Balanced Depth Completion between Dense Depth Inference and Sparse Range Measurements via KISS-GP

Sungho Yoon[1] and Ayoung Kim[2*]

Abstract— Estimating a dense and accurate depth map is the key requirement for autonomous driving and robotics. Recent advances in deep learning have allowed depth estimation in full resolution from a single image. Despite this impressive result, many deep-learning-based monocular depth estimation (MDE) algorithms have failed to keep their accuracy yielding a meter-level estimation error. In many robotics applications, accurate but sparse measurements are readily available from Light Detection and Ranging (LiDAR). Although they are highly accurate, the sparsity limits full resolution depth map reconstruction. Targeting the problem of dense and accurate depth map recovery, this paper introduces the fusion of these two modalities as a depth completion (DC) problem by dividing the role of depth inference and depth regression. Utilizing the state-of-the-art MDE and our Gaussian process (GP) based depth-regression method, we propose a general solution that can flexibly work with various MDE modules by enhancing its depth with sparse range measurements. To overcome the major limitation of GP, we adopt Kernel Interpolation for Scalable Structured (KISS)-GP and mitigate the computational complexity from $O(N^3)$ to $O(N)$. Our experiments demonstrate that the accuracy and robustness of our method outperform state-of-the-art unsupervised methods for sparse and biased measurements.

I. INTRODUCTION

Depth sensing is one of the most essential features for autonomous driving and mobile robots so as to avoid obstacles and to build localization and mapping. Toward this objective, industries seek efficient methods to build a dense depth-sensing system with low prices.

Toward this objective, monocular depth estimation (MDE) is used to predict a dense depth map only with a single image. It has been actively studied in computer vision and robotics area. Recent researches have found that the depth scale can be learned and predicted even in the single-camera setup. Using such property, the recent simultaneous localization and mapping (SLAM) studies have utilized MDE to reduce the uncertainty of the depth initialization, improving its performance [1, 2, 3]. However, it tends to have some limitations in its accuracy because it needs to learn many types of monocular cues and to reason the dense depth map based on the overall features. Moreover, MDE does not use any geometric relationship such as triangulation, which signifies that it is a fundamental ill-posed problem.

Another topic, depth completion (DC), is used to boost depth inference performance using both an image and depth

[1]S. Yoon is with the Robotics Program, KAIST, Daejeon, S. Korea `sungho.yoon@kaist.ac.kr`

[2]A. Kim is with the Department of Civil and Environmental Engineering, KAIST, Daejeon, S. Korea `ayoungk@kaist.ac.kr`
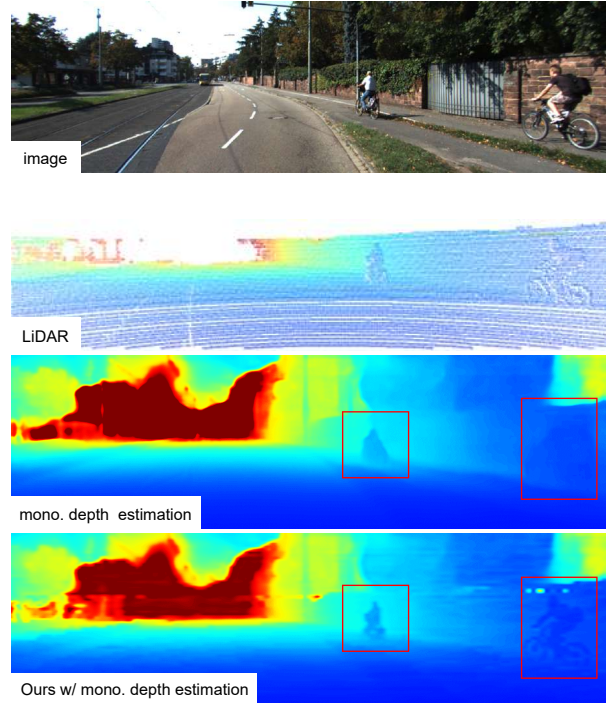
Fig. 1: Depth completion in the KITTI dataset. The dot size is inflated for visualization (the second row). The third-row image shows that the shapes of men and bicycles are blurry as a result of the monocular depth estimation using an image. Our method improves the exact shapes using LiDAR and offers robustness on sparse or biased depth measurements.

measurements. While accurate dense mapping has not been achieved only using the 3D feature points of SLAM or the sparse-point cloud of a lightweight LiDAR, DC accompanied by an image has presented a more powerful performance. Similar to MDE, this topic has some challenging problems for the three reasons. First, it has difficulties handling multiple sensing modalities such as image and depth measurements. Second, the 3D measurements are highly sparse and irregularly spaced depending on the environment and sensor characteristics. Third, finding a feasible dataset with well-labeled ground truth is still lacking.

In this work, we propose a back-end module wherein the result of MDE can be directly improved using sparse depth measurements, as shown at Fig. 1.

