

Towards Domain-agnostic Depth Completion*

Wei Yin*, Jianming Zhang[†], Oliver Wang[‡], Simon Niklaus[‡], Simon Chen[‡], Chunhua Shen^{*†}

* DJI

* Zhejiang University, China

[‡] Adobe Research

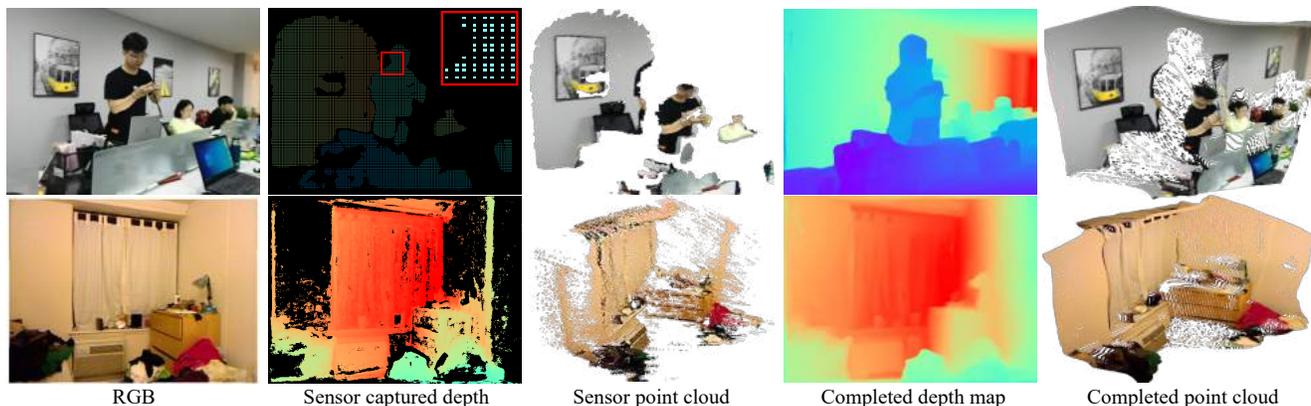


Figure 1. Our method fills in missing information in different types of sparse depth maps. A single model can be used to complete the sparse depth from different methods, e.g. Huawei Mate30 Time-of-Flight sensor (top) and a multi-view stereo algorithm [31] (bottom).

Abstract

Existing depth completion methods are often targeted at a specific sparse depth type, and generalize poorly across task domains. We present a method to complete sparse/semi-dense, noisy, and potentially low-resolution depth maps obtained by various range sensors, including those in modern mobile phones, or by multi-view reconstruction algorithms. Our method leverages a data driven prior in the form of a single image depth prediction network trained on large-scale datasets, the output of which is used as an input to our model. We propose an effective training scheme where we simulate various sparsity patterns in typical task domains. In addition, we design two new benchmarks to evaluate the generalizability and the robustness of depth completion methods. Our simple method shows superior cross-domain generalization ability against state-of-the-art depth completion methods, introducing a practical solution to high quality depth capture on a mobile device. Code is available at: <https://github.com/YvanYin/FillDepth>.

*Part of this work was done when WY and CS were with The University of Adelaide.

[†]Correspondence should be addressed to C. Shen.

1. Introduction

Accurate metric depth is important for many computer vision applications, in particular 3D perception [34, 41] and reconstruction [24, 25]. Typically, depth is obtained by using direct range sensors such as LiDAR and Time-of-Flight (ToF) sensors on modern mobile phones, or computed by multi-view stereo methods [31, 46]. However, usually these sources can only provide incomplete and/or sparse depth information. For example, LiDAR sensors capture depth in a linear scanning pattern, ToF sensors are lower resolution and fail at specular or distant surfaces, and multi-view reconstruction methods [31, 46, 53] only provide confident depth at textured regions and are range limited by the camera baseline (the iPhone rear stereo camera has a maximum depth of 2.5 meters).

Existing depth completion methods can be classified into two categories according to the input sparsity pattern: depth inpainting methods that fill large holes [14, 33, 54], and sparse depth densification methods that densify sparsely distributed depth measurements [5, 7, 26, 27, 44]. When working on a specific sparsity pattern, e.g., on either NYU [35] or KITTI [38], recent approaches [5, 6, 16, 26, 27] can obtain impressive performance. However, in real-world scenarios, the sparsity pattern may be subject to change or unknown at training time, as it is a function of hardware, software, as well as the configuration of the scene itself.

In this paper, we revisit existing methods and analyze the gap between the performance on the training setup and downstream applications, and we find that existing depth completion methods suffer from the following key limitations. First, their methods work best on one specific sparsity pattern, but they are sensitive to mild perturbations of the depth sampling and generalize poorly to other types of sparse depth (e.g., from Kinect depths to smart phone depths). Second, they are sensitive to noise and outliers produced by the depth capture process. To address these issues, we propose a simple yet effective method towards robust cross-domain depth completion.

Our method provides the following improvements. First, inspired by domain randomization methods [36, 37, 51], we analyze the existing set of common sparsity patterns and create a diverse set of synthetic sparsity patterns to train our model. To improve the cross-domain generalization ability, we follow recent monocular depth prediction methods [28, 49] and utilize a diverse training dataset which consist of multiple depth sources. Furthermore, to make our method robust to noise, we leverage the depth map predicted by a well-trained single image depth prediction method as a data-driven scene prior. Such approaches learn a strong prior on diverse scenes [28], but their predicted depths have unknown shift and scale due to the training data used (stereo images with unknown baseline). By incorporating sparse metric depth cues and a single image relative depth prior, our method is able to robustly produce a dense metric depth map.

Furthermore, as current benchmarks on NYU [35] and Matterport3D [2] cannot evaluate a method’s robustness to different sparsity types, noises, and diverse scenes, we redesign a benchmark based on existing datasets. We combine 3 datasets, i.e. NYU [35], ScanNet [8], and DIODE [39], which include both indoor and outdoor scenes. We simulate 3 different sparsity types for model evaluation based on typical application scenarios. Furthermore, we also set up another benchmark to evaluate a model’s robustness to noisy sparse depth inputs, which can be produced by SLAM, SfM methods or multi-view stereo methods. Therefore, we collect 16 NYU videos and employ COLMAP [30] to create noisy sparse depth for completion.

Our simple approach sets up a strong baseline on our new benchmarks for robust domain-agnostic depth completion. We also show that our method can generalize well on sparse depth captured by various smart phones, providing a practical solution to high quality depth sensing on low-cost mobile devices.

In conclusion, our main contributions are as follows.

- We analyze existing depth completion methods in terms of generalization across different domains and robustness to noise.
- We propose a method that incorporates a data driven

single-image prior and effective data augmentation techniques for domain-agnostic depth completion.

- We design two new synthetic benchmarks for evaluating the robustness and the generalizability of a depth completion method. Our new benchmark design simulates challenges in various real-world scenarios.

We hope that our new paradigm of training and benchmarking depth completion models can also evoke new thoughts on improving single depth estimation methods using sparse depth measurements [15, 28, 49].

2. Related Work

Sparse Depth Completion. Depth completion [4, 16] aims to densify a sparse depth input. As the sparse depths captured by different solutions have varying sparsity patterns, several methods are proposed to solve these problems. Depth maps captured with low-cost LiDAR only have a few hundred or thousand measurements per image. Several methods [3, 5, 7, 26, 45] propose to leverage texture information to complete these types of sparsity patterns. Besides such very sparse depth types, commodity-level RGB-D cameras such as Kinect, RealSense, and Tango produce depth images that are semi-dense but missing certain regions. This often happens due to objects with low reflective properties and objects beyond the maximum supported distance. Several methods treat this as a depth inpainting task and leverage smoothness priors [13], background surface extrapolation [23], and surface normals [54]. These methods have shown promising results, but they focus on only a single sparse depth type. In contrast, we design a unified solution for all these depth sparsity patterns. Additionally, we propose to use a pretrained scale-shift-invariant depth prediction model as a scene prior to improve the depth completion quality.

Monocular Depth Estimation. Monocular depth estimation aims to predict the depth from a single RGB image. Eigen *et al.* [9] proposed the first multi-scale neural network to predict depth from a single image. Subsequently, many supervised [20, 43, 47] and self-supervised methods [1, 11] have been proposed to address this task. One problem is the limited availability of training data (especially metric depth), so recent methods [28, 48, 50] have shown that much better generalization properties can be obtained by training on diverse large-scale datasets that include data with an unknown shift and scale (e.g., from stereo images). As a result, these methods can only predict relative depth, as opposed to metric depth. We propose to leverage these advancements and use the predicted depth as a guidance map for our completion network.

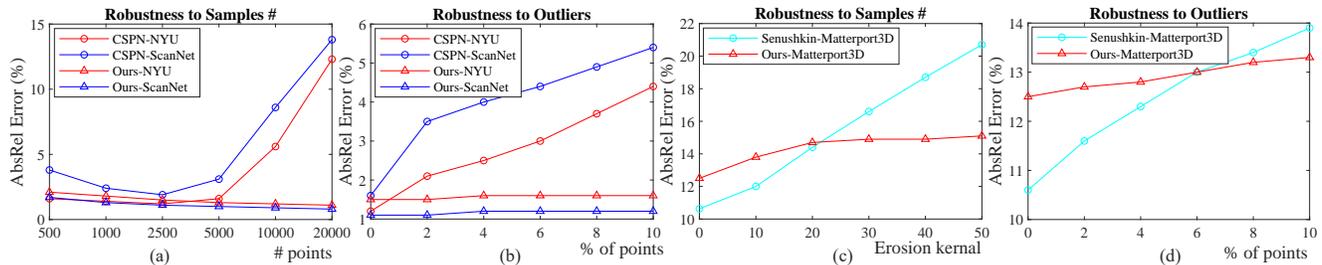


Figure 2. Robustness analysis. We analyze the performance of CSPN [5] (completion) and Senushkin *et al.* [33] (inpainting) in terms of input point numbers/patterns (a, c) and outlier ratios (b, d). CSPN is trained on NYU [35], and we evaluate it on both NYU and ScanNet [8]. Senushkin *et al.* is trained and evaluated on Matterport3D [2].

Sparsity type 1	Uniform (2500 points, AbsRel (%)↓)		
	NYU	ScanNet	Matterport3D
CSPN [6]	1.1	2.0	27.9
Senushkin <i>et al.</i> [33]	77.3	62.1	69.5
Ours	1.5	1.1	2.2
Sparsity type 2	Features (2500 points, AbsRel (%)↓)		
	NYU	ScanNet	Matterport3D
CSPN [6]	1.9	13.8	51.9
Senushkin <i>et al.</i> [33]	76.9	60.8	66.7
Ours	1.5	1.6	2.7

Table 1. Robustness to different sparse depth patterns (AbsRel error). CSPN is trained on NYU to complete very sparse depth, while Senushkin *et al.* [33] is trained on Matterport3D to complete holes. They are tested on 3 datasets with two different sparse depth types. Sparse depth of the ‘Uniform’ type are sampled uniformly from the ground truth depth, while ‘Features’ are sampled from FAST feature points. Both methods are sensitive to the sparsity pattern and dataset domain.

