the integration of both multi-scale architecture and multi-scale loss enhances system performance while merely introducing the multi-scale architecture without incorporating multi-scale loss decreases the performance. In BP-Net, multi-scale architecture facilitates coarse-to-fine depth estimation, where low-resolution results are upsampled to guide the propagation in high-resolution. Consequently, multi-scale loss provides direct supervision at each resolution level, making the low-resolution results meaningful and improving the guidance effectively. Omitting the multi-scale loss might make the network ignore low-resolution results, and increase training difficulty.

### 4.3.4 Effect of 3-stage Depth Estimation

At each scale, dense depth is estimated via three sequential stages, consisting of pre-processing (*Pre.*), multi-modal fusion (*MF.*), and post-processing (*Post.*). We train different variants with different combinations of these three stages and analyze the performance gain on RMSE and extra cost on parameters (*Params.*), Multiply-adds operations (*Madds.*), GPU RAM and runtime. As shown in Tab. 3, combining *MF.* with *Pre.* can achieve 2.12mm performance gain on RMSE at the cost of 1.11M more *Params.* and 7.39G more *Madds.* As a comparison, combining *MF.* with *Post.* produces less performance gain with more cost on *Params.* and *Madds.* In addition, *Pre.+MF.* achieves better RMSE than *MF.+Post.* with less GPU RAM and runtime. And combining all three stages achieves the best RMSE. These comparison verify the effectiveness of our 3-stage design and prove the importance of depth propagation at the preprocessing stage in contrast to the refinement stage.

### 4.4. Input Sparsity Analysis

In real applications, the depth sparsity level might change, i.e. the number of points in the input sparse depth map might vary. We compare our method with others under various sparsity levels on the NYUv2 dataset. In this comparison, all models are trained with 500 depth points and then




Figure 6. The performance under various sparsity levels.  $\delta_{1.25}$  is plotted in the upper part, and RMSE is in the lower part. All comparing methods are directly evaluated without retraining.

evaluated under various sparsity levels ranging from 250 to 750 depth points. For a thorough evaluation, given a sparsity level, each test image is sampled 100 times with different random seeds to generate the input sparse depth map. The performance of each method is averaged on these 100 randomly sampled inputs to reduce the potential bias due to random sampling. As shown in Fig. 6, our method consistently achieves the lowest RMSE and the highest  $\delta_{1.25}$  across various sparsity levels. The differences are more significant when the sparsity level is higher. This comparison demonstrates the strong generalization capability of our method across various sparsity levels.

## 5. Conclusion

This paper introduces a Bilateral Propagation Network (BP-Net) for depth completion. BP-Net propagates depth at the earliest stage, rather than the refinement stage, avoiding the following multi-modal stage from the sparsity problem. The proposed bilateral propagation module can dynamically predict propagation coefficients conditioned on both *radiometric difference* and *spatial distance*, to enable depth propagation with the preference of nearest values on both domain and range. Experimental results demonstrate the outstanding performance of BP-Net and also suggest the importance of propagation at the earliest stage in contrast to the refinement stage. BP-Net, mostly consisting of local propagation and convolution operations, may be limited to local structure and have difficulty with long-range information delivery. Combining BP-Net with global and non-local operations, such as Transformer for multi-modal fusion [53] and non-local propagation for depth refinement [31], are directions for future works.

**Acknowledgements:** This work was supported in part by National Natural Science Foundation of China (NSFC) under Grants 61973311, 62273353, and 62103431.