- We propose a general solution for the DC problem, which can be used with various state-of-the-art MDE modules by means of monochrome, thermal, RGB, or multiple cameras.
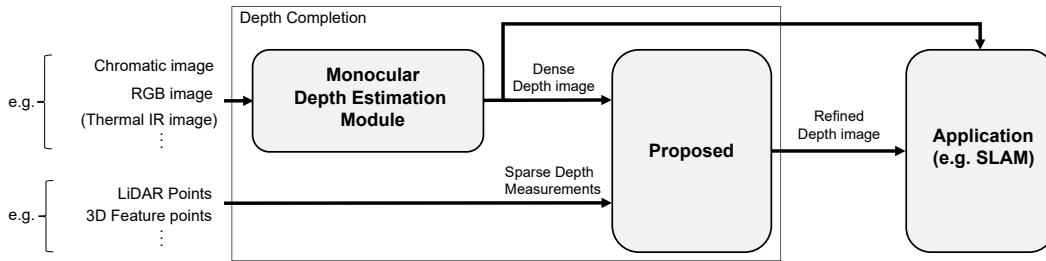
Fig. 2: System diagram. The proposed method utilizes a state-of-the-art depth prediction module specifically designed for each imaging modality. Then together with sparse depth measurements such as LiDAR, a depth regression is performed to make a refined depth image using the Gaussian process (GP).

- The proposed method is especially robust in the cases where the sparsity or bias of range measurements can vary or become unexpectedly deformed. Compared to the deep-learning-based DC methods that work appropriately only in their trained situation, our GP-based method maintains stable performance even in changing sensor conditions.
- Our evaluation method reveals the effect of changing sparsity or bias of range measurements, which we may have overlooked in the DC problem.

II. LITERATURE REVIEW

Monocular Depth Estimation (MDE): MDE aims to predict dense depth from a single image. The classical methods utilize meaningful cues on an image such as perspective projection, texture gradients, shadow, focus/defocus, or relative size [4, 5, 6]. Recently, the performance of depth estimation has been significantly improved by exploiting deep neural networks. Fu et al. [7] proposed a supervised depth estimation using a spacing-increasing discretization (SID) strategy and an ordinal regression loss. Godard et al. [8] trained monocular depths in a self-supervised manner using a left-right depth consistency term with a stereo camera. They also proposed another self-supervised method, called Monodepth2, using minimum reprojection loss and a multi-scale sampling method [9]. Liu et al. [10] estimated not only its depths but also its uncertainty. Yin et al. [11] improved their depth-prediction accuracy by enforcing the geometric constraints of virtual normal direction to their loss function. Kuznietsov et al. [12] introduced a training method for their semi-supervised approach using sparse LiDAR measurements for a supervised depth cue and using stereo image pairs for an unsupervised training cue. As another semi-supervised approach, Ramirez et al. [13] showed that learning depth with semantic labels improved the estimation performance. Chen et al. [14] presented a neural network inferring geometric depth and semantic segments at the same time as their content consistency check.

Depth Completion (DC): DC utilizes sparse range measurements to predict an image with a denser depth. Ma et al. [15] approached the DC problem within the compressive sensing technique. By minimizing the number of edge candidates given sparse depth measurements, they recovered a 3D map within its error bounds. As more recent works have started using deep learning, they have introduced a supervised auto-encoder architecture that uses an RGB image and sparse depth measurements as inputs [16]. DeepLi-DAR [17] presented a method using deep-surface normals as an intermediate representation toward dense depth in an outdoor scene. Their network also predicted a confidence mask to select the reliable raw data of LiDAR. Cheng et al. [18] proposed a convolutional spatial propagation network (CSPN) that learned the affinity matrix for depth estimation. The CSPN has been effective in recovering structural details of the depth map. They [19] further developed CSPN++, which improved their effectiveness and efficiency through context and resource-aware CSPN. The CSPN++ learned two hyperparameters to select adequate convolutional kernel sizes and the number of iterations from the predefined architecture. Gansbeke et al. [20] implemented a global and local depth map, and their depth maps were merged using each confidence map for the final depth.

The recent studies have reported limitations in supervised methods in terms of generalizability over various environments. As an alternative, unsupervised or self-supervised methods were also actively studied in DC. Ma et al. [21] proposed a self-supervised method that adds a photometric loss and a second-order smoothness prior. As another unsupervised work, Yang et al. [22] exploited a conditional prior network (CPN) that estimated a probability of depths at each pixel. They trained the CPN in a virtual dataset that is similar to and has been tested in the KITTI depth completion dataset. However, the CPN failed in generalizing to another indoor dataset, NYU-Depth-V2.

In contrast to these previous unsupervised works, our method is a general DC solution with improved stability. Most of the DC research highly coupled the RGB image and depth measurements by concatenation at the early steps of their architecture. The existing approaches may lose flexibility when another imaging modality becomes additionally available. Furthermore, under sparse or biased range measurements, the robustness of performance is broken easily, which critically deteriorates safety and stability. Differing from these previous studies and overcoming limitations, we propose a loosely coupled method that divides a DC problem to MDE and depth regression. Our method overcomes such sparsity or bias issues of range measurements by propagating the precise depths within a limited length using GP.

III. PROPOSED METHOD

This paper reports a solution to combine dense depth estimation from deep learning with sparse-but-accurate depth measurement using GP as a regression back-end. The overall system architecture of our method is illustrated in Fig. 2. We also discuss a depth propagation strategy in the case of measurements like LiDAR that are regularly distanced.

A. Depth Measurements

1) Deep-Learning-Based Monocular Depth Estimation: Recently introduced MDE predicts a depth map of a single image through the convolutional neural network (CNN) architecture. Estimating the depth map from a single image is truly an advantage; however, depth accuracy is still very limited compared to the existing model-based triangulation or direct measurement. In our pipeline, we adopted monodepth2 [9] to obtain depth estimation in a full-size image. For example, the depth estimation root mean square error (RMSE) of the monodepth2 ranges from 3 to 4 meters when tested over the KITTI odometry and DC dataset.