3. Analysis of Existing Methods

In this section, we evaluate two state-of-the-art depth completion methods to assess their performance with respect to different sparsity patterns, dataset generalization, and robustness to noise. To do this, we perturb the sparsity pattern of the input depth in various ways and add additional noise to it. We also evaluate methods on their zero-shot cross-dataset generalization performance [28] (evaluating on a different dataset than what the models were trained on). We chose two methods for this analysis, CSPN [7] and Senushkin *et al.* [33]. The former is designed to complete very sparse depth with only hundreds of sparse points, while the latter is designed to complete contiguous holes. We use the code and the model weights provided by the authors for this evaluation. As a reference, we also report the performance of our method, which is going to be presented in Sec. 4. Note that our method is not trained on any of the datasets mentioned below.

For CSPN [7] which is trained on NYU [35], we use NYU and ScanNet [8] for testing, and we vary the number of measured/input points from 500 to 20000. Senushkin *et al.* [33] is trained on Matterport3D with the task of complet-

ing holes of depth maps. We use Matterport 3D for testing, and erode the valid depth regions with different kernel sizes to control the number of valid points on Matterport3D [2]. Note that these perturbations do not significantly change the sparsity patterns, and a robust model should be able to handle these mild variations. However, from Fig. 2 (a) and (c), we observe that their methods are sensitive to such mild perturbations. It is especially quite counterintuitive that the performance of CSPN degrades with the increase of the input points, showing that it generalizes poorly beyond the sampling density it was trained on.

Besides, as outliers are unavoidable in many applications, we also simulate depth noise for both methods by sampling 0%–10% points from the sparse depth and multiplying the original depth with a random factor from 0.1–2. Fig. 2 (b) and (d) show that their performance also degrades more significantly with the increase in outliers than our method.

Furthermore, we test their robustness to different sparse depth patterns and cross-domains generalization (Tab. 1). For both methods, we input two kinds of sparse depth, *i.e.* uniform sparse depth (Uniform) and sparse depth based on points that are detected by the FAST [29] feature detector (Features). It shows that CSPN, which was trained on the uniform sparse depth, is sensitive to the point distribution (performance degrades a lot from ‘Uniform’ to ‘Features’), while Senushkin *et al.* [33] fails on both sparsity patterns. Besides, we can observe that CSPN is sensitive to the dataset domains. It performs much better on the training dataset (NYU) than other 2 datasets. Results are summarized in Tab. 1. By contrast, our method is more robust to cross domains and different sparsity patterns.

4. Robust Depth Completion

We now introduce our approach, a sparse depth completion method that is designed to be robust to noise, applicable to different types of sparse depth, and to generalize well to unseen datasets.

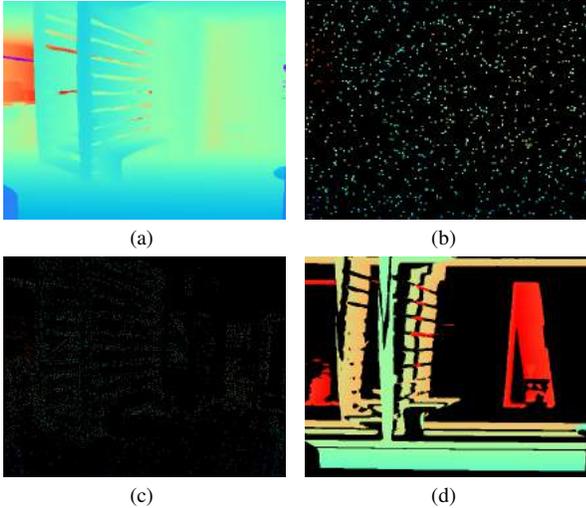


Figure 3. Visualization of sampled sparse depths. We simulate three different patterns from (a) the dense depth to train models: (b) random uniform sampling, (c) feature point based sampling, and (d) region-based sampling.

4.1. Model Architecture and Training

Our depth completion model takes as input the RGB image, sparse depth, as well as a guidance depth map, and it outputs a dense completed depth map. We use the ESANet-R34-NBt1D network with the ResNet-34 backbone proposed by Seichter *et al.* [32] for depth completion, and we use the affine-invariant depth predicted by the method from Yin *et al.* [49] as the guidance depth map. More details of framework is shown in the supplementary material. We employ the virtual normal loss [47], pair-wise normal regression loss [49], ranking edge loss [43], and L1 loss to supervise training. The overall loss is defined as follows.

$$L = L_{vnl} + L_{pwn} + L_{rel} + L_1 \quad (1)$$

We sample 36000 images from Taskonomy [52], DIML [18], and TartainAir [40] as the training data.

4.2. Sparsity Pattern Generation

As we cannot access enough data to cover the diverse sparsity patterns of all possible downstream applications during training, we instead simulate a set of various sparsity patterns. This approach is motivated by domain randomization methods [36,37,51] that train models on simulated data and show that the domain gap to real data can be reduced by randomizing the rendering in the simulator.

We categorize the sparse depth patterns into three main classes, which are illustrated in Fig. 3. During training, we sample sparse points from the dense ground truth depth and try to recover the dense depth map.

- **Uniform.** We sample uniformly distributed points, from hundreds to thousands of points, to simulate the

	Dataset	Sparsity Pattern	Scenes	Images #
GeneralSparsity Benchmark	NYUv2 [35]	Sparse ToF / Unpaired FoV / Short Range	Indoor	654
	ScanNet [8]	Sparse ToF / Unpaired FoV / Short Range	Indoor	700
	DIODE [39]	Sparse ToF / Unpaired FoV / Short Range	Indoor /Outdoor	771
NoisySparsity Benchmark	NYUv2	Holes, Noisy	Indoor	4046

Table 2. Details of proposed two new benchmarks. ‘GeneralSparsity’ aims to evaluate the model’s robustness to different sparsity patterns across domains, while ‘NoisySparsity’ aims to evaluate the robustness to noises.

sparsity pattern from the low-resolution depths, e.g., those captured by ToF sensors on mobile phones.

- **Features.** In order to simulate the sparsity pattern from structure-from-motion and multi-view stereo methods, where high confidence depth values are produced only at regions with distinct/matchable features, we apply the FAST [29] feature detector that samples points on textured regions and particularly image corners.
- **Holes.** Commodity-grade depth sensors cannot capture depth on bright, transparent, reflective and distant surfaces. Therefore, multiple large coherent regions may be missing. We simulate this by 1) masking the depth using a random polygonal region, 2) masking regions at a certain distance, and 3) masking the whole image with the exception of a polygonal region.

To increase the diversity of these patterns, we augment each type of sparse depth by randomizing its parameters (e.g., the number of valid points, mask size, feature thresholds), and then by combining sparsity patterns together.

Improving the Robustness to Outliers. Outliers and depth sensor noise are unavoidable in any depth acquisition method. Most of previous methods only take an RGB image and a sparse depth as the input, and they do not have any extra source of information with which it could distinguish the outliers. However, our method leverages a data prior from the single image depth network, which can help resolve incorrect constraints when there is a significant discrepancy between the two. In order to encourage the network to learn this, we add outliers during training. Specifically, we randomly sample several points and scale their depth by a random factor from 0.1 to 2.

4.3. Benchmark Redesign

Existing benchmarks such as NYU and Matterport3D only consider a specific type of sparsity, and thus they are not suitable for benchmarking a model’s robustness and generalization ability across different task domains. Therefore, we propose two new benchmark designs based on typical types of depth sparsity and depth noise in real-world applications. Critically, we note that the methods we use for generating sparse and noisy data are *not* the same as

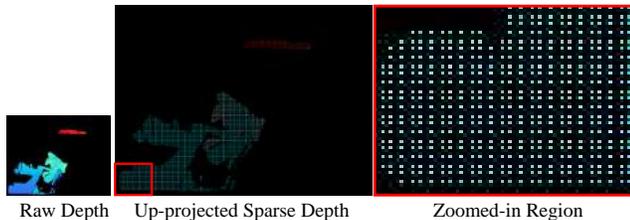


Figure 4. RGBD capture using a Huawei phone and its up-projected depth map. The depth and the RGB sensor have the same field of view but the resolution is different. RGB resolution is in 1280×960 but for depth it is in 240×180 . The raw depth map is up-projected to the RGB resolution, thus yielding the resulting sparse depth.

used for training our approach, this allows us to evaluate the generalization of our training sparsity patterns on these (more realistic) sparsity patterns. Additionally, neither of these datasets were used when training our method.

The first one aims to evaluate the generalization ability of depth completion methods to various sparsity situations, termed *GeneralSparsity*. We collect 3 datasets (NYU [35], ScanNet [8], as well as DIODE [39]) and propose 3 sparse depth patterns to simulate commodity-level RGBD sensors: 1) *Unpaired FOV*. Some RGBD sensors have different field of view of between the RGB camera and the depth sensor, such as KinectV2 (RGB: $84.1^\circ \times 53.8^\circ$; Depth: $70.6^\circ \times 60^\circ$) [42]. We propose to mask a 25% region along the 4 borders of the ground truth depth to simulate this situation. 2) *Sparse ToF*. Time-of-light sensors (ToF) embedded on mobile devices have much lower resolution than that of a RGB camera. For example, the RGB sensor and the depth sensor on a Huawei phone have the same field of view but the resolutions for the RGB camera and the ToF sensor are 1280×960 and 240×180 respectively. We up-project the raw depth to the RGB size and obtain the sparse depth shown in Fig. 4. In our benchmark, we downsample the ground truth depth map to a low resolution, up-project to the original size and mask the distant regions to obtain the sparse depth. 3) *Short Range*. To simulate the short measuring range of depth sensors (such as RealSense), we mask the 50% most distant regions of ground truth to obtain the incomplete depth.

Furthermore, we set up another benchmark to evaluate the robustness to noisy input, termed *NoisySparsity*. We sample 16 videos with over 4000 images from NYU. To simulate the noise pattern in a real scene reconstruction application, COLMAP [30] is used to reconstruct the scene, and we mask the most noisy regions based on the inconsistency between ground truth depth and the COLMAP depth. Note that the masked incomplete depths are still very noisy. Some examples are shown in Fig. 5. More details of these two benchmarks are shown in Tab. 2.

Methods	RMSE↓	AbsRel↓	$\delta_1 \uparrow$	$\delta_2 \uparrow$	$\delta_3 \uparrow$
S2D [22]	0.230	0.044	97.1	99.4	99.8
S2D+SPN [21]	0.172	0.031	98.3	99.7	99.9
DepthCoeff [17]	0.118	0.013	99.4	99.9	-
CSPN [6]	0.117	0.016	99.2	99.9	100.0
CSPN++ [5]	0.116	-	-	-	-
DeepLiDAR [27]	0.115	0.022	99.3	99.9	100.0
DepthNormal [44]	0.112	0.018	99.5	99.9	100.0
NLSP [26]	0.092	0.012	99.6	99.9	100.0
Lee <i>et al.</i> [19]	0.104	0.014	99.4	99.9	100.0
Huynh <i>et al.</i> [15]	0.090	0.014	99.6	99.9	100.0
MiDaS [28]	0.513	0.110	88.6	98.1	99.6
Yin <i>et al.</i> [49]	0.402	0.090	91.3	98.0	99.5
Ours-baseline	0.210	0.036	98.4	99.6	99.9
Ours-W MiDaS [28]	0.199	0.024	98.6	99.6	99.9
Ours-W Yin <i>et al.</i> [49]	0.183	0.022	98.7	99.7	99.9

Table 3. Depth completion results on the NYU dataset. Following [6, 26], we uniformly sample 500 points from ground truth as the sparse depth. Our method is *not trained on NYU* but is comparable with state-of-the-art methods show here that *are trained on NYU*. MiDaS and Yin *et al.* are monocular depth estimation methods that are used as input to our method.

