

Depth Completion via Inductive Fusion of Planar LIDAR and Monocular Camera

Chen Fu¹, Chiyu Dong¹, Christoph Mertz² and John M. Dolan^{1,2}

Abstract—Modern high-definition LIDAR is expensive for commercial autonomous driving vehicles and small indoor robots. An affordable solution to this problem is fusion of planar LIDAR with RGB images to provide a similar level of perception capability. Even though state-of-the-art methods provide approaches to predict depth information from limited sensor input, they are usually a simple concatenation of sparse LIDAR features and dense RGB features through an end-to-end fusion architecture. In this paper, we introduce an inductive late-fusion block which better fuses different sensor modalities inspired by a probability model. The proposed demonstration and aggregation network propagates the mixed context and depth features to the prediction network and serves as a prior knowledge of the depth completion. This late-fusion block uses the dense context features to guide the depth prediction based on demonstrations by sparse depth features. In addition to evaluating the proposed method on benchmark depth completion datasets including NYUDepthV2 and KITTI, we also test the proposed method on a simulated planar LIDAR dataset. Our method shows promising results compared to previous approaches on both the benchmark datasets and simulated dataset with various 3D densities.

I. INTRODUCTION

A robust perception system that understands the surrounding scene is crucial for autonomous vehicles. Instead of simply detecting the obstacles in 3D world coordinates, this system should provide additional information such as the heading angle, distance, and 3D shape of the obstacles [1], [2]. To achieve this goal, current autonomous vehicles are usually equipped with high-definition LIDAR (Light detection and Ranging) (e.g. Velodyne LIDAR). However, this family of sensors has certain limitations. Firstly, high-definition LIDAR is not affordable for commercial autonomous vehicles. Even though it provides rich depth information, it lacks detailed texture of the objects, which makes object detection and classification difficult. Relying on a single LIDAR is dangerous owing to potential sensor failure. Finally, compared to a RGB image, the point cloud of a LIDAR is sparse, especially for those objects that are far away from the ego-vehicle.

Fusing multiple planar LIDAR with an RGB camera is an affordable solution to these limitations. This system can provide the autonomous vehicle with sparse 3D information of the surrounding scene and semantic understanding. However, this type of system has a few challenges: 1) It

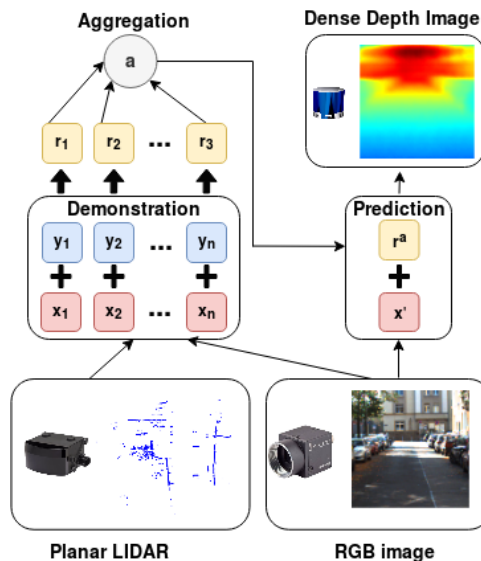


Fig. 1: We propose a novel late fusion architecture which better combines dense context features with sparse depth features by inductive fusion block. Our network takes data from affordable RGB camera and planar LIDAR and generates a depth image with similar density to that of high-definition LIDAR.

is difficult to detect and classify objects using the sparse point clouds from planar LIDAR. 2) Though a single RGB image has rich texture and semantic information, it does not give direct information about the depth of the surrounding environment. 3) How to fuse the sparse depth information from planar LIDAR with dense RGB semantic features from camera is still an unsolved problem. In this paper, we complete a sparse depth map from planar LIDAR by the guidance of semantic and texture information from the RGB image. To achieve this goal, we improve the traditional fusion architecture by introducing an inductive late-fusion technique which achieves a more accurate depth map with more detailed and smoother features, as shown in Figure 1.