2) Sparse Depth Measurement: Unlike the depth inference of MDE, 3D depth can be measured directly and indirectly. An active depth sensor like LiDAR is a direct method that utilizes illuminating laser light and measures reflected light. As an indirect method with a monocular camera, depth can be measured through triangulation between corresponding feature points and it is refined by optimization like bundle adjustment (BA). However, both direct and indirect methods provide sparse measurements. The densest LiDAR sensor is 128 rays which is still quite sparse when projected on an image plane with a significant cost for a sensor.

B. Depth Integration via KISS-GP

1) Problem Formulation using Exact GP: To create a more precise map than MDE and a denser map than sparse depth measurement, two different properties should be integrated effectively. The proposed method using GP estimates the posterior depth map $\mathbf{Y}_*$ as a mean function μ_* and its variance Σ_*. We assume the joint probability of Gaussian random variables between test data $\mathbf{X}_* = \{\mathbf{x}_{1:M}\}$ and training data $X = \{\mathbf{x}_{1:N}\}$, where $\mathbf{x}_i$ is an i-th pixel coordinate on an image $\mathbf{x}_i = (\mathbf{u}_i, \mathbf{v}_i)$ and $\mathbf{M}$ and $\mathbf{N}$ are numbers of pixels. In our problem, we set the size $\mathbf{M}$ as the number of pixels in a refined depth map and the size $\mathbf{N}$ as the number of pixels in both imaging depth map and ranging measurements. The size $\mathbf{M}$ of $\mathbf{X}_*$ can be larger than the size $\mathbf{N}$. This shows that we can enlarge the output depth map without linear extrapolation within GP. The predictive mean and covariance of $\mathbf{Y}_*$ for a test input $\mathbf{X}_*$ is obtained within a conditional Gaussian distribution $p(Y_*|\mathbf{X}_*, X, Y) = \mathcal{N}(\mu_*, \Sigma_*)$:

$$
\begin{aligned}
\mu_* &= K(\mathbf{X}_*, X)^\top (K + \sigma_n^2 I)^{-1} \mathbf{Y} \\
\Sigma_* &= K(\mathbf{X}_*, \mathbf{X}_*) - K(\mathbf{X}_*, X)^\top (K + \sigma_n^2 I)^{-1} K(\mathbf{X}_*, X).
\end{aligned}
\tag{1}
$$

In the equation, K represents a kernel function $K(X, X)$, and σ_n is the observation noise of Y.

Algorithm 1 Depth Completion with KISS-GP
_____
1: **Input:** Dense depth image $I_{m \times n}$, Sparse measurement image $J_{m \times n}$, Observation noise σ_{DL} and σ_{meas}, and length-scale l
2: Set of training inputs $X = \{(u_i, v_i)_{1:mn+p}\}$,
3: Set of training outputs $Y = \{(d_i)_{1:mn+p}\}$,
4: Set of testing inputs $X_* \leftarrow$ MATTOARRAY($D_{in} = I_{m' \times n'}$)
5: $K_{X_* X} \leftarrow$ COMPUTE-KERNELMAT(X_*, X, l)
6: $Y_* \leftarrow$ COMPUTE GP-MEAN($K_{X_* X} \{K_{SKI} + \sigma_{mn+p} I\}^{-1} Y$)
7: $D_{out} \leftarrow$ ARRAYTOMAT(Y_*)
8: **Output:** Refined depth image D_{out}
_____

2) Training Set: For targeting a full-resolution depth image of $m \times n$, we used incoming deep learning-based depth estimation ($X_{DL} = \{\mathbf{x}_{1:mn}\}$) and the p number of measured sparse-but-accurate depth ($X_{Meas} = \{\mathbf{x}_{1:p}\}$) for training data. This training set of $X = \{X_{DL}, X_{Meas}\}$ is the grid coordinates of depth image $(u, v)_{i=1:mn+p}$, and $y_{i=1:mn+p}$ is its depth values $D(u, v)_{i=1:mn+p}$.

3) KISS-GP Implementation: The main challenge in using exact GP is its high computation $O(n^3)$ and storage cost $O(n^2)$, along with the fact that it requires calculating a large inverse matrix of kernels. As shown in equation (1), we have been limited to handle large data such as a high-resolution image. To alleviate this issue, we adapted KISS-GP [23] that exploits kernel approximation strategy. By combining an inducing point method, such as a Subset of Regression (SoR) with kernel interpolation, the size of the kernel is approximated into a small size:

$$
\begin{aligned}
K(X, X) &\overset{\text{SoR}}{\approx} K(X, U) K(U, U)^{-1} K(U, X) \\
&\overset{mat.intrpl}{\approx} W K(U, U) K(U, U)^{-1} K(U, U) W^\top \\
&= W K(U, U) W^\top = K_{\text{SKI}}.
\end{aligned}
\tag{2}
$$

The inducing point set U represents the effectively summarized data among the original input data by SoR, and the interpolation weights matrix W contains the relative distances from original input data X and the inducing point U. Since the W is extremely sparse, K_{SKI} approximates the original GP kernel function K with $O(n + m^2)$ computation and storage. In addition to structured kernel interpolation (SKI), KISS-GP in 2D inputs uses a Kronecker structure, and it finally reduces computation $O(2m^{1+1/2})$ and storage $O(n + 2m)$. In practice, we exploited a GPyTorch implementation [24] to enhance computational speed further. The GPyTorch is a modern GP toolbox using PyTorch [25], which reduces computational complexity using Blackbox Matrix-Matrix multiplication (BBMM) with GPU acceleration. The BBMM uses a modified conjugated gradients algorithm and enables the effective use of GPU hardware [24].