5. Experiments

In this section, we conduct several experiments to demonstrate the effectiveness of our approach.

Implementation Details. We employ the network proposed by Seichter *et al.* [32] with a pretrained ResNet-34 [12] backbone in our experiments. We use SGD for optimization with an initial learning rate of 0.02 for all layers. The learning rate is decayed every 40000 iterations with the ratio 0.1, and we used 24 samples per batch. We sampled 12000 images from three datasets respectively for training, i.e. Taskonomy [52], DIML [18], and TartanAir [40]. During training, images are randomly flipped horizontally, and resized to 448×448 .

Evaluation Metrics. Following previous methods [5, 26], we employ multiple metrics to evaluate the results, including absolute relative error (AbsRel), mean absolute error (MAE), root mean squared error (RMSE), and percentage of pixels satisfying $\delta_\tau = \max\left(\frac{d_{pred}}{d_{gt}}, \frac{d_{gt}}{d_{pred}}\right) < \tau$.

5.1. Depth Completion Evaluation

In this experiment, we first compare our method with previous methods on the old NYU and Matterport3D benchmarks to show that our method can achieve performance comparable to previous methods, although it has not been trained on either of these datasets. Second, we compare on our new sparse depth completion benchmarks. Finally, we test our method on the DualPixel dataset [10] captured by a smart phone to further show the generalization of our method to mobile sensors.

Existing Benchmarks. Here, we compare to current state-of-the-art methods on NYU [35] and Matterport3D [2]. Although these two benchmarks have different sparse depth

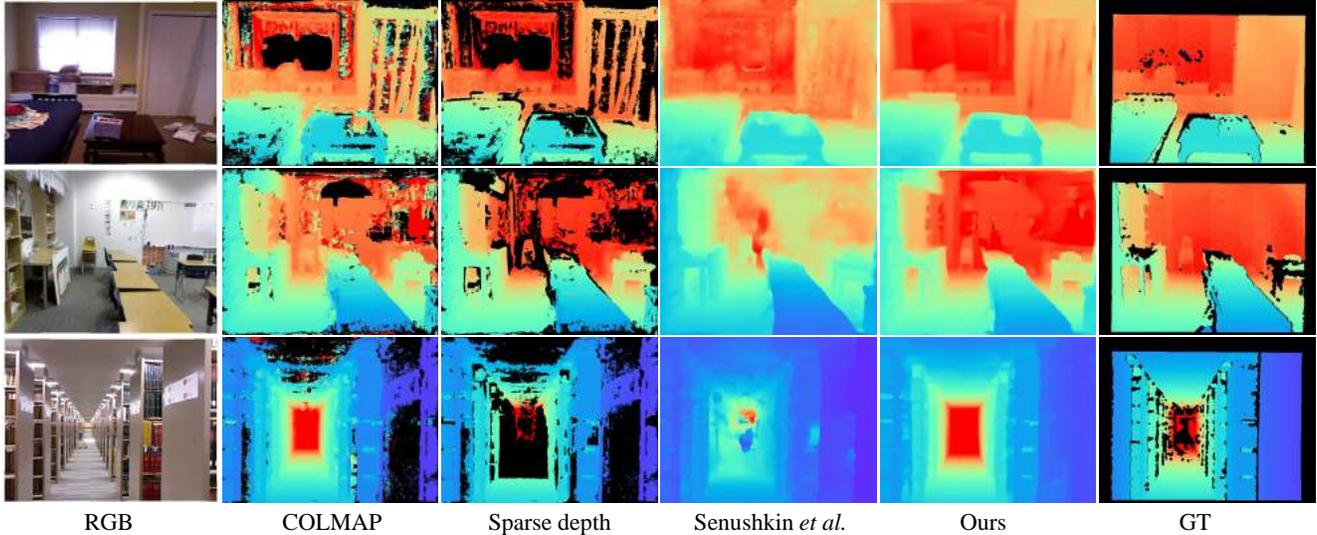


Figure 5. Qualitative comparison for completing noisy sparse depth. The noisy sparse depths are obtained by masking COLMAP [30] depths. Our completed results have less outliers and errors.

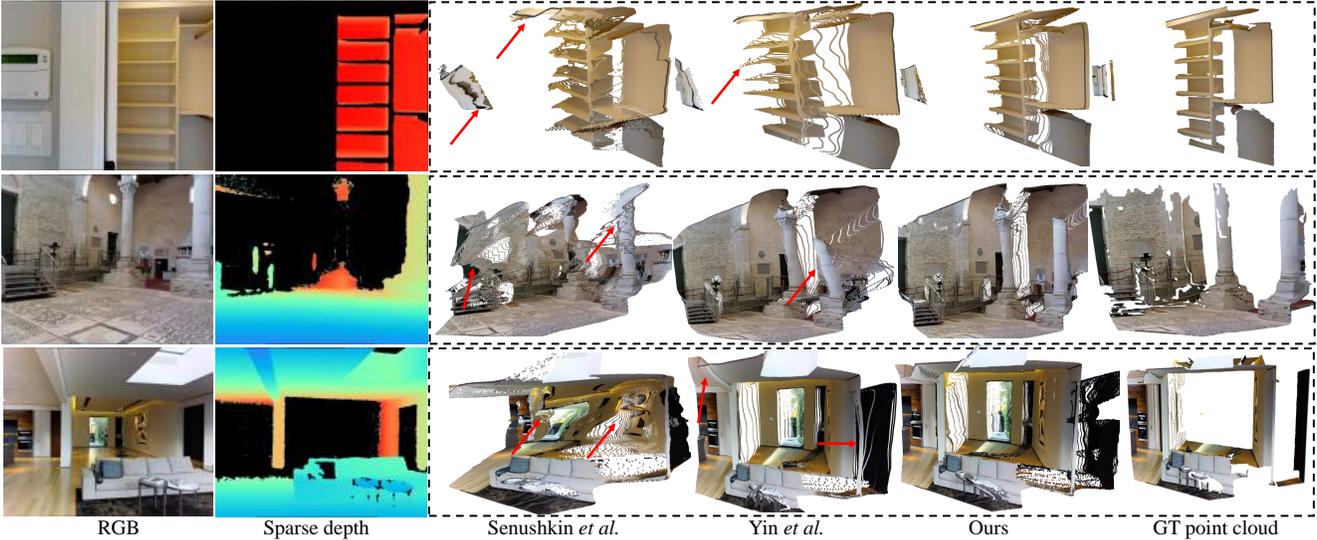


Figure 6. Qualitative comparison of depth and reconstructed 3D shape. Our completed metric depth has finer details and the reconstructed 3D shape is more accurate than Senushkin *et al.* [33] and Yin *et al.* [49]. Note that Yin *et al.* cannot predict metric depth.

types, we use a single model for evaluation on both of them. We include a baseline method (Ours-baseline) where the model does not make use of depth guidance, that is, it directly predicts the complete depth from RGB and sparse depth. Results are shown in Tab. 3 and Tab. 4, which indicate that our method achieves performance comparable to current methods, despite not being trained on these datasets.

Following [6, 26], we uniformly sample 500 points from the ground truth as the sparse depth on NYU. We can see that our method achieves performance on par with previous methods, and better accuracy than the baseline. Comparing with our baseline (Ours-baseline), using a pretrained guidance model further improves performance. When compared

Methods	RMSE \downarrow	MAE \downarrow	$\delta_{1.05}\uparrow$	$\delta_{1.1}\uparrow$	$\delta_1\uparrow$	$\delta_2\uparrow$	$\delta_3\uparrow$
Huang <i>et al.</i> [14]	1.092	0.342	66.1	75.0	85.0	91.1	93.6
Zhang <i>et al.</i> [54]	1.316	0.461	65.7	70.8	78.1	85.1	88.8
Senushkin <i>et al.</i> [33]	1.028	0.299	71.9	80.5	89.0	93.2	95.0
Yin <i>et al.</i> [49]	2.06	1.13	17.9	29.8	50.7	72.3	83.4
MiDaS [28]	3.45	2.01	13.2	21.8	37.5	54.8	66.4
Ours-baseline	2.35	0.574	68.9	78.6	86.1	91.5	96.0
Ours-W [28]	1.49	0.448	67.8	76.3	85.0	91.0	94.5
Ours-W [49]	1.03	0.320	71.2	79.0	87.1	93.1	96.0

Table 4. Depth completion results on the Matterport3D dataset. Our method is *not trained on Matterport3D* but is comparable with state-of-the-art methods that are trained on Matterport3D. RMSE and MAE are given in meters. MiDaS and Yin *et al.* are monocular depth estimation methods that are used as input to our method.

Methods	NYU			ScanNet			DIODE		
	AbsRel↓			AbsRel↓			AbsRel↓		
	Unp. FOV	Sparse ToF	Short Range	Unp. FOV	Sparse ToF	Short Range	Unp. FOV	Sparse ToF	Short Range
NLSP [26]	0.150	0.190	0.114	0.716	1.413	0.202	6.684	11.370	1.005
Senushkin <i>et al.</i> [33]	0.224	0.615	0.093	0.255	0.793	0.166	0.687	6.120	0.623
Ours-baseline	0.046	0.018	0.041	0.049	0.022	0.047	0.150	0.143	0.144
Ours-W Guidance	0.031	0.013	0.030	0.028	0.014	0.037	0.139	0.111	0.137

Table 5. Comparison of state-of-the-art methods on our proposed GeneralSparsity benchmark. This benchmark mainly analyzes the robustness of depth completion methods to different sparsity types. Although these testing datasets and sparsity types are unseen to our methods during training, we can achieve better performance than current methods.