## References

- [1] Yun Chen, Bin Yang, Ming Liang, and Raquel Urtasun. Learning joint 2d-3d representations for depth completion. In *ICCV*, pages 10023–10032, 2019. 3
- [2] Zhao Chen, Vijay Badrinarayanan, Gilad Drozdov, and Andrew Rabinovich. Estimating depth from rgb and sparse sensing. In *ECCV*, pages 167–182, 2018. 2
- [3] Xinjing Cheng, Peng Wang, and Ruigang Yang. Learning depth with convolutional spatial propagation network. *IEEE TPAMI*, 42(10):2361–2379, 2019. 1, 3, 4, 5, 6
- [4] Xinjing Cheng, Peng Wang, Chenye Guan, and Ruigang Yang. Cspn++: Learning context and resource aware convolutional spatial propagation networks for depth completion. In *AAAI*, pages 10615–10622, 2020. 3, 4, 5, 6
- [5] Andrea Conti, Matteo Poggi, and Stefano Mattoccia. Sparsity agnostic depth completion. In *WACV*, pages 5871–5880, 2023. 2
- [6] James Diebel and Sebastian Thrun. An application of markov random fields to range sensing. *NeurIPS*, 18, 2005. 2
- [7] Abdelrahman Eldesokey, Michael Felsberg, and Fahad Shahbaz Khan. Confidence propagation through cnns for guided sparse depth regression. *IEEE TPAMI*, 42(10):2423–2436, 2019. 2
- [8] David Ferstl, Christian Reinbacher, Rene Ranftl, Matthias R  ther, and Horst Bischof. Image guided depth upsampling using anisotropic total generalized variation. In *ICCV*, pages 993–1000, 2013. 2
- [9] Andreas Geiger, Philip Lenz, and Raquel Urtasun. Are we ready for autonomous driving? the kitti vision benchmark suite. In *CVPR*, 2012. 1
- [10] Kaiming He, Jian Sun, and Xiaoou Tang. Guided image filtering. *IEEE TPAMI*, 35(6):1397–1409, 2012. 2
- [11] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *CVPR*, pages 770–778, 2016. 5, 6
- [12] Dan Hendrycks and Kevin Gimpel. Gaussian error linear units (gelus). *arXiv preprint arXiv:1606.08415*, 2016. 4
- [13] Mu Hu, Shuling Wang, Bin Li, Shiyu Ning, Li Fan, and Xiaojin Gong. Penet: Towards precise and efficient image guided depth completion. In *ICRA*, pages 13656–13662. IEEE, 2021. 6
- [14] Zixuan Huang, Junming Fan, Shenggan Cheng, Shuai Yi, Xiaogang Wang, and Hongsheng Li. Hms-net: Hierarchical multi-scale sparsity-invariant network for sparse depth completion. *IEEE TIP*, 29:3429–3441, 2019. 2
- [15] Benjamin Huhle, Timo Schairer, Philipp Jenke, and Wolfgang Stra  ber. Fusion of range and color images for denoising and resolution enhancement with a non-local filter. *Computer vision and image understanding*, 114(12):1336–1345, 2010. 2
- [16] Saif Imran, Yunfei Long, Xiaoming Liu, and Daniel Morris. Depth coefficients for depth completion. In *CVPR*, pages 12438–12447. IEEE, 2019. 1, 2
- [17] Sergey Ioffe and Christian Szegedy. Batch normalization: Accelerating deep network training by reducing internal covariate shift. In *International conference on machine learning*, pages 448–456. pmlr, 2015. 4