We organize this paper as follows: Section II briefly reviews the prior work in depth prediction and depth completion. Section III introduces the proposed network architecture. Section III-A details the proposed inductive late-fusion network. The experimental setup and dataset are explained in Sections IV-B. Finally, Sections IV-D and V discuss the experimental results and give conclusions.

II. RELATED WORK

In robotics perception and especially in autonomous driving, depth estimation is a difficult regression problem. It provides enhanced input for other perception tasks such as vehicle segmentation, tracking and state estimation [2]. The

¹Chen Fu and Chiyu Dong are with the Department of Electrical and Computer Engineering, Carnegie Mellon University, Pittsburgh, PA 15213, USA {cful, chiyud}@andrew.cmu.edu

²Christoph Mertz and John M. Dolan are with the Robotics Institute, Carnegie Mellon University, Pittsburgh, PA 15213, USA {mertz, jmd}@cs.cmu.edu

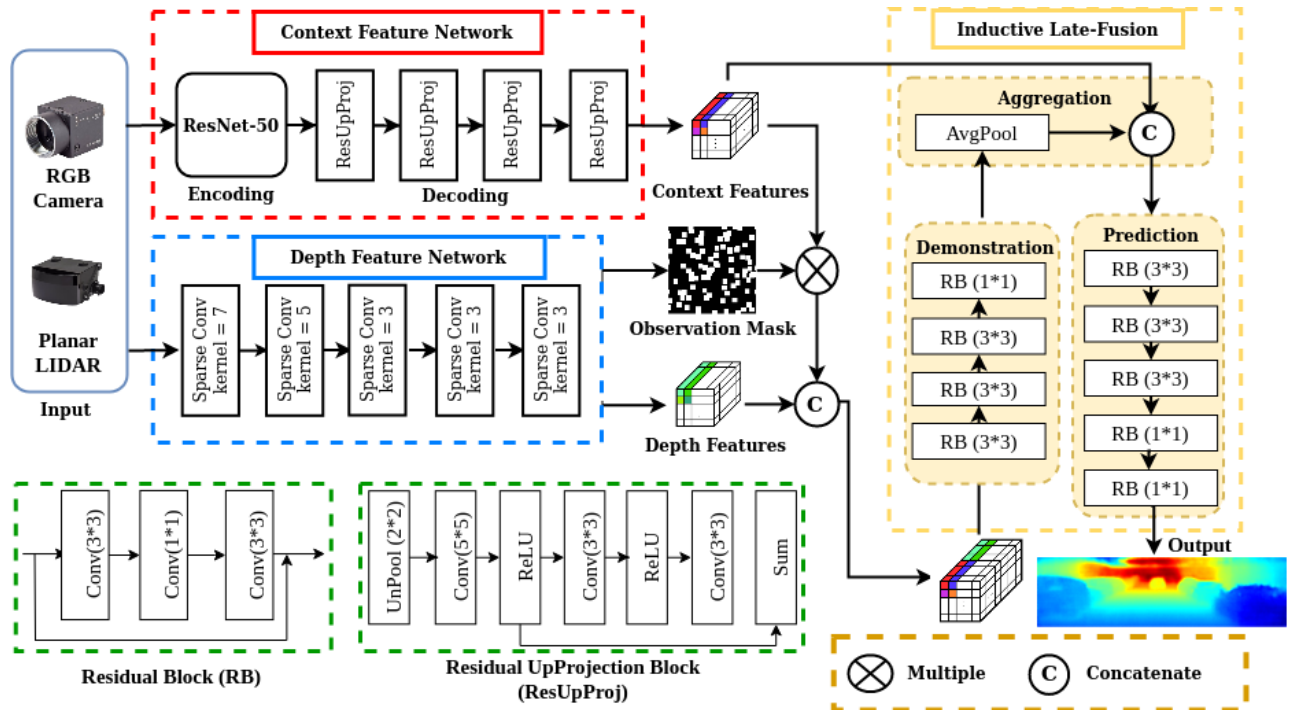


Fig. 2: Detailed architecture of proposed inductive late-fusion network. The red and blue blocks indicate the context feature extraction network and depth feature extraction network. The corresponding Residual Block (RB) and Residual UpProjection Block (ReUpProj) are shown in the green blocks. The proposed inductive late-fusion block is shown in the yellow blocks.

depth estimation problem can be further divided into different categories according to different sensor inputs. It includes depth in-painting for RGB-D camera and depth completion for planar LIDAR [3], [4].