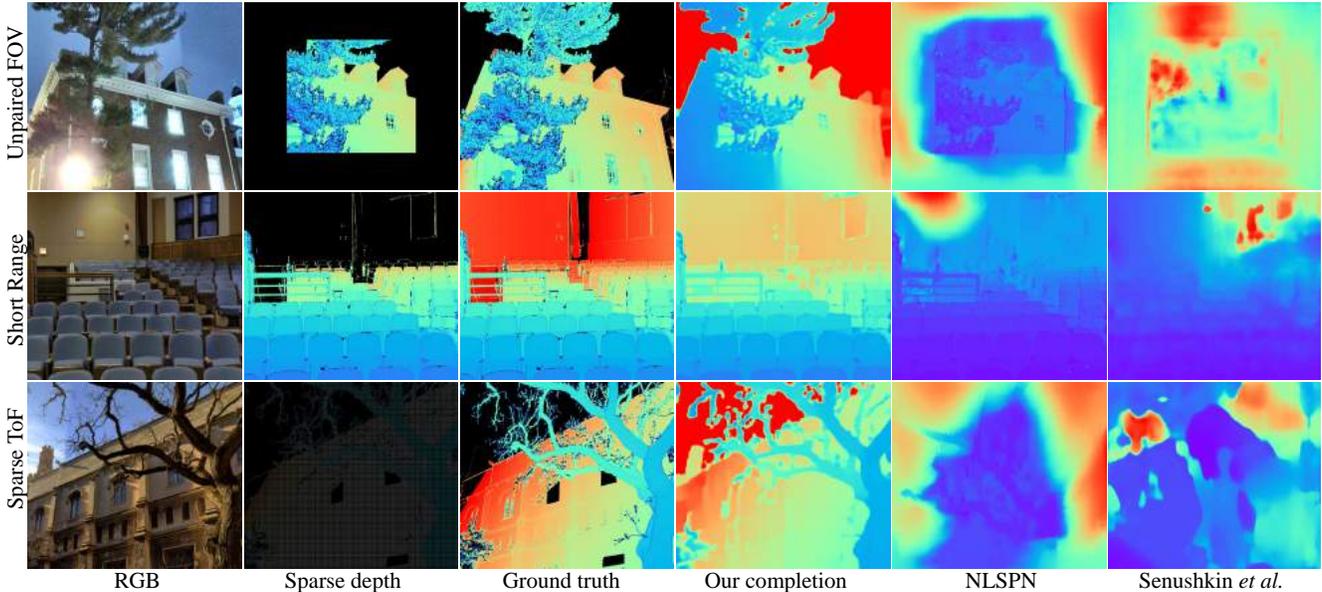


Figure 7. Qualitative completion results on the DIODE [39] dataset. Note that none of the methods are trained on this dataset. We compare our method with Senushkin *et al.* [33] and NLSP [26] using 3 different unseen sparsity patterns.

to Yin *et al.* [49] (which provides our guidance input), we see that even with only 500 sparse points, the performance is improved significantly. Furthermore, we observe that our method is robust with respect to different guides, see ‘Ours-W MiDaS’ (using the MiDaS depth as the input during test) and ‘Ours-W Yin *et al.* [49]’. As depths from Yin *et al.* [49] are employed for the guidance map during training, taking their depth can achieve better results in test.

Moreover, the qualitative comparison on Matterport3D is illustrated in Fig. 6. Although Senushkin *et al.* [33] can achieve better accuracy than ours (note that it was trained on Matterport3D) we find that our reconstructed scene structure is better.

Proposed New Benchmarks. We now compare to current state-of-the-art methods on the proposed GeneralSparsity and NoisySparsity datasets.

We compare our approach to NLSPN [26] and Senushkin *et al.* [33] since they have achieved the most promising results on current benchmarks. Note that the NLSPN method is trained on NYU, while Senushkin *et al.* [33] is trained on Matterport3D. In contrast, our method is not trained on either of these datasets. Furthermore, the spe-

cific sparsity patterns used for evaluation are not utilized in our training. Results are summarized in Tab. 5. We find that NLSP [26] and Senushkin *et al.* [33] do not generalize well to different types of sparse depth, or datasets, which is reasonable as they are trained on one dataset and one sparsity type only. In contrast, our method can achieve comparable performance on both datasets. Furthermore, comparing to our baseline method, we see that using the guidance map can consistently improve the performance over all datasets and sparse depth types. The quantitative comparison of completing 3 different patterns on the DIODE dataset is demonstrated in Fig. 7.

On the NoisySparsity benchmark, we evaluate robustness to noisy sparse inputs. In Tab. 6, we observe that our method can achieve better accuracy than existing methods. Comparing with our baseline approach, using the guidance map further boosts performance. Qualitative comparisons are shown in Fig. 5, we can see that our completed depths have much less outliers and noise (see the wall). More comparisons can be found in the supplementary. m

Completing Mobile Phone Sensor Depth. Many smartphones have been equipped with cheap 3D sensors. To eval-

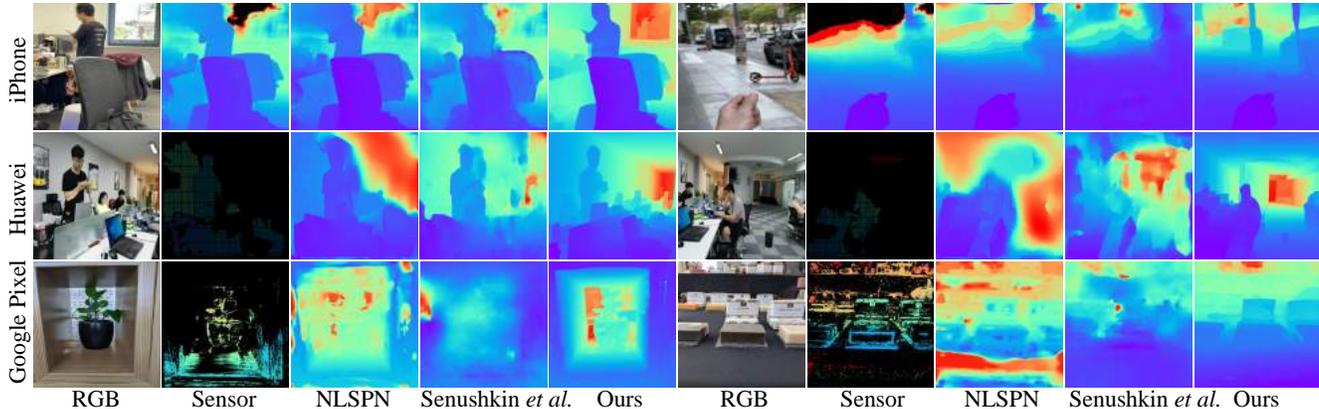


Figure 8. Completion of the phone-captured depths. Our method is more robust to different depth sensors than previous methods.

Metrics	NLSPN [26]	Senushkin <i>et al.</i> [33]	Ours (baseline)	Ours (W Guidance)
AbsRel (%)↓	24.9	18.26	6.41	5.25
δ_1 (%)↑	54.7	70.3	93.6	95.2

Table 6. Comparison on the proposed NoisySparsity benchmark, showing robustness to noisy input.

Methods	AbsRel ↓		
	NYU (Sparse ToF)	ScanNet (Unpaired FOV)	Matterport3D (Original pattern)
W/o Features	0.036	0.047	0.122
W/o Uniform	0.021	0.044	0.121
W/o Holes	0.012	0.12	0.165
Ours-Full	0.013	0.028	0.120

Table 7. Ablation of different training sparsity patterns. The model is trained on the simulated patterns except the one specified by ‘W/o’ and evaluated on zero-shot datasets.

uate the robustness of our completion method with respect to noisy phone-captured depths, we perform a quantitative comparison on a Google Pixel 3 captured RGBD dataset (DualPixel) [10]. We apply the provided confidence map to sample the most confident 36% regions as the sparse depths and evaluate the depth on the most confident 80% regions. On this dataset, our method achieves **1.6%** AbsRel (↓), while NLSPN [26] and Shenshkin [33] only achieve 4.1% and 24.86% respectively.

Furthermore, we also qualitatively test on 2 more phones, i.e. iPhone 12, and Huawei Mate 30. In each case, the depths are acquired differently. iPhone uses a stereo matching method to obtain the depth, which is normalized to 0 – 255 and saved as the inverse depth. Huawei provides a low-resolution dense depth, i.e. 180×240 pixels, which is captured by a ToF sensor. Google Pixel uses the dual-pixel sensor, which has a very small baseline, to capture the depth. The comparison is illustrated in Fig. 8. We can see that our method generalizes well to different real world sparse data acquisition devices with a single trained model. More examples are in the supplementary material.

Ablation of Synthetic Sparsity Patterns. This study aims to investigate the effectiveness of different simulated sparsity patterns during training. We remove one of the proposed patterns during training and evaluate them on 3 zero-shot datasets with different patterns. On NYU and ScanNet, we simulate the pattern ‘Sparse ToF’ and ‘Unpaired FOV’, while we use the provided sparse depth on Matterport3D. All models have been trained with the same number of epochs, and identical parameters other than the sparsity patterns. Results are summarized in Tab. 7. We observe that when missing the simulated sparse depth pattern, the performance on the most-related testing data will decrease. For example, the model trained without the ‘Holes’ pattern has worse performance than others on ScanNet and Matterport3D. Therefore, we see how combining different sparsity patterns improves cross-domain generalization.

6. Discussion

Limitations. Our method takes as input a depth map from the monocular depth estimation model as the guidance. While such single image depth methods are improving, it is possible for the guidance map to have significant errors still. This leads to adverse effects on the depth completion. Furthermore, although our method is robust to outliers and noise, we observe that the completion quality will decrease significantly in very sparse scenarios, e.g., where over 50% of sparse depths are outliers.

Conclusion. In this paper, we analyze exiting depth completion methods and show that they cannot generalize to different data domains and are sensitive to noise. We propose a simple yet robust system for depth completion. Our method leverages a single image depth prior and is trained with diverse data augmentation. In addition, we re-design two depth completion benchmarks to better evaluate a depth completion method’s generalization ability and its robustness to noise, different sparsity patterns, and diverse scenes. We show that our approach achieves promising performance on our new benchmarks, works with a variety of

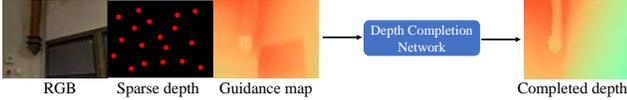


Figure 9. Our method takes an RGB image, sparse depth, and guidance map as input, and it outputs dense depth. We can iterate the network several times, replacing the guidance map with the output of the previous iteration.

sensor types and is robust to the sensor noise, with a single trained model. Our method provides a solution to high quality mobile depth capture, which can improve downstream depth related applications on mobile devices.

7. Appendix

7.1. Details of Framework

Our depth completion model takes as input the RGB image, sparse depth, and a guidance depth map. It outputs a completed dense depth map. In our experiments, we use the ESANet-R34-NBt1D network with the ResNet-34 backbone proposed by Seichter *et al.* [32] for depth completion, and we use the affine-invariant depth predicted by the method from Yin *et al.* [49] as the guidance depth map. The framework is illustrated 9. During training, we sample 36000 images from Taskonomy [52], DIML [18], and TartainAir [40] as the training data.

7.2. More Examples on Our Proposed GeneralSparsity Benchmark

In this section, we demonstrate more qualitative comparisons on our proposed GeneralSparsity benchmark. There are three different sparsity patterns: 1) Unpaired FOV. We mask 25% region around 4 borders of the ground truth depth, see Fig. 10. 2) Sparse ToF. We simulate Huawei ToF sensor to sample the depth from the ground truth, see Fig. 12. 3) Short Range. We mask the most 50% distant regions, see Fig. 11. These three types never existed in our training. We compare with NLSPN [26] and Senushkin *et al.* [33]. We can observe that both NLSPN and Senushkin *et al.* [33] methods cannot handle other unseen sparse depth types. Although ‘Short Range’ type is similar to the training type of Senushkin *et al.* method, they may also fail. We conjunct that their methods lack the generalization to diverse scenes.

7.3. More Examples on Our Proposed NoisySparsity Benchmark

In Fig. 15 and Fig. 16, we compare with Senushkin *et al.* [33] methods. Our method is more robust to noise.

7.4. Complete Phone Sensor Captured Depth

Many smart phones have been equipped with the depth sensors. To evaluate the generalization ability of our methods, we test on Pixel (see Fig. 13), iPhone (see the first 3 rows of Fig. 14), and Huawei (see the last 2 rows of Fig. 14) captured depth. We compare with NLSPN [26] and Senushkin *et al.* [33] methods. Yin *et al.* [49] is input as the guidance map. For the Pixel sensor, we employ the provided confidence map to remove the most noisy regions. We can see that most regions of such depths have valid value. NLSPN can predict good results on most of examples except the third row. We conjunct this is the reason that the third row example is very noisy and lost most of depth information. By contrast, Senushkin *et al.* method fail on most cases. Comparing with the guidance map, our completed results can recover finer details. On the iPhone depth, NLSPN can also perform much better than that of Senushkin *et al.* method.