4) Kernel and Hyperparameters: We have selected the Matern kernel for our method, which is a stationary kernel to utilize grid interpolation in KISS-GP. Among stationary kernels like squared exponential, the Matern kernel is more beneficial in incorporating potentially abrupt change near the edges. In terms of hyperparameters, the observation noises and length scale should be defined. The depth observation noises from two different modalities are different in our case.

(a) RGB image

(b) full scan

1/8

(c) uniformly sampled scan

1/4x
1/2x
1x

(d) horizontally sampled scan

1x   1/2x   1/4x 1/8x ....

(e) vertically sampled scan

Fig. 3: Three different sampling strategies for the evaluation of sparsity and biased measurements. (a) and (b) The original RGB image and LiDAR raw data. (c) The uniform sampling. (d) and (e) The full LiDAR data is subsampled into eight levels of biased measurements. The performance of depth extrapolation can be evaluated with horizontally and vertically sampled data.

While the sparse measurements are precise (i.e., low noise), the dense depth image would have higher noise. Since the different sensors have dissimilar accuracy, the observation noise should be uniquely defined depending on each depth model. We used two observation noise terms from different sensor modalities, namely σ_{DL} and σ_{Meas}.

Together with these observation noises, we needed to determine the length scale. Length scale l defines the amount of depth propagation between each pixel. Although it is related to the measurements' sparsity, it is correlated with the noise factor because it is better for the length scale to be large when the observation noise of a dense depth image is

low enough. In the LiDAR case, in which measurements are almost regularly distributed, the length scale can be easily assigned with the distance between each measurement point. In this paper, we used $\sigma_{DL} = 0.05$, $\sigma_{Meas} = 0.001$ and $l = 1.5$ for KITTI.

IV. EXPERIMENTS AND RESULTS

In this section, we present a quantitative and qualitative evaluation of the proposed method when tested on both indoor and outdoor. The result is also available from https://youtu.be/x8n0lvjvorg.

A. Datasets

For the evaluation of our method, we used two different datasets on outdoor and indoor: KITTI depth completion [26] and NYU-Depth-V2 [27].

In the KITTI depth completion dataset, we resized the original image to one third. When projecting LiDAR onto an image plane, occlusion in LiDAR may cause duplicated depth measurement for a pixel. Existing DC studies resolved this problem by designing a LiDAR confidence map [17, 20, 28]. By learning the LiDAR's confidence map from a training dataset, they filtered out such unreliable duplicated depth data. Instead of using the trained mask, we alleviated such occlusion effects by selecting the closest distance in a downsampling procedure. For sparsity and bias of depth measurements, we evaluated the performance over three sampling strategies, as shown in Fig. 3. In the NYU-Depth-V2 dataset, we resized the original images to be half their size and center-cropped them to be 304×228, performed in previous works [16, 17]. In the experiment, we used a PC with an Intel Core i7-6700 CPU with 3.4GHz, 64GB of RAM, as well as an NVidia GeForce RTX 2080Ti GPU.

As a result of the evaluation, the proposed method directly improves the frond-end MDE when additional sparse measurements are given. The accuracy of the proposed method outperforms a state-of-the-art method that uses unsupervised learning in terms of sparsity and bias of LiDAR measurements, and it is also comparable with the full LiDAR scan case. In the following section, we present experimental results at each dataset. Table I lists the error metrics used in the evaluation. The RMSE is the main evaluation metric in both datasets.

B. KITTI Depth Completion

The KITTI depth completion dataset provides RGB images, LiDAR raw data, and ground-truth depth data of outdoor scenes. The ground-truth data were generated by accumulating 11 sparse laser, and the outliers of laser scans were