In the dense completion problem, a low-resolution or sparse depth image is completed or super-resolved into a high-quality pixel-wise depth image. In this case, more depth points are augmented for the surrounding scenes and objects, given an RGB image and sparse LIDAR point cloud [5]. This depth augmentation algorithm enhances the accuracy of the vehicle tracking task. However, this approach only considers the dense completion task as interpolation, which does not predict a pixel-wise dense depth image. To further improve the prediction, [6] adapts a novel data term into the MRF and improves the resolution of the depth image qualitatively and quantitatively. As an improvement [7] considers the relationship between the boundaries of image segmentation and depth image and achieves a better depth prediction for aggressively down-sampled RGB images. In order to recover dense depth information from cropped input, [8] introduced a novel super-resolved architecture, the Sparsity Invariant Convolution network. However, this method is insufficient for the perception system of autonomous vehicles, as multi-modal sensing data need to be fused together. To better fuse the input from the RGB image and sparse LIDAR, [3], [9] proposed a self-supervised training pipeline that takes the sequential information of the RGB image as a geometry constraint of the optimization. However, the performance of this architecture is highly reliant on the accuracy of the transition relation between nearby frames, which can be influenced by moving objects within the scene. In contrast, [10], [11] fuses the LIDAR and RGB image through deep CNN. With the

guidance of the surface normal and intention map, a more accurate depth image is predicted. [12] even proposed an architecture which handles potential sensor failures in real autonomous driving cases. However, these depth completion methods [10], [11], [13] take dense LIDAR instead of planar LIDAR as input to predict a pixel-wise depth image. Instead of simply concatenating different feature modulations via an end-to-end network architecture, we propose a inductive late-fusion block which fuses the planar LIDAR depth information with context information. This architecture formulates the fusion problem as a conditional distribution. With the prior knowledge of the observed LIDAR depth feature and RGB context feature, the underlying distribution of the depth is learnt by an induction framework.

III. METHOD

In this section, we explain in detail how to tackle this sensor fusion problem. Compared to state-of-the-art methods which fuse LIDAR depth images with RGB images [10], [11], [13], fusing sparse (planar) LIDAR and RGB images is quite difficult. The challenge is to combine multi-modal features, where the RGB features are dense but the LIDAR features are sparse. Instead of fusing these two streams of information by an end-to-end deep network, we propose a novel architecture which fuses the depth and RGB features through induction. We combine a depth pathway into our previously proposed depth completion network with inductive fusion units adapted from the conditional neural process technique [14]. This fusion architecture fits the planar LIDAR and RGB image fusion scenario even with an extremely low sparsity. Our proposed network has three parts: depth feature extraction network, context feature ex-

traction network and the inductive fusion block. The detailed architecture is discussed in the following sections.




Fig. 3: Assuming this is an image plane with a dimension of $4 * 4$, the figure shows an example of the observation set and target set shown in purple and orange respectively. The black dots represent the pixels where we have depth features.

A. Inductive Late-level Fusion

Due to the sparsity of the input depth image, we cannot directly achieve dense depth features from the network. In this case, instead of directly fusing the sparse depth features with dense context features, we first formulate this fusion problem by a probability model. We currently have two sets of input, the observation set ($N = n$) and target set ($N = m$), where $n + m$ equals the total number of pixels in the image. In the observation set, each pixel location has a context feature vector and a depth feature vector associated with it. However, in the target set, we only have the context features. We define the observation set as $O = \{(x_i, y_i)\}_{i=1}^n$, where x_i represents the context feature vector and y_i represents the depth feature generated by the feature extraction network. The target set is $T = \{x_i\}_{i=n+1}^{n+m}$. An example observation set and target set is shown in Fig. 3. We now define a mapping function $f : X \rightarrow Y$ and let P_θ be a probability distribution over the function f . For $f \sim P_\theta$ we set $\hat{y} = f(x_i)$, where \hat{y} is the ground truth depth for each pixel. So P_θ defines a joint probability distribution over $\{x_i\}_{i=n+1}^{n+m}$ and the depth completion problem is formulated as a Gaussian Process (GP) as follows:

$$P_\theta(f(T)|O, T) \quad (1)$$

To solve this GP problem, an appropriate prior knowledge is required. However, it is difficult to model the complex correlation structure between depth value and RGB image using a simple kernel function. Moreover, the GP is computationally expensive, as the scale depends on the input and output dimension, which is large in the depth completion problem. In this case, we created a novel adaptation of the Conditional Neural Process (CNP) [14] architecture that addresses the above limitations.