References

- [1] Jia-Wang Bian, Huangying Zhan, Naiyan Wang, Zhichao Li, Le Zhang, Chunhua Shen, Ming-Ming Cheng, and Ian Reid. Unsupervised scale-consistent depth learning from video. *IJCV*, 2021. 2
- [2] Angel Chang, Angela Dai, Thomas Funkhouser, Maciej Halber, Matthias Niessner, Manolis Savva, Shuran Song, Andy Zeng, and Yinda Zhang. Matterport3d: Learning from rgb-d data in indoor environments. *Int. Conf. 3D. Vis.*, 2017. 2, 3, 5
- [3] Weifeng Chen, Shengyi Qian, and Jia Deng. Learning single-image depth from videos using quality assessment networks. In *CVPR*, pages 5604–5613, 2019. 2
- [4] Yun Chen, Bin Yang, Ming Liang, and Raquel Urtasun. Learning joint 2d-3d representations for depth completion. In *ICCV*, pages 10023–10032, 2019. 2
- [5] Xinjing Cheng, Peng Wang, Chenye Guan, and Ruigang Yang. Cspn++: Learning context and resource aware convolutional spatial propagation networks for depth completion. In *AAAI*, volume 34, pages 10615–10622, 2020. 1, 2, 3, 5
- [6] Xinjing Cheng, Peng Wang, and Ruigang Yang. Depth estimation via affinity learned with convolutional spatial propagation network. In *ECCV*, pages 103–119, 2018. 1, 3, 5, 6
- [7] Xinjing Cheng, Peng Wang, and Ruigang Yang. Learning depth with convolutional spatial propagation network. *IEEE Trans. Pattern Anal. Mach. Intell.*, 2019. 1, 2, 3
- [8] Angela Dai, Angel X Chang, Manolis Savva, Maciej Halber, Thomas Funkhouser, and Matthias Nießner. Scannet: Richly-annotated 3d reconstructions of indoor scenes. In *CVPR*, pages 5828–5839, 2017. 2, 3, 4, 5
- [9] David Eigen, Christian Puhrsch, and Rob Fergus. Depth map prediction from a single image using a multi-scale deep network. In *NeurIPS*, pages 2366–2374, 2014. 2

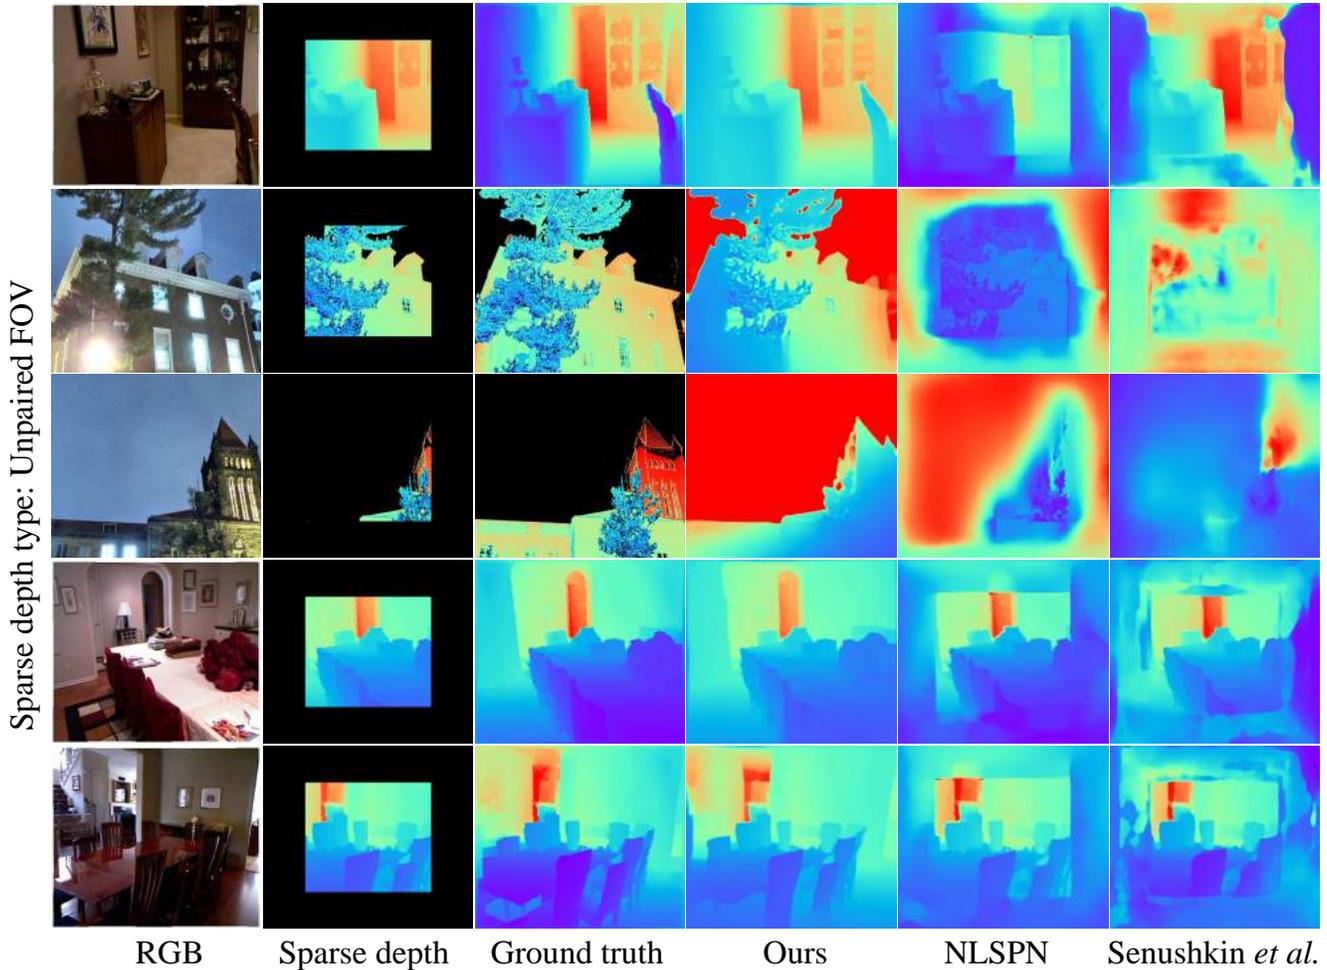


Figure 10. The comparison on completing ‘Unpaired FOV’ sparse depth. We create this type by masking 25% region around 4 borders of the ground truth depth for evaluation. We can see that our method is more robust on this type.

- [10] Rahul Garg, Neal Wadhwa, Sameer Ansari, and Jonathan T. Barron. Learning single camera depth estimation using dual-pixels. *ICCV*, 2019. 5, 8
- [11] Clément Godard, Oisín Mac Aodha, Michael Firman, and Gabriel J. Brostow. Digging into self-supervised monocular depth prediction. In *ICCV*, 2019. 2
- [12] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *CVPR*, pages 770–778, 2016. 5
- [13] Daniel Herrera, Juho Kannala, Janne Heikkilä, et al. Depth map inpainting under a second-order smoothness prior. In *Scandinavian Conference on Image Analysis*, pages 555–566. Springer, 2013. 2
- [14] Yu-Kai Huang, Tsung-Han Wu, Yueh-Cheng Liu, and Winston H Hsu. Indoor depth completion with boundary consistency and self-attention. In *Int. Conf. Comput. Vis. Worksh.*, pages 0–0, 2019. 1, 6
- [15] Lam Huynh, Phong Nguyen, Jiri Matas, Esa Rahtu, and Janne Heikkilä. Boosting monocular depth estimation with lightweight 3d point fusion. In *ICCV*, pages 12767–12776, 2021. 2, 5
- [16] Saif Imran, Xiaoming Liu, and Daniel Morris. Depth completion with twin surface extrapolation at occlusion boundaries. *CVPR*, 2021. 1, 2
- [17] Saif Imran, Yunfei Long, Xiaoming Liu, and Daniel Morris. Depth coefficients for depth completion. In *CVPR*, pages 12438–12447. IEEE, 2019. 5
- [18] Youngjung Kim, Hyungjoo Jung, Dongbo Min, and Kwanghoon Sohn. Deep monocular depth estimation via integration of global and local predictions. *IEEE TIP*, 27(8):4131–4144, 2018. 4, 5, 9
- [19] Byeong-Uk Lee, Kyunghyun Lee, and In So Kweon. Depth completion using plane-residual representation. In *CVPR*, pages 13916–13925, 2021. 5
- [20] Fayao Liu, Chunhua Shen, Guosheng Lin, and Ian Reid. Learning depth from single monocular images using deep convolutional neural fields. *IEEE TPAMI*, 38(10):2024–2039, 2015. 2
- [21] Sifei Liu, Shalini De Mello, Jinwei Gu, Guangyu Zhong, Ming-Hsuan Yang, and Jan Kautz. Learning affinity via spatial propagation networks. *arXiv: Comp. Res. Repository*, page 1710.01020, 2017. 5