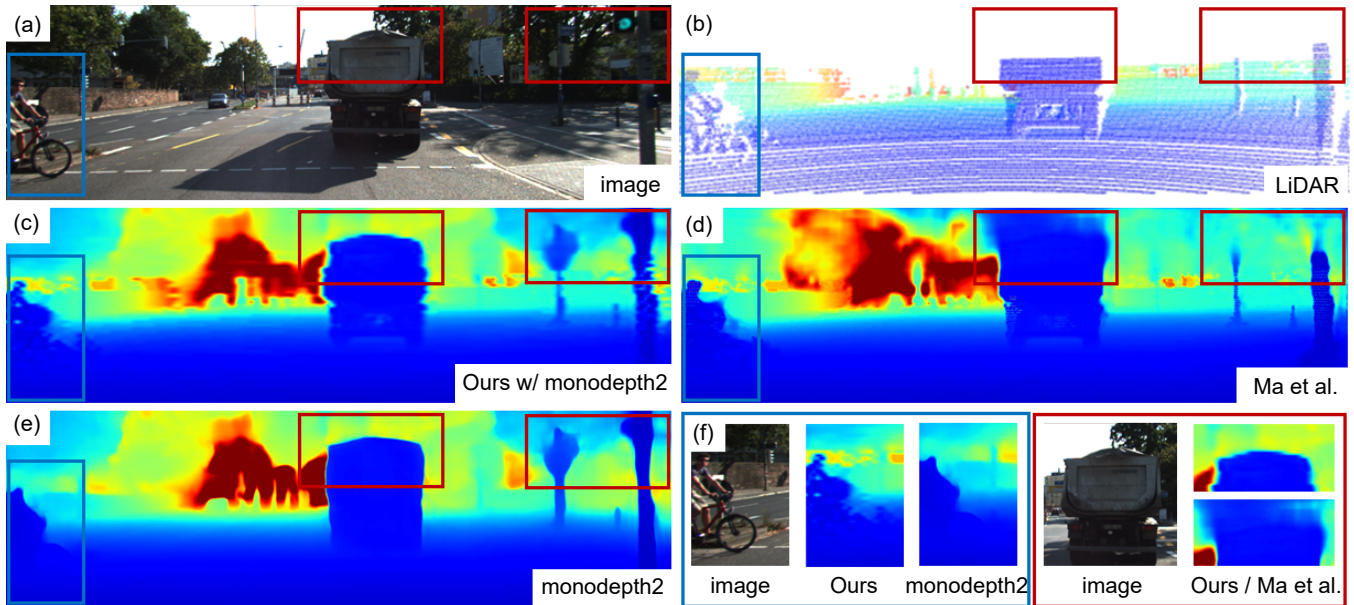| Metric	Details	KITTI	NYU-Depth-V2
RMSE	root mean squared error	✓	✓
MAE	mean absolute error	✓	
InvRMSE	inversed root mean squared error	✓	
InvMAE	inversed mean absolute error	✓	
REL	mean absolute relative error		✓
δ_i	% of pixel count where its relative error is within 1.25^i		✓

TABLE I: Evaluation metrics

Fig. 4: Qualitative comparison to monodepth2 [9] and Ma et al. [21] over the KITTI dataset. In all sample figures, the warmer color denotes farther distance. (a) and (b) The RGB image and LiDAR raw data. Depth map from (c) the proposed method, (d) Ma et al. [21], and (e) monodepth2 [9] are shown respectively. The blue box represents our improvement from monodepth2. The red box illustrates that our method keeps the shape of a truck better than Ma et al. [21] in the region without LiDAR data.
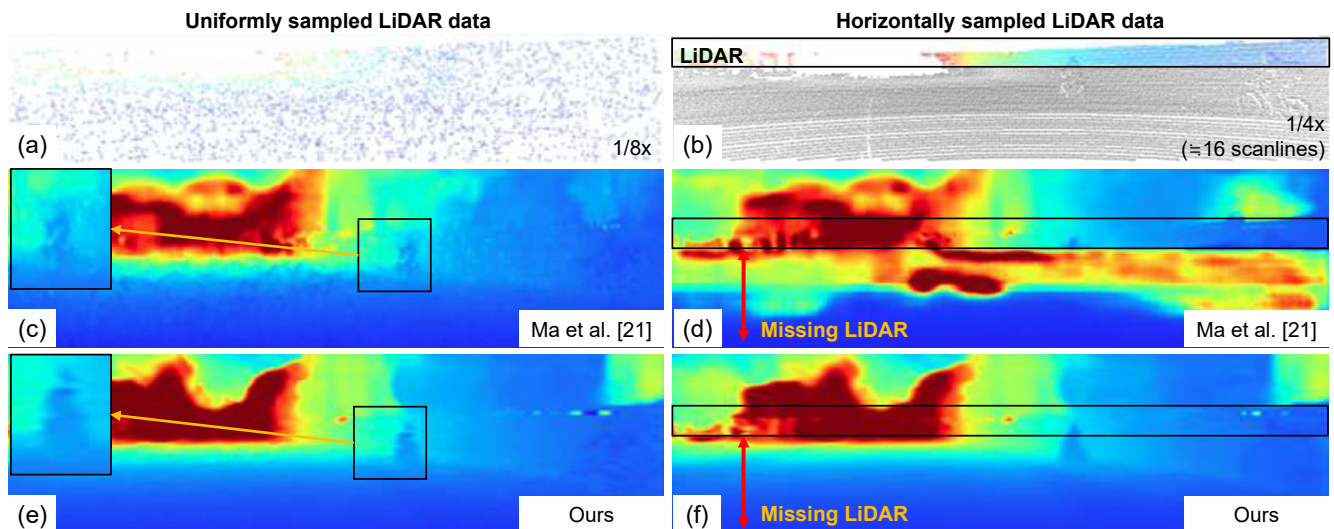

Fig. 5: Qualitative comparison to Ma et al. [21] given the sparse and bias LiDAR measurements in the KITTI dataset. (a) and (b) Uniformly sampled LiDAR data and horizontally sampled LiDAR data. The depth map from (c) and (d) result from Ma et al. [21]. And the depth maps from (e) and (f) result from the proposed method. The results show that the proposed method is more robust on the sparsity or bias change of range measurements than Ma et al. [21].

removed using stereo reconstruction [26]. The evaluation of DC is performed with these semi-dense ground-truth points. This dataset includes 85,898 training data, 1,000 validation data, and 1,000 test data. We evaluated the proposed method and a state-of-the-art method in the validation set.