- [18] Johannes Kopf, Michael F Cohen, Dani Lischinski, and Matt Uyttendaele. Joint bilateral upsampling. *ACM TOG*, 26(3):96–es, 2007. 2
- [19] Jason Ku, Ali Harakeh, and Steven L Waslander. In defense of classical image processing: Fast depth completion on the cpu. In *2018 15th Conference on Computer and Robot Vision (CRV)*, pages 16–22. IEEE, 2018. 2
- [20] Gustav Larsson, Michael Maire, and Gregory Shakhnarovich. Fractalnet: Ultra-deep neural networks without residuals. *arXiv preprint arXiv:1605.07648*, 2016. 6
- [21] Yijun Li, Jia-Bin Huang, Narendra Ahuja, and Ming-Hsuan Yang. Joint image filtering with deep convolutional networks. *IEEE TPAMI*, 41(8):1909–1923, 2019. 2
- [22] Yuankai Lin, Tao Cheng, Qi Zhong, Wending Zhou, and Hua Yang. Dynamic spatial propagation network for depth completion. In *AAAI*, pages 1638–1646, 2022. 1, 3, 6
- [23] Lina Liu, Yiyi Liao, Yue Wang, Andreas Geiger, and Yong Liu. Learning steering kernels for guided depth completion. *IEEE TIP*, 30:2850–2861, 2021. 2
- [24] Sifei Liu, Shalini De Mello, Jinwei Gu, Guangyu Zhong, Ming-Hsuan Yang, and Jan Kautz. Learning affinity via spatial propagation networks. *NeurIPS*, 30, 2017. 3
- [25] Xin Liu, Xiaofei Shao, Bo Wang, Yali Li, and Shengjin Wang. Graphcspn: Geometry-aware depth completion via dynamic gcns. In *ECCV*, pages 90–107. Springer, 2022. 3
- [26] Ilya Loshchilov and Frank Hutter. Decoupled weight decay regularization. In *ICLR*, 2018. 6
- [27] Fangchang Ma and Sertac Karaman. Sparse-to-dense: Depth prediction from sparse depth samples and a single image. In *ICRA*, pages 4796–4803. IEEE, 2018. 1, 2, 6
- [28] Fangchang Ma, Guilherme Venturilli Cavalheiro, and Sertac Karaman. Self-supervised sparse-to-dense: Self-supervised depth completion from lidar and monocular camera. In *ICRA*, pages 3288–3295. IEEE, 2019. 2
- [29] Jeffrey Mahler, Matthew Matl, Vishal Satish, Michael Danielczuk, Bill DeRose, Stephen McKinley, and Ken Goldberg. Learning ambidextrous robot grasping policies. *Science Robotics*, 4(26):eaau4984, 2019. 1
- [30] 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. In *ECCV*, 2020. 5
- [31] Jinsun Park, Kyungdon Joo, Zhe Hu, Chi-Kuei Liu, and In So Kweon. Non-local spatial propagation network for depth completion. In *ECCV*, pages 120–136. Springer, 2020. 1, 3, 6, 7, 8
- [32] Adam Paszke, Sam Gross, Francisco Massa, Adam Lerer, James Bradbury, Gregory Chanan, Trevor Killeen, Zeming Lin, Natalia Gimelshein, Luca Antiga, Alban Desmaison, Andreas Kopf, Edward Yang, Zachary DeVito, Martin Raison, Alykhan Tejani, Sasank Chilamkurthy, Benoit Steiner, Lu Fang, Junjie Bai, and Soumith Chintala. Pytorch: An imperative style, high-performance deep learning library. In *NeurIPS*, pages 8024–8035. Curran Associates, Inc., 2019. 6