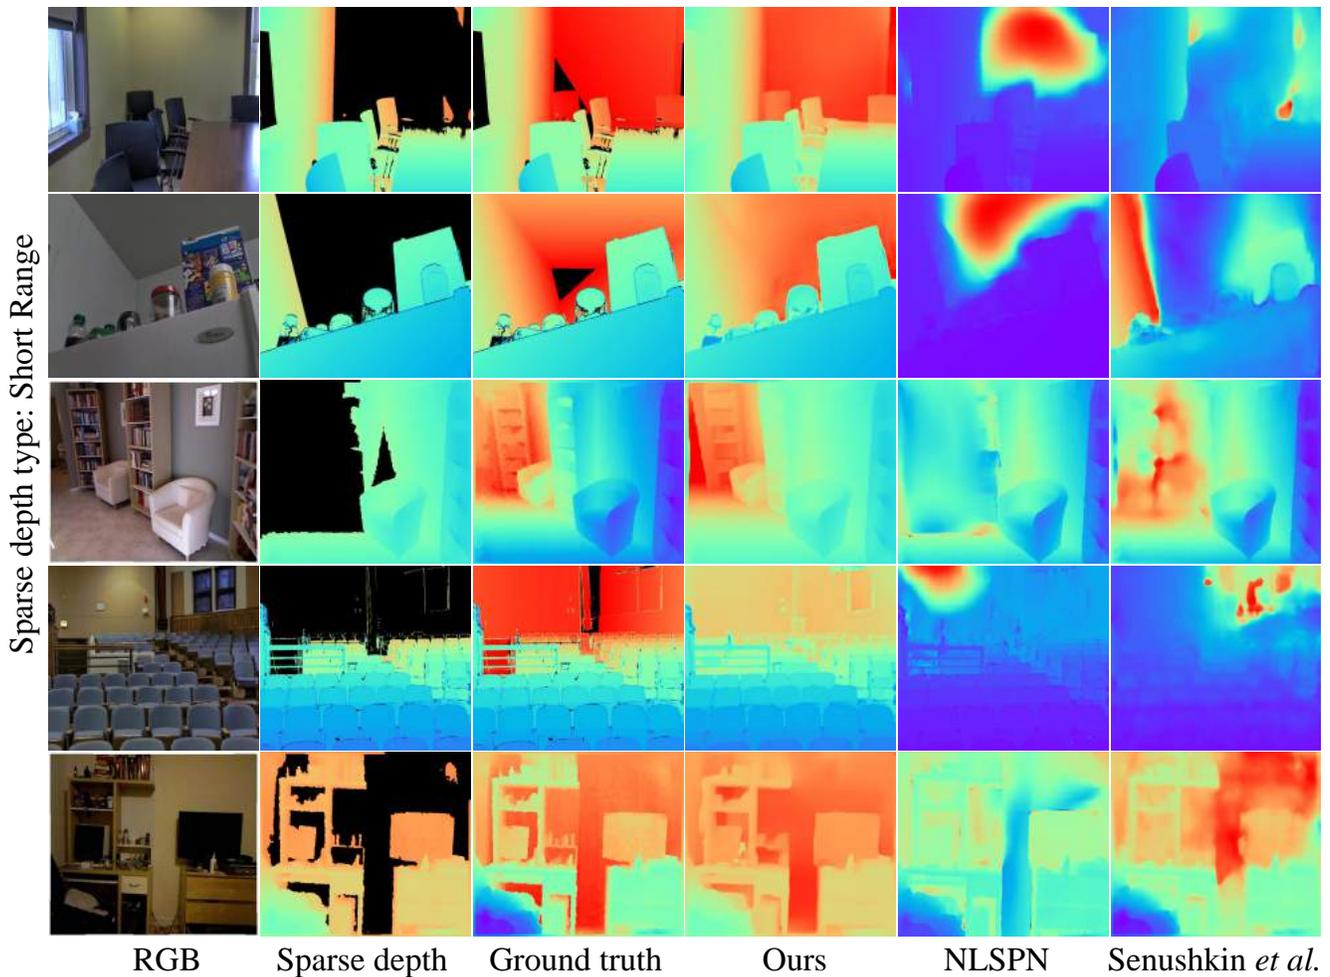


Figure 11. The comparison on completing ‘Short Range’ sparse depth. We create this type by masking the most 50% distant regions. We can see that our method is more robust on this type.

- [22] Fangchang Ma and Sertac Karaman. Sparse-to-dense: Depth prediction from sparse depth samples and a single image. In *IEEE Int. Conf. on Robot. and Auto.*, pages 4796–4803. IEEE, 2018. 5
- [23] Kiyoshi Matsuo and Yoshimitsu Aoki. Depth image enhancement using local tangent plane approximations. In *CVPR*, pages 3574–3583, 2015. 2
- [24] Raúl Mur-Artal and Juan D. Tardós. ORB-SLAM2: an open-source SLAM system for monocular, stereo and RGB-D cameras. *IEEE Trans. Robot.*, 33(5):1255–1262, 2017. 1
- [25] Richard A Newcombe, Shahram Izadi, Otmar Hilliges, David Molyneaux, David Kim, Andrew J Davison, Pushmeet Kohi, Jamie Shotton, Steve Hodges, and Andrew Fitzgibbon. Kinectfusion: Real-time dense surface mapping and tracking. In *IEEE International Symposium on Mixed and Augmented Reality*, pages 127–136. IEEE, 2011. 1
- [26] Jinsun Park, Kyungdon Joo, Zhe Hu, Chi-Kuei Liu, and In-So Kweon. Non-local spatial propagation network for depth completion. In *ECCV. European Conference on Computer Vision*, 2020. 1, 2, 5, 6, 7, 8, 9
- [27] Jiaxiong Qiu, Zhaopeng Cui, Yinda Zhang, Xingdi Zhang, Shuaicheng Liu, Bing Zeng, and Marc Pollefeys. DeepLi-
- dar: Deep surface normal guided depth prediction for outdoor scene from sparse lidar data and single color image. In *CVPR*, pages 3313–3322, 2019. 1, 5
- [28] René Ranftl, Katrin Lasinger, David Hafner, Konrad Schindler, and Vladlen Koltun. Towards robust monocular depth estimation: Mixing datasets for zero-shot cross-dataset transfer. *IEEE TPAMI*, 2020. 2, 3, 5, 6
- [29] Edward Rosten and Tom Drummond. Machine learning for high-speed corner detection. In *ECCV*, pages 430–443. Springer, 2006. 3, 4
- [30] Johannes Lutz Schönberger, Enliang Zheng, Marc Pollefeys, and Jan-Michael Frahm. Pixelwise view selection for unstructured multi-view stereo. In *ECCV*, 2016. 2, 5, 6, 15, 16
- [31] Thomas Schops, Johannes L Schönberger, Silvano Galliani, Torsten Sattler, Konrad Schindler, Marc Pollefeys, and Andreas Geiger. A multi-view stereo benchmark with high-resolution images and multi-camera videos. In *CVPR*, pages 3260–3269, 2017. 1
- [32] Daniel Seichter, Mona Köhler, Benjamin Lewandowski, Tim Wengelfeld, and Horst-Michael Gross. Efficient rgb-d seman-

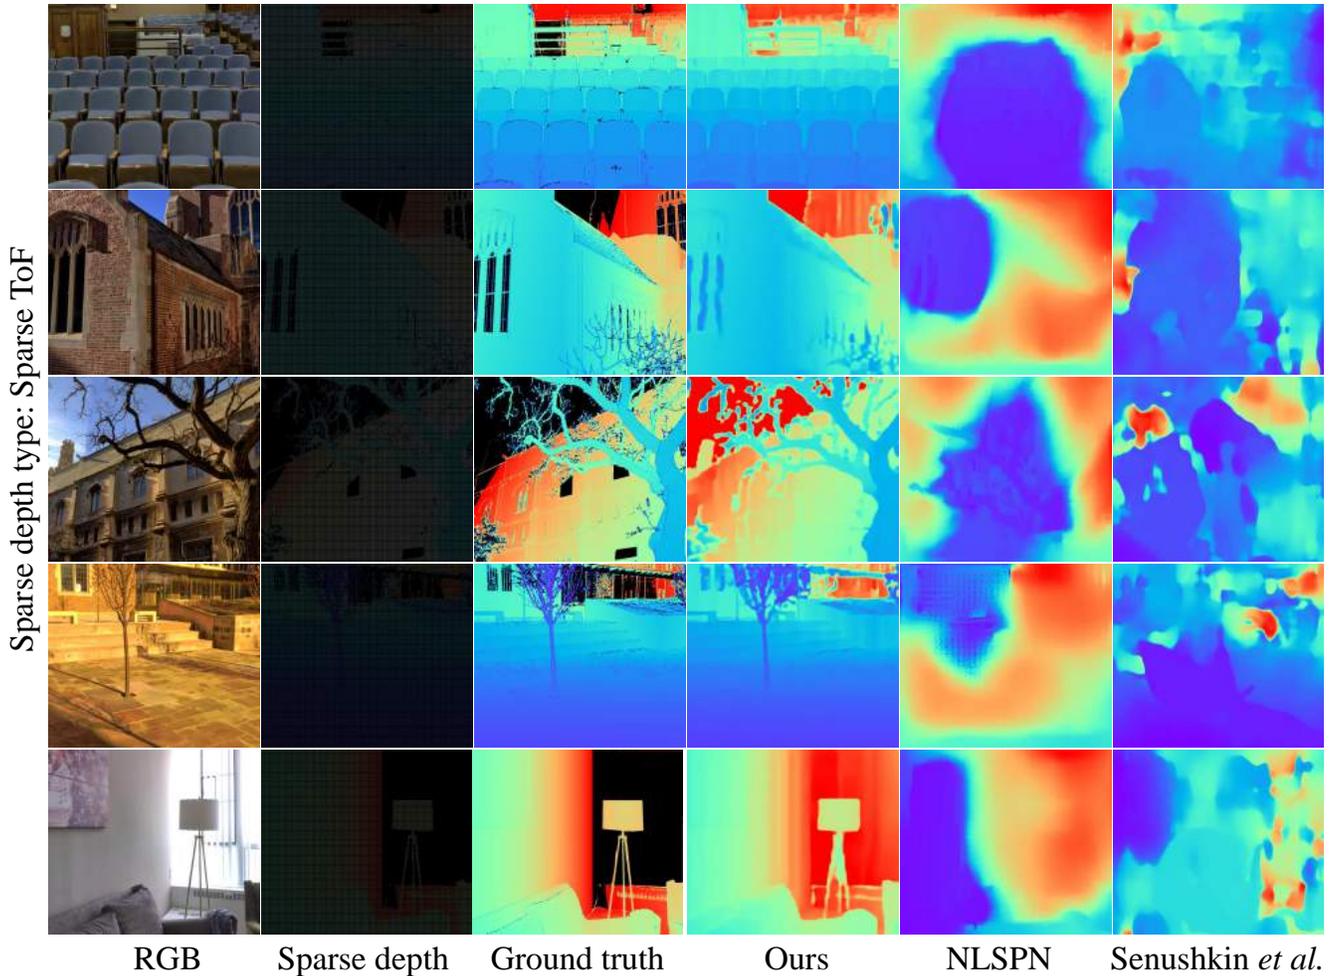


Figure 12. The comparison on completing ‘Sparse ToF’ sparse depth. We create this type by sampling the depth from the ground truth every 3 pixels horizontally and vertically. We can see that our method can achieve better results.

- tic segmentation for indoor scene analysis. *arXiv: Comp. Res. Repository*, 2020. 4, 5, 9
- [33] Dmitry Senushkin, Iliia Belikov, and Anton Konushin. Decoder modulation for indoor depth completion. *arXiv: Comp. Res. Repository*, page 2005.08607, 2020. 1, 3, 6, 7, 8, 9
- [34] Shaoshuai Shi, Chaoxu Guo, Li Jiang, Zhe Wang, Jianping Shi, Xiaogang Wang, and Hongsheng Li. Pv-rcnn: Point-voxel feature set abstraction for 3d object detection. In *CVPR*, pages 10529–10538, 2020. 1
- [35] Nathan Silberman, Derek Hoiem, Pushmeet Kohli, and Rob Fergus. Indoor segmentation and support inference from rgb-d images. In *ECCV*, pages 746–760. Springer, 2012. 1, 2, 3, 4, 5
- [36] Josh Tobin, Lukas Biewald, Rocky Duan, Marcin Andrychowicz, Ankur Handa, Vikash Kumar, Bob McGrew, Alex Ray, Jonas Schneider, Peter Welinder, et al. Domain randomization and generative models for robotic grasping. In *Proc. IEEE/RSJ Int. Conf. Intelli. Robot. & Sys.*, pages 3482–3489. IEEE, 2018. 2, 4
- [37] Josh Tobin, Rachel Fong, Alex Ray, Jonas Schneider, Wojciech Zaremba, and Pieter Abbeel. Domain randomization for transferring deep neural networks from simulation to the real world. In *Proc. IEEE/RSJ Int. Conf. Intelli. Robot. & Sys.*, pages 23–30. IEEE, 2017. 2, 4
- [38] Jonas Uhrig, Nick Schneider, Lukas Schneider, Uwe Franke, Thomas Brox, and Andreas Geiger. Sparsity invariant cnns. In *International Conference on 3D Vision (3DV)*, 2017. 1
- [39] Igor Vasiljevic, Nick Kolkin, Shanyi Zhang, Ruotian Luo, Haochen Wang, Falcon Z Dai, Andrea F Daniele, Mohamadreza Mostajabi, Steven Basart, Matthew R Walter, et al. Diode: A dense indoor and outdoor depth dataset. *arXiv: Comp. Res. Repository*, page 1908.00463, 2019. 2, 4, 5, 7
- [40] Wenshan Wang, DeLong Zhu, Xiangwei Wang, Yaoyu Hu, Yuheng Qiu, Chen Wang, Yafei Hu, Ashish Kapoor, and Sebastian Scherer. Tartanair: A dataset to push the limits of visual slam. *arXiv: Comp. Res. Repository*, 2020. 4, 5, 9
- [41] Yan Wang, Wei-Lun Chao, Divyansh Garg, Bharath Hariharan, Mark Campbell, and Kilian Q Weinberger. Pseudo-lidar from visual depth estimation: Bridging the gap in 3d object detection for autonomous driving. In *CVPR*, pages 8445–8453, 2019. 1