Fig. 4 and Fig. 5 show the qualitative performance of the proposed method. The proposed method utilizes monodepth2 [9] as a front-end module. As shown in the blue box in Fig. 4(e), a blurry shape of a man riding a bicycle takes the output of MDE. However, the proposed method can be used as a back-end module given sparse depth measurements so that the man's shape becomes sharper and more accurate, as shown in the blue box in Fig. 4(c). The proposed method also outperforms a state-of-the-art DC method [21], as shown in the red boxes in Fig. 4(c) and (d). Our module effectively sustains the state-of-the-art depth inference of monodepth2 at

(a) RMSE comparison to [21] against LiDAR points' sparsity and bias

point density	Ma et al.[21] (uniform)	Ma et al.[21] (horizontal)	Ma et al.[21] (vertical)	Ours w/ [9] (uniform)	Ours w/ [9] (horizontal)	Ours w/ [9] (vertical)
0 (MDE)	24092.25	24092.25	24092.25	**3951.26**	3951.26	3951.26
1/64x	7367.46	23937.48	23924.91	**3700.71**	3920.50	3916.37
1/32x	5466.10	23769.53	23826.57	**3508.57**	3916.17	3910.60
1/16x	4261.62	23290.30	23542.19	**3200.06**	3890.24	3905.22
1/8x	3139.47	22475.57	22837.19	**2760.77**	3633.68	3872.23
1/4x	**2249.94**	20584.628	20962.07	2252.38	2614.04	3716.86
1/2x	**1646.55**	5741.87	16050.62	1822.76	1696.34	3103.81
1x(full)	**1343.33**	1343.33	1343.33	1593.37	1593.37	1593.37

(b) RMSE metric result (mm)

Point density (# scan lines)	Ours				Ma et al. [21]			
	RMSE [mm]	MAE [mm]	iRMSE [1/km]	iMAE [1/km]	RMSE [mm]	MAE [mm]	iRMSE [1/km]	iMAE [1/km]
1/4x (16)	**2614.04**	**1141.02**	**14.89**	**5.68**	20584.63	13550.55	54.97	43.45
1/2x (32)	**1696.34**	**705.02**	**18.17**	**4.36**	5741.87	3132.26	45.32	30.65
full (64)	1593.37	547.00	27.98	2.36	**1343.33**	**358.66**	**4.28**	**1.64**

(c) RMSE metric result (mm) in the horizontally sampled scan

Fig. 6: Quantitative comparison to Ma et al. [21] in the KITTI depth completion dataset regarding sparse and biased range data. Best performance is marked in bold. The point density of 0% represents the result of monocular depth estimation [9] without the proposed method.

the region where the LiDAR is not received. Fig. 4(f) shows the overall comparison of the proposed method to monodepth2 and Ma et al. [21]. Furthermore, Fig. 5 represents the DC results under uniformly and horizontally sampled range data. This shows that Ma et al. [21] found it difficult to cope with the changes in various range data, whereas our proposed module enhances MDE effectively with the given LiDAR data.

Next, we present the quantitative results of the proposed methods in terms of sparsity, biased measurements, and modality change.

1) Sparsity of LiDAR Measurements: We uniformly sampled the eight-level ratio from the full-scan data (Fig. 3) to see the effects of sparsity. We compared the proposed method using unsupervised MDE [9] against a self-supervised DC method [21]. We use the open-sourced version of the Ma et al. [21] for this evaluation. As shown in Fig. 6(a) and (b), the results reveal that the proposed method presents comparability with one-fourth of the samples of full LiDAR scans and that it outperforms after one-eighth of the LiDAR scan, whereas [21] performed better than us by 250.04 mm on the RMSE metric when the full LiDAR scans are used. As aforementioned, the deep-learning-based methods are able to train the LiDAR point's confidence in a given LiDAR range map, whereas our method only picks the closest point in the resized range map. We consider that this effect deteriorates our results when the number of LiDAR points increases. Note that the RMSE error of our method increases until the level of MDE, whereas the RMSE error of Ma et al. [21] increases as a power function cx^p as described in their paper. This means that our method performs more stably in sparsity variance, which is especially important in robotics usage cases.

2) Biased LiDAR Measurements: The proposed method is not only robust toward data sparsity but also toward the distribution pattern of bias. This may involve installing

configuration and various field of view (FOV) specifications of the perceptual sensors. For validation, we sampled the full LiDAR data horizontally and vertically to see the effects of biased FOV. The KITTI dataset originally provided biased LiDAR data in terms of RGB images for the limited FOV of a LiDAR sensor. As a result, we were unable to evaluate the depth accuracy where the LiDAR points could not reach. As alternative methods, we divided the LiDAR's FOV horizontally and vertically and used LiDAR raw data on only a portion of the area so that we could evaluate the remaining areas that the LiDAR points did not exist. As can be seen in Fig. 6, the proposed method also outperforms the state-of-the-art method [21] except for the case of a full LiDAR scan. Compared to the uniformly sampled sets, the RMSE of Ma et al. [21] increased rapidly. From the experimental results, we consider that separating a depth inference module and its depth regression would be advantageous for the stability of DC against eccentric measurements.

3) Robustness to Modality Variance: In Fig. 7, we show the qualitative result that the proposed method can be adapted to another front-end MDE used with a gray-scaled input image. When MDE is designed, it is considered within the properties of its own imaging modality. For example, the MDE using thermal infrared (IR) images [29] discusses the unique features of thermal images, such as blurry imaging and thermal time variance, in its architecture. By exploiting its own best performance on imaging depth inference, the proposed method only focuses on the depth-to-depth regression, in which GP can perform strongly in the theoretically proven way.