In equation 1, θ represents all parameters of the inductive fusion block, which defines the distribution P . In the depth completion task, the pixels in each image patch do not have the problem of ordering. In other words, the depth value of a particular pixel is highly dependent on the surrounding pixels but not the order of the pixels. As a result, the prediction is permutation-invariant and has the following property, which

fits the assumption of CNPs:

$$P_\theta(f(T)|O, T) = P_\theta(f(T')|O, T') = P_\theta(f(T)|O', T) \quad (2)$$

where O' and T' are permutations of O and T . So, we can apply the following equations from CNPs:

$$r_i = h_\theta(x_i, y_i), \quad \forall (x_i, y_i) \in O \quad (3)$$

$$r_j = \frac{1}{T} \sum_{i=1}^T r_i \quad (4)$$

$$\phi_i = g_\theta(x_i, r), \quad \forall (x_i) \in T \quad (5)$$

So, the inductive fusion block contains three subnets, the demonstration network, aggregation network and prediction network. For the demonstration network, as shown in 3, it fuses the depth feature and context features and encodes the information from two feature extraction networks for each pixel location and achieves a higher-level description of the prior knowledge of the regression task. Instead of directly passing the higher-level descriptors to the decoder network, which is also the prediction network, we aggregate them into an intermediate feature representation r_j for each image patch. By virtue of the aggregation network, the features are invariant to translation and rotation, which results in a robust network. In the prediction network, target depth values are estimated based on aggregated feature and context feature.

B. Depth Feature Extraction Network

To better extract the features from sparse depth image input, we apply the Sparsity Invariant CNN (SCNN) [8] to generate the feature representation of the depth. Different from the vanilla convolution layer, SCNN introduces a sparse convolution layer which propagates the features as well as an observation mask along the forward path. The sparse depth value is diluted to the nearby pixels and results in blurred predictions. During this process, the observation mask is also expanded by the max-pooling layer. To generate a depth feature for each pixel, we delete the last layer from the network in the proposed model. The detailed network architecture is shown in the blue blocks in Fig. 2.

C. Context Feature Extraction Network

For the context feature extraction network, we follow our previously proposed depth prediction architecture described in [4]. In this architecture, we concatenate the popular Resnet model-proposed Residual Up-Projection blocks shown in the green blocks of Fig. 2. This model takes advantage of residual learning techniques and captures more detailed features from the input RGB image features [15]. In the depth completion task, we find that the detailed features within the depth images show similarity to the low-level features in the RGB images. For example, the RGB image and depth image have similar edges and corners. So, low-level features from the shallow network contain useful features for the depth prediction and reconstruction in higher layers. To further improve the prediction results, we add three skip connections to pass low-level features to the residual Up-Projection blocks. Instead of using the whole network as

a context feature extraction network, we get rid of the top layers and the network outputs a context feature tensor.

IV. RESULTS AND DISCUSSION

In this section, we concentrate on explaining our experimental setup and results. To verify the performance of the proposed architecture and compare to the state of the art, we evaluate our model on two public datasets: the indoor scenes NYUdepthV2 dataset and the KITTI depth completion dataset, which focuses on real on-road scenes.

A. Network Implementation and Training Strategy

For the demonstration network, we concatenate 4 Residual Blocks with kernel size of $3 * 3$. In the prediction network, we concatenate 5 Residual Blocks and finally predict the depth map. The detailed architecture is shown in Fig. 2. To train the network, we start from the ImageNet pretrained Resnet model and fine-tune the context feature network by the NYUdepthV2 and KITTI dataset. For the depth feature extractor, we pretrain the network by the NYUdepthV2 and KITTI dataset as well. Finally, we combine the depth feature network with the context feature network through the inductive fusion unit. In order to merely compare the network architecture with previous works, we use the L-1 norm as our loss function throughout the training procedure, both the pre-training and whole network architecture. The whole network architecture is trained by the SGD optimizer [24], with the default parameter setting from the Pytorch optimization package [25].