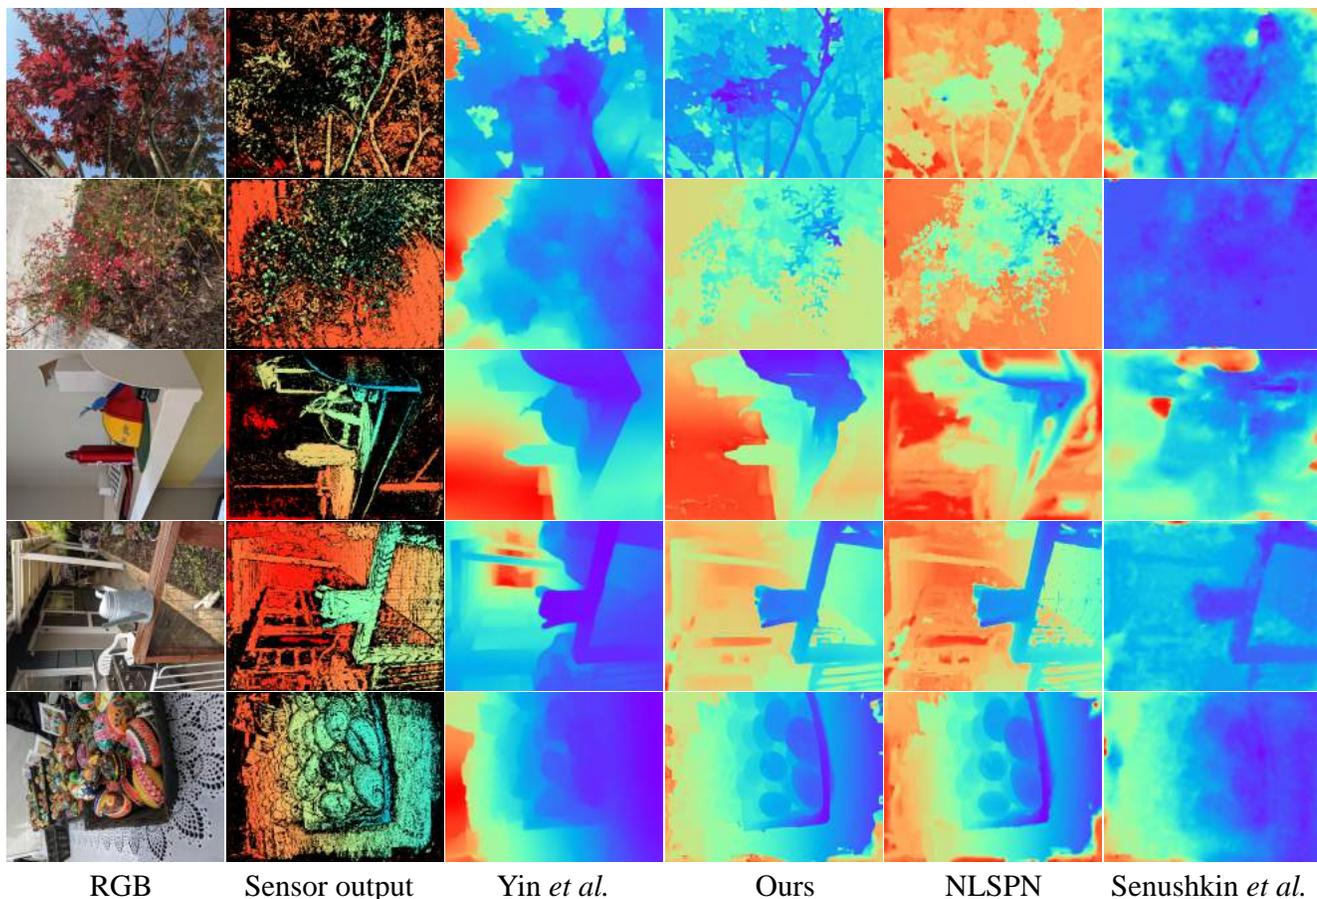


Figure 13. The comparison of completing Pixel phone captured depths. We can see that our method can recover better results than other methods.

- [42] Oliver Wasenmüller and Didier Stricker. Comparison of kinect v1 and v2 depth images in terms of accuracy and precision. In *ACCV*, pages 34–45. Springer, 2016. [5](#)
- [43] Ke Xian, Jianming Zhang, Oliver Wang, Long Mai, Zhe Lin, and Zhiguo Cao. Structure-guided ranking loss for single image depth prediction. In *CVPR*, pages 611–620, 2020. [2](#), [4](#)
- [44] 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. [1](#), [5](#)
- [45] Yanchao Yang, Alex Wong, and Stefano Soatto. Dense depth posterior (ddp) from single image and sparse range. In *CVPR*, pages 3353–3362, 2019. [2](#)
- [46] Yao Yao, Zixin Luo, Shiwei Li, Jingyang Zhang, Yufan Ren, Lei Zhou, Tian Fang, and Long Quan. Blendedmvs: A large-scale dataset for generalized multi-view stereo networks. *CVPR*, 2020. [1](#)
- [47] Wei Yin, Yifan Liu, Chunhua Shen, and Youliang Yan. Enforcing geometric constraints of virtual normal for depth prediction. In *ICCV*, 2019. [2](#), [4](#)
- [48] Wei Yin, Xinlong Wang, Chunhua Shen, Yifan Liu, Zhi Tian, Songcen Xu, Changming Sun, and Dou Renyin. Diversedepth: Affine-invariant depth prediction using diverse data. *arXiv: Comp. Res. Repository*, page 2002.00569, 2020. [2](#)
- [49] Wei Yin, Jianming Zhang, Oliver Wang, Simon Niklaus, Long Mai, Simon Chen, and Chunhua Shen. Learning to recover 3d scene shape from a single image. In *CVPR*, 2021. [2](#), [4](#), [5](#), [6](#), [7](#), [9](#)
- [50] Wei Yin, Jianming Zhang, Oliver Wang, Simon Niklaus, Long Mai, Simon Chen, and Chunhua Shen. Learning to recover 3d scene shape from a single image. *CVPR*, 2021. [2](#)
- [51] Sergey Zakharov, Wadim Kehl, and Slobodan Ilic. Deceptionnet: Network-driven domain randomization. In *ICCV*, pages 532–541, 2019. [2](#), [4](#)
- [52] Amir Zamir, Alexander Sax, William Shen, Leonidas Guibas, Jitendra Malik, and Silvio Savarese. Taskonomy: Disentangling task transfer learning. In *CVPR*. IEEE, 2018. [4](#), [5](#), [9](#)
- [53] Feihu Zhang, Victor Prisacariu, Ruigang Yang, and Philip Torr. Ga-net: Guided aggregation net for end-to-end stereo matching. In *CVPR*, pages 185–194, 2019. [1](#)
- [54] Yinda Zhang and Thomas Funkhouser. Deep depth completion of a single rgb-d image. *CVPR*, 2018. [1](#), [2](#), [6](#)

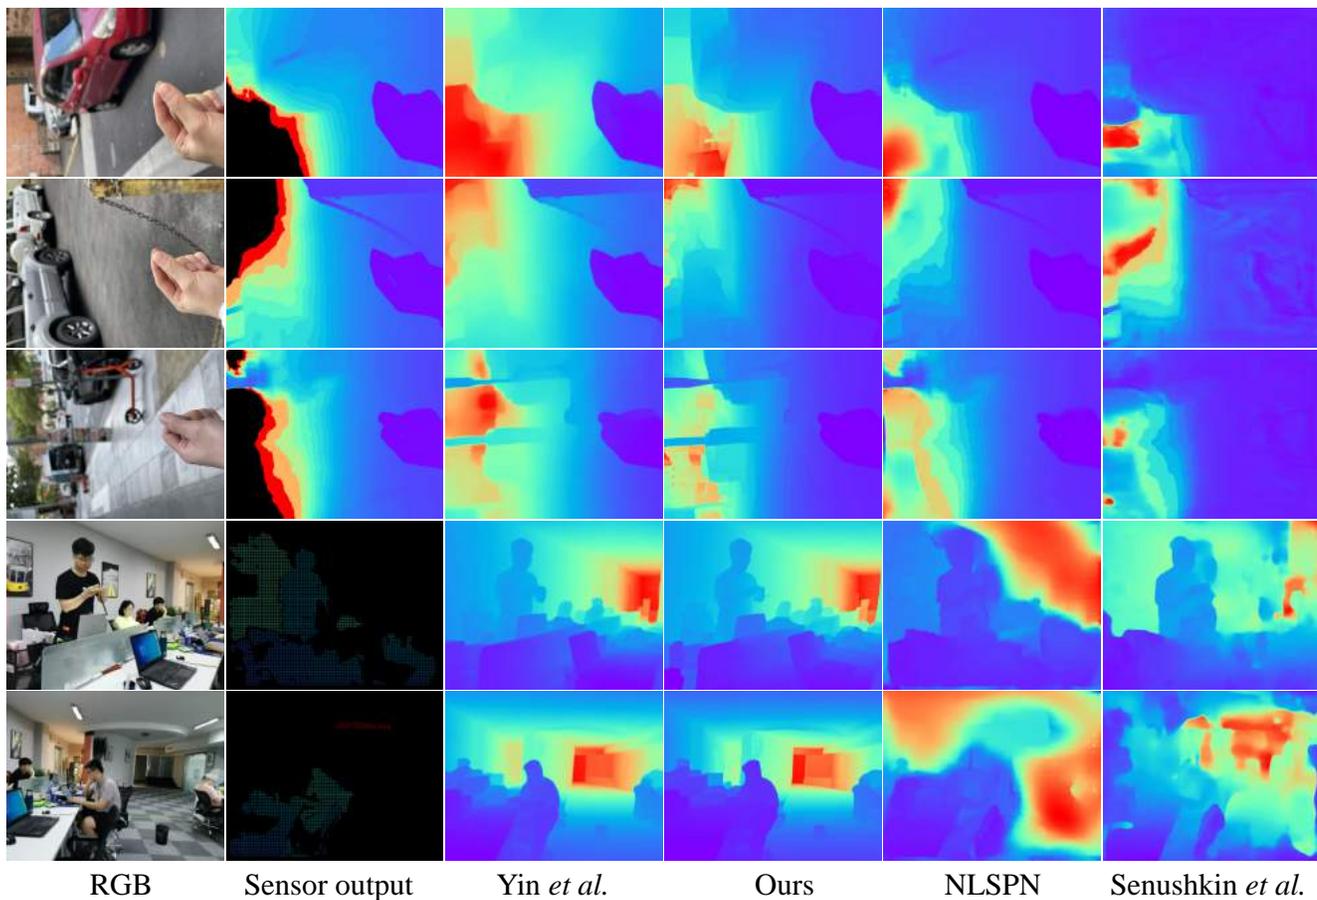


Figure 14. The comparison of completing iPhone and Huawei phone captured depths. The first 3 rows examples are captured by iPhone, while the others are obtained by Huawei. We can see that our method can recover better results than other methods.