C. NYU-Depth-V2

The NYU-Depth-V2 dataset provides RGB and depth image, as well as accelerometer data collected by Microsoft Kinect. As a subset of the raw data, the labeled data consist

Fig. 7: Qualitative result of the proposed method that used a different imaging modality. For a given gray image on the top left, the red box represents a man riding a bicycle. The depth image is worse than its original output (bottom left) due to the different modality. The proposed method successfully recovers the shape of a man riding a bicycle (bottom right).

(a) RMSE comparison to [16] and [22]

point density (%) (#points)	Ma et al.[16] (sup.)	Yang et al.[22] (unsup.)	Ours w/ [11]
0 (0)	514	-	500.31
0.01 (7)	-	3829	506.3
0.02 (14)	-	2623	500.92
0.03 (20)	351	-	481.6
0.06 (42)	-	1391	435.85
0.07 (50)	281		435.85
0.3 (200)	230	569	308.89

(b) RMSE metric result (mm)

Fig. 8: Quantitative comparison to Ma and Karaman [16] and Yang et al. [22] in the NYU-Depth-V2 dataset. The point density of 0% represents the result of monocular depth estimation [11] without the proposed method.

of synchronized pairs of RGB and depth images with dense labels. The dataset is officially split into 795 training images and 654 test images. We evaluated the proposed method in the test set. As sparse range measurements in this experiment, we randomly selected a uniform and assigned number of depth points from the ground-truth depth map.

Since we could not find an open-sourced and unsupervised MDE method implemented in NYU-Depth-V2, we inevitably used a supervised MDE method, VNL [11] as our front-end module. Therefore, we compared our results with a supervised method [16] and an unsupervised method [22], as illustrated in Fig. 8. The proposed method outperforms the state-of-the-art unsupervised method [22] in every point density case, and it is also comparable with a recent supervised method [16]. Note that the proposed methods may have the potential for improvement considering the ability to adapt with state-of-the-art MDE, although the RMSE of Ma and Karaman [16] is 79 mm higher in the case of 200 samples' case (Table. II). In the overall results of our

method, we experienced that the depth map of MDE becomes more blurred after passing our method. We consider this is due to the fact that GP does not directly understand object shape. Therefore, we remain that our method can be improved with semantic segmentation or image-guided edge-preserving methods for future work. The quantitative and qualitative results are shown in Table. II, Fig. 8, and Fig. 9.

V. CONCLUSION

In this paper, we introduced a depth-regression method for DC using sparse-but-accurate depth measurements and dense-but-inaccurate depth inference from deep-learning-based estimation. Unlike the other DC methods, we divided the DC architecture into two steps (i.e., MDE and depth regression using GP), so that we could achieve robust and accurate performance even when given sparse and biased range measurements. We demonstrated that the proposed method outperformed state-of-the-art unsupervised methods in KITTI depth completion and NYU-Depth-V2 datasets with regard to the sparsity and bias of measurements. We also showed our method's flexibility, which can be used as a plugin module in place of other depth-estimation modules that may have a different imaging modality. We believe that the proposed methods can be used in a SLAM application that needs a dense visualized map and that only observes a few biased feature points in its environment.

# points	RMSE	REL	δ_1	δ_2
0	500.31	10.79	87.94	97.52
50	435.85	8.89	90.77	98.18
200	308.89	5.76	95.22	99.06
500	230.66	3.97	97.40	99.54
1000	186.85	3.13	98.40	99.77

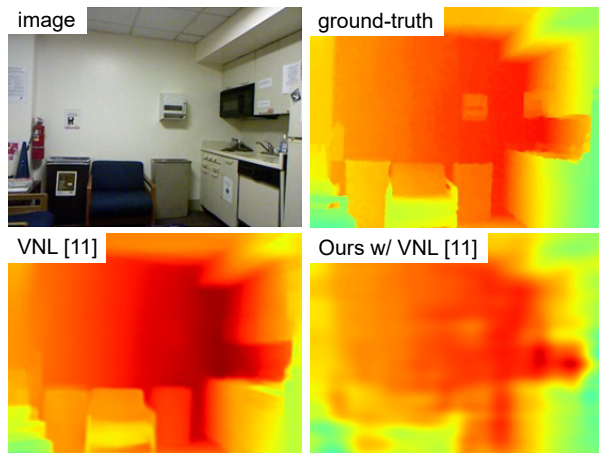TABLE II: Overall results with VNL[11] in the NYU-Depth-V2

Fig. 9: Qualitative result in the NYU-Depth-V2 dataset. Our method successfully recovers an indoor depth image given 500 uniformly distributed samples (point density of 0.7%).

REFERENCES

[1] K. Tateno, F. Tombari, I. Laina, and N. Navab, "CNN-SLAM: real-time dense monocular SLAM with learned depth prediction," CoRR, vol. abs/1704.03489, 2017.

[2] S. Y. Loo, A. J. Amiri, S. Mashohor, S. H. Tang, and H. Zhang, "CNN-SVO: improving the mapping in semi-direct visual odometry using single-image depth prediction," CoRR, vol. abs/1810.01011, 2018.

[3] N. Yang, R. Wang, J. Stückler, and D. Cremers, "Deep virtual stereo odometry: Leveraging deep depth prediction for monocular direct sparse odometry," CoRR, vol. abs/1807.02570, 2018.

[4] M. Subbarao and G. Surya, "Depth from defocus: A spatial domain approach," Intl. J. of Comput. Vision, vol. 13, pp. 271–294, 1994.