B. Experiment Setup and Evaluation Metric

Dataset: We conduct the experiments on several benchmark datasets introduced by state-of-the-art methods, including NYUdepthv2, KITTI Odometry and the reduced-resolution KITTI dataset [4], [9]. We apply the same experiment setup as in our previous paper [4], with data augmentation including scale, flip, rotation and color translation [9]. We apply a uniform sampling strategy to generate different input densities for evaluation.

Evaluation Metrics: In order to compare the proposed method with benchmark results of the state-of-the-art methods, we apply the following standard evaluation metrics. We not only directly measure the average error over all pixels by root mean square error (RMSE), but also compare the Mean Absolute Relative Error (REL), which prevents the scaling problem, as follows:

$$e_{rel} = \frac{1}{N} \sum_i \frac{|\hat{y}_i - y_i|}{y_i} \quad (6)$$

In order to count the percentage of pixels within a certain threshold, we also consider the δ_j metric, defined as follows:

$$\delta_j = \frac{\text{count}(\{y_i : \max(\frac{\hat{y}_i}{y_i}, \frac{y_i}{\hat{y}_i}) \leq 1.25^j\})}{\text{count}(y_i)} \quad (7)$$

where \hat{y}_i and y_i are the prediction and ground truth depth for each pixel. As a result, the method performs better if it achieves a low RMSE and REL or a higher δ_j .




Fig. 4: Visualization result on NYUdepthV2 dataset. The images from left to right are: RGB image, prediction results from [4], prediction results by proposed method and ground truth. We are using 200 depth points in this experiment setting.

C. Ablation Study

In order to study the function of the proposed late-fusion block and prove its advantage, we also provide an ablation study of different fusion architectures and their performance over different input densities. We compare our late-fusion block with a simple late-fusion network that uses 9 residual blocks as the fusion network. We also change the size of the depth feature network by varying the number of sparse convolution layers (SCN). In this experiment, we compare the depth feature network with 3 SCN and 5 SCN. The result is shown in Fig. 5.

In comparison, the context feature network only achieves a RMSE of 0.46m with a single RGB image as input. In this case, the fusion method has much better performance than the single RGB network. Even with an early-fusion input as in [4], the context feature network only achieves a RMSE of 0.20m on the NYUdepthV2 dataset. As a result, the proposed late-fusion architecture benefits the performance of the depth completion task. As we can see in Fig. 5, as we increase the number of input depth samples, we achieve better performance and the proposed inductive late-fusion block outperforms the traditional late-fusion network concatenated by convolution layers. As the prediction network is estimating the depth value based on the demonstration and observed context features, it gives smoother prediction at image pixels which have similar RGB patterns and geometry constraints. In other words, pixels with similar context should have a similar depth value. Moreover, pixels that are physically nearby but have distinct context features should have different depth values. As a result, we achieve a smoother prediction on flat surfaces and sharper edges on contours of objects. Instead, traditional network architectures are merely learning the distribution of depth based on RGB image and depth samples, which results in a blurred prediction on both surfaces and edges.

TABLE I: Performance comparison of proposed network with previous single RGB-based and fusion methods on NYUdepthV2 dataset with 200 depth points

Input	#Sample	Methods	RMSE (m)	REL	δ_1 (%)	δ_2 (%)	δ_3 (%)
RGB	0	Roy et al. [16]	0.744	0.187	-	-	-
	0	Eigen et al. [17]	0.641	0.158	76.9	95.0	98.8
	0	Laina et al. [18]	0.573	0.127	81.1	95.3	98.8
RGB+D	225	Liao et al. [19]	0.442	0.104	87.8	96.4	98.9
	200	Ma et al. [9]	0.230	0.044	97.1	99.4	99.8
	200	Lee et al. [20]	0.225	0.046	97.2	-	-
	200	Fu et al. [4]	0.203	0.040	97.6	99.5	99.9
	200	Abdelrahman et al. [21]	0.192	0.030	97.9	99.5	99.8
	200	Proposed	0.169	0.028	98.4	99.9	99.9