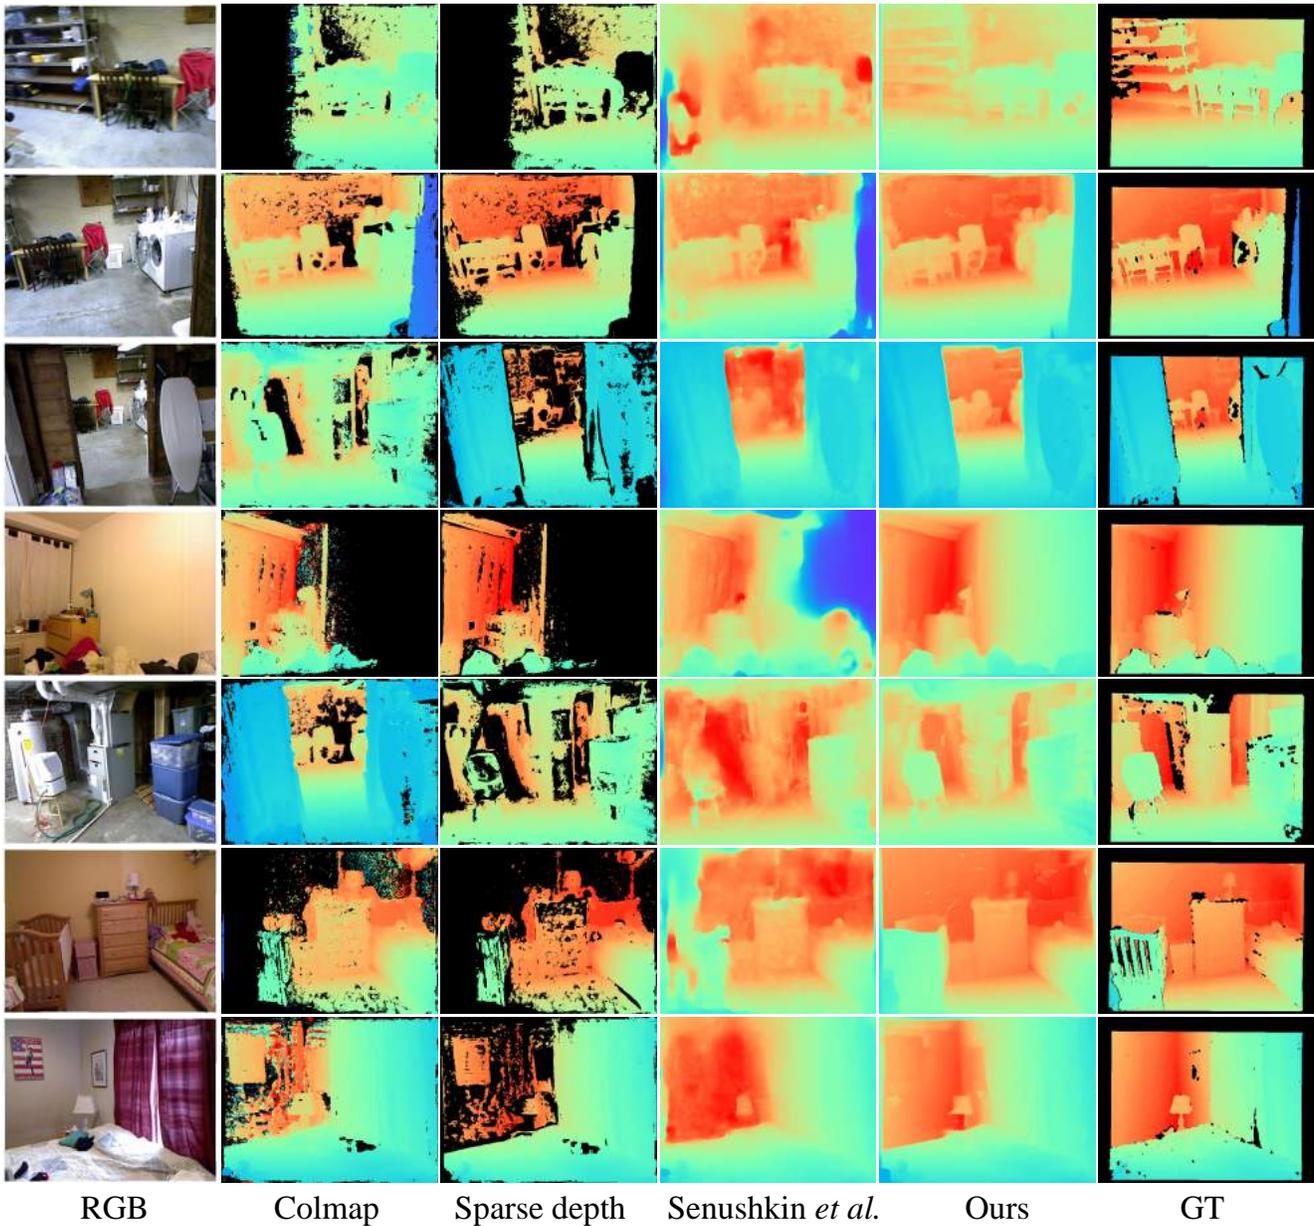


Figure 15. Complete noisy depth maps on sampled 16 NYU scenes. We use the colmap [30] to create the noisy sparse depth for completion. Our method is consistently better than other methods on all scenes.

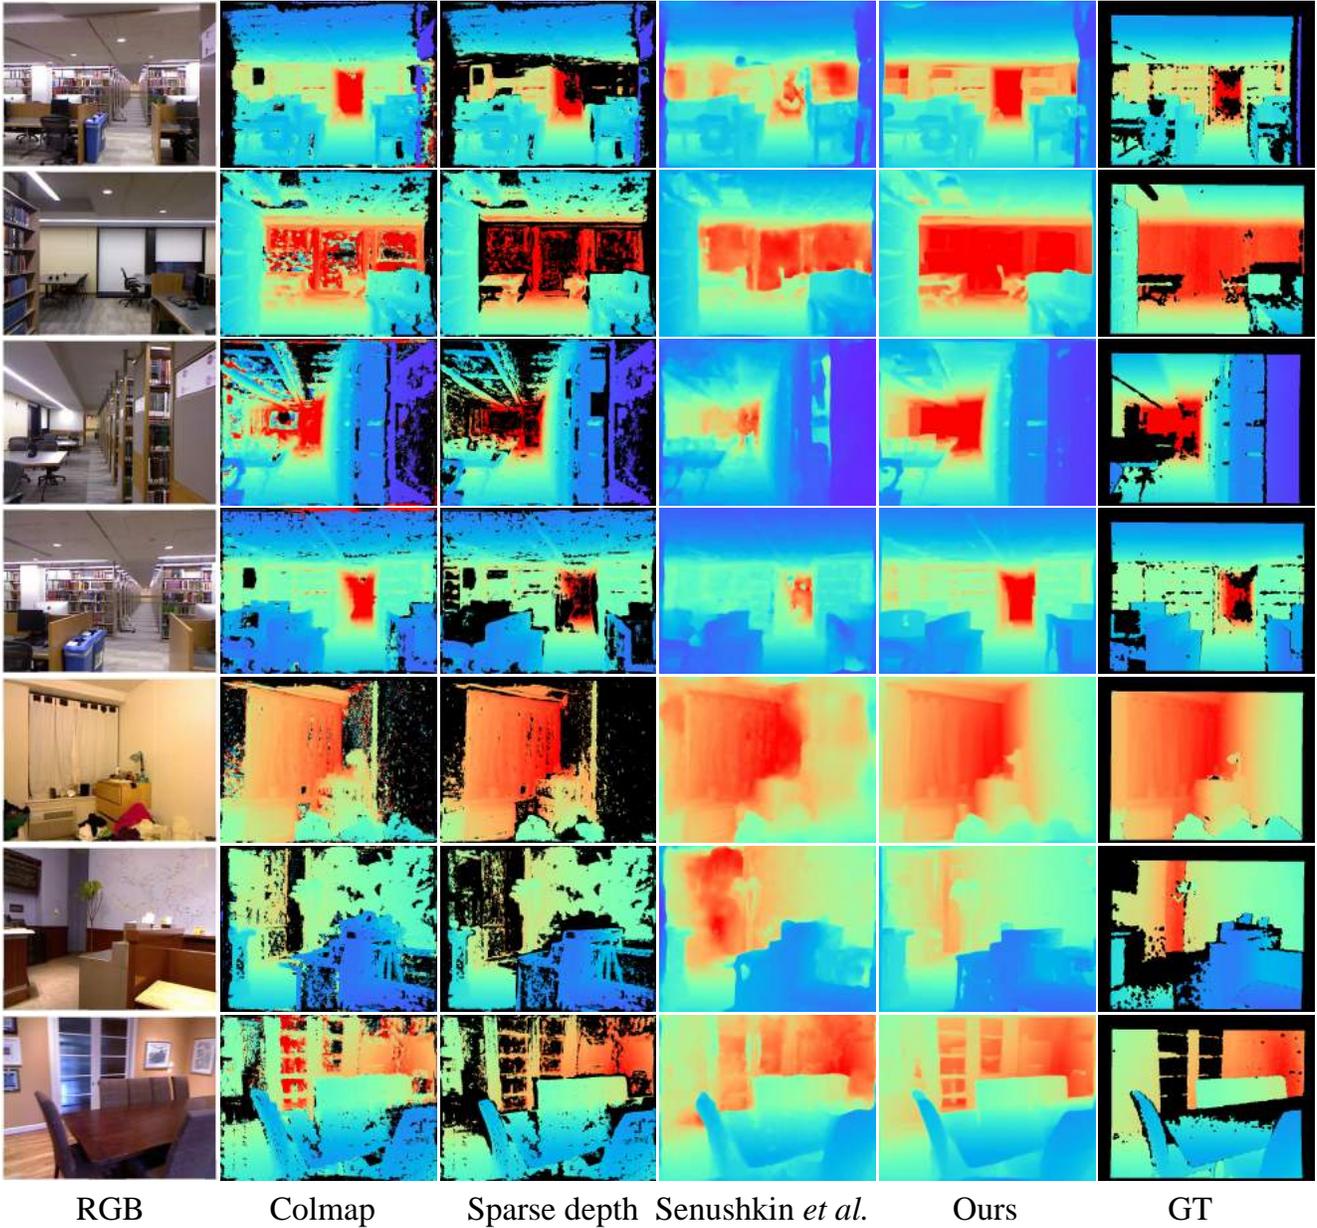


Figure 16. Complete noisy depth maps on sampled 16 NYU scenes. We use the colmap [30] to create the noisy sparse depth for completion. Our method is consistently better than other methods on all scenes.