- [33] Jiaxiong Qiu, Zhaopeng Cui, Yinda Zhang, Xingdi Zhang, Shuaicheng Liu, Bing Zeng, and Marc Pollefeys. DeepLidar: Deep surface normal guided depth prediction for outdoor scene from sparse lidar data and single color image. In *CVPR*, pages 3313–3322, 2019. 2, 3, 6
- [34] Chao Qu, Ty Nguyen, and Camillo Taylor. Depth completion via deep basis fitting. In *WACV*, pages 71–80, 2020. 2
- [35] Chao Qu, Wenxin Liu, and Camillo J Taylor. Bayesian deep basis fitting for depth completion with uncertainty. In *ICCV*, pages 16147–16157, 2021. 2
- [36] Kyeongha Rho, Jinsung Ha, and Youngjung Kim. Guideformer: Transformers for image guided depth completion. In *CVPR*, pages 6250–6259, 2022. 3
- [37] Olaf Ronneberger, Philipp Fischer, and Thomas Brox. U-net: Convolutional networks for biomedical image segmentation. In *International Conference on Medical image computing and computer-assisted intervention*, pages 234–241. Springer, 2015. 5
- [38] Johannes L Schonberger and Jan-Michael Frahm. Structure-from-motion revisited. In *CVPR*, pages 4104–4113, 2016. 1
- [39] Nathan Silberman, Derek Hoiem, Pushmeet Kohli, and Rob Fergus. Indoor segmentation and support inference from rgb-d images. In *ECCV*, pages 746–760, 2012. 6
- [40] Leslie N Smith and Nicholay Topin. Super-convergence: Very fast training of neural networks using large learning rates. In *Artificial intelligence and machine learning for multi-domain operations applications*, pages 369–386. SPIE, 2019. 6
- [41] Jie Tang, Fei-Peng Tian, Wei Feng, Jian Li, and Ping Tan. Learning guided convolutional network for depth completion. *IEEE TIP*, 30:1116–1129, 2020. 1, 2, 4, 6, 7
- [42] Carlo Tomasi and Roberto Manduchi. Bilateral filtering for gray and color images. In *ICCV*, pages 839–846. IEEE, 1998. 2, 4
- [43] Ivana Tomic and Sarah Drewes. Learning joint intensity-depth sparse representations. *IEEE TIP*, 23(5):2122–2132, 2014. 2
- [44] Jonas Uhrig, Nick Schneider, Lukas Schneider, Uwe Franke, Thomas Brox, and Andreas Geiger. Sparsity invariant cnns. In *IEEE International Conference on 3D Vision (3DV)*, pages 11–20, 2017. 2, 6
- [45] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. Attention is all you need. *NeurIPS*, 30, 2017. 3, 5
- [46] Yufei Wang, Bo Li, Ge Zhang, Qi Liu, Tao Gao, and Yuchao Dai. Lrru: Long-short range recurrent updating networks for depth completion. In *ICCV*, pages 9422–9432, 2023. 2, 3, 6
- [47] Yu-Shiang Wong, Changjian Li, Matthias Niessner, and Niloy J Mitra. Rigidfusion: Rgb-d scene reconstruction with rigidly-moving objects. In *Computer Graphics Forum*, pages 511–522. Wiley Online Library, 2021. 1
- [48] Yan Xu, Xinge Zhu, Jianping Shi, Guofeng Zhang, Hujun Bao, and Hongsheng Li. Depth completion from sparse lidar data with depth-normal constraints. In *ICCV*, pages 2811–2820, 2019. 2
- [49] Zhiqiang Yan, Kun Wang, Xiang Li, Zhenyu Zhang, Jun Li, and Jian Yang. Rignet: Repetitive image guided network for depth completion. In *ECCV*, pages 214–230. Springer, 2022. 2, 4, 6
- [50] Jingyu Yang, Xinchen Ye, Kun Li, Chunping Hou, and Yao Wang. Color-guided depth recovery from rgb-d data using an adaptive autoregressive model. *IEEE TIP*, 23(8):3443–3458, 2014. 2
- [51] Yanchao Yang, Alex Wong, and Stefano Soatto. Dense depth posterior (ddp) from single image and sparse range. In *CVPR*, pages 3353–3362, 2019. 2
- [52] Yinda Zhang and Thomas Funkhouser. Deep depth completion of a single rgb-d image. In *CVPR*, pages 175–185, 2018. 2
- [53] Youmin Zhang, Xianda Guo, Matteo Poggi, Zheng Zhu, Guan Huang, and Stefano Mattoccia. Completionformer: Depth completion with convolutions and vision transformers. In *CVPR*, pages 18527–18536, 2023. 3, 6, 7, 8
- [54] Shanshan Zhao, Mingming Gong, Huan Fu, and Dacheng Tao. Adaptive context-aware multi-modal network for depth completion. *IEEE TIP*, 30:5264–5276, 2021. 3, 6
- [55] Yiming Zhao, Lin Bai, Ziming Zhang, and Xinming Huang. A surface geometry model for lidar depth completion. *IEEE Robotics and Automation Letters*, 6(3):4457–4464, 2021. 2
- [56] Wending Zhou, Xu Yan, Yinghong Liao, Yuankai Lin, Jin Huang, Gangming Zhao, Shuguang Cui, and Zhen Li. Bev@dc: Bird’s-eye view assisted training for depth completion. In *CVPR*, pages 9233–9242, 2023. 6