TABLE II: Performance comparison of proposed network with previous single RGB-based and fusion methods on KITTI Odometry dataset

Input	#Sample	Methods	RMSE (m)	REL	δ_1 (%)	δ_2 (%)	δ_3 (%)
RGB	0	Laina et al. [16]	8.73	0.280	60.1	82.0	92.6
	0	Mancini et al. [22]	7.51	-	31.8	61.7	81.3
	0	Eigen et al. [17]	6.16	0.190	69.2	89.9	96.7
RGB+D	200	Liao et al. [19]	4.50	0.113	87.4	96.0	98.4
	200	Ma et al. [9]	3.85	0.083	91.9	97.0	98.9
	200	Fu et al. [4]	3.67	0.072	92.3	97.3	98.9
	200	Proposed	3.11	0.058	93.9	97.6	98.9
RGB+D	500	Ma et al [9]	3.38	0.073	93.5	97.6	98.9
	500	Cheng et al. [23]	2.98	0.044	95.7	98.0	99.1
	500	Proposed	2.84	0.045	95.3	98.1	99.2

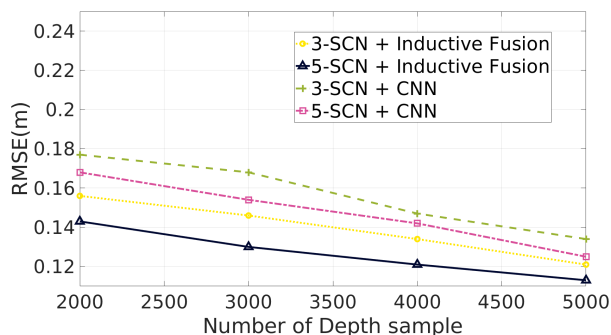


Fig. 5: Ablation study of different network architectures and performance with various input 3D densities. In this figure, the proposed late-fusion block is compared with vanilla convolution networks. We also study how different size of depth feature network influences the performance of depth estimation.

D. Results and Discussion on Benchmark Dataset

In this section, we mainly discuss the experimental results on the benchmark dataset, NYUDepthV2 and KITTI dataset. We take the reported accuracy of state-of-the-art depth completion and depth prediction methods from their original papers.

NYUDepthV2 Dataset: In general, taking multi-modal sensor data as input improves the estimation result compared to using a single sensor. In methods [9], [19], these architectures apply the end-to-end network, which concatenates multiple residual blocks and deep decoder networks. However, this family of methods loses detailed features and context information through the encoding and decoding architecture. On the NYUDepthV2 dataset, we achieve a 12.5% improvement with 200 dense inputs in terms of

RMSE. The detailed comparison is shown in TABLE I. By applying our proposed architecture, we achieve an 11.9% better performance compared with the baseline method [21] on the NYUDepthV2 dataset. In Fig. 4, we show the visualization result on the NYUDepthV2 dataset. Comparing with ground truth, the prediction results of the proposed method have much clearer context details. In all, combining REL and RMSE in TABLE I and Fig. 4, we achieve a better estimation on the boundaries as well as flat surfaces in the scenes.

KITTI Odometry Dataset and Reduced-Resolution

KITTI Dataset: In TABLE II, we notice that fusion of sparse LIDAR input with dense RGB image has better depth completion results. In comparing with our previous best depth completion result using 200 depth samples as input, we achieve a 14.7% improvement in terms of RMSE. The visualization result is shown in Fig. 6. We also conducted a performance comparison on the reduced-resolution KITTI dataset and compared with a few state-of-the-art single RGB images with LIDAR fusion algorithms. We take 200 depth points to simulate the density of planar LIDAR in the band. As we see from TABLE III, the proposed method performs better than previous methods by around 20% and 25% regarding RMSE and REL, respectively.

TABLE III: Performance comparison of proposed network on reduced-resolution KITTI odometry dataset. We take a band of the LIDAR