[5] D. Hoiem, A. A. Efros, and M. Hebert, "Recovering surface layout from an image," Intl. J. of Comput. Vision, vol. 75, no. 1, October 2007.

[6] L. Ladicky, J. Shi, and M. Pollefeys, "Pulling things out of perspective," in Proc. IEEE Conf. on Comput. Vision and Pattern Recog., June 2014.

[7] H. Fu, M. Gong, C. Wang, K. Batmanghelich, and D. Tao, "Deep ordinal regression network for monocular depth estimation," CoRR, vol. abs/1806.02446, 2018.

[8] C. Godard, O. Mac Aodha, and G. J. Brostow, "Unsupervised monocular depth estimation with left-right consistency," in Proc. IEEE Conf. on Comput. Vision and Pattern Recog., 2017.

[9] ——, "Digging into self-supervised monocular depth estimation," CoRR, vol. abs/1806.01260, 2018.

[10] C. Liu, J. Gu, K. Kim, S. G. Narasimhan, and J. Kautz, "Neural rgb->d sensing: Depth and uncertainty from a video camera," CoRR, vol. abs/1901.02571, 2019.

[11] W. Yin, Y. Liu, C. Shen, and Y. Yan, "Enforcing geometric constraints of virtual normal for depth prediction," CoRR, vol. abs/1907.12209, 2019.

[12] Y. Kuznietsov, J. Stückler, and B. Leibe, "Semi-supervised deep learning for monocular depth map prediction," CoRR, vol. abs/1702.02706, 2017.

[13] P. Z. Ramirez, M. Poggi, F. Tosi, S. Mattoccia, and L. di Stefano, "Geometry meets semantics for semi-supervised monocular depth estimation," CoRR, vol. abs/1810.04093, 2018.

[14] P.-Y. Chen, A. H. Liu, Y.-C. Liu, and Y.-C. F. Wang, "Towards scene understanding: Unsupervised monocular depth estimation with semantic-aware representation," in Proc. IEEE Conf. on Comput. Vision and Pattern Recog., June 2019.

[15] F. Ma, L. Carlone, U. Ayaz, and S. Karaman, "Sparse depth sensing for resource-constrained robots," arXiv preprint arXiv:1703.01398, 2017.

[16] F. Ma and S. Karaman, "Sparse-to-dense: Depth prediction from sparse depth samples and a single image," CoRR, vol. abs/1709.07492, 2017.

[17] J. Qiu, Z. Cui, Y. Zhang, X. Zhang, S. Liu, B. Zeng, and M. Pollefeys, "Deeplidar: Deep surface normal guided depth prediction for outdoor scene from sparse lidar data and single color image," CoRR, vol. abs/1812.00488, 2018.

[18] X. Cheng, P. Wang, and R. Yang, "Depth estimation via affinity learned with convolutional spatial propagation network," CoRR, vol. abs/1808.00150, 2018.

[19] X. Cheng, P. Wang, C. Guan, and R. Yang, "CSPN++: learning context and resource aware convolutional spatial propagation networks for depth completion," CoRR, vol. abs/1911.05377, 2019.

[20] W. V. Gansbeke, D. Neven, B. D. Brabandere, and L. V. Gool, "Sparse and noisy lidar completion with RGB guidance and uncertainty," CoRR, vol. abs/1902.05356, 2019.

[21] F. Ma, G. V. Cavalheiro, and S. Karaman, "Self-supervised sparse-to-dense: Self-supervised depth completion from lidar and monocular camera," CoRR, vol. abs/1807.00275, 2018.

[22] Y. Yang, A. Wong, and S. Soatto, "Dense depth posterior (DDP) from single image and sparse range," CoRR, vol. abs/1901.10034, 2019.

[23] A. G. Wilson and H. Nickisch, "Kernel interpolation for scalable structured gaussian processes (KISS-GP)," CoRR, vol. abs/1503.01057, 2015.

[24] J. R. Gardner, G. Pleiss, D. Bindel, K. Q. Weinberger, and A. G. Wilson, "Gpytorch: Blackbox matrix-matrix gaussian process inference with GPU acceleration," CoRR, vol. abs/1809.11165, 2018.

[25] A. Paszke, S. Gross, F. Massa, A. Lerer, J. Bradbury, G. Chanan, T. Killeen, Z. Lin, N. Gimelshein, L. Antiga, A. Desmaison, A. Kopf, E. Yang, Z. DeVito, M. Raison, A. Tejani, S. Chilamkurthy, B. Steiner, L. Fang, J. Bai, and S. Chintala, "Pytorch: An imperative style, high-performance deep learning library," in Advances in Neural Information Processing Sys. Conf., 2019, pp. 8026–8037.

[26] J. Uhrig, N. Schneider, L. Schneider, U. Franke, T. Brox, and A. Geiger, "Sparsity invariant cnns," in Proc. Intl. Conf. on 3D Vis., 2017.

[27] P. K. Nathan Silberman, Derek Hoiem and R. Fergus, "Indoor segmentation and support inference from rgbd images," in Proc. European Conf. on Comput. Vision, 2012.

[28] A. Eldesokey, M. Felsberg, and F. S. Khan, "Confidence propagation through cnns for guided sparse depth regression," CoRR, vol. abs/1811.01791, 2018.

[29] N. Kim, Y. Choi, S. Hwang, and I.-S. Kweon, "Multispectral transfer network: Unsupervised depth estimation for all-day vision," in Proc. AAAI National Conf. on Art. Intell., 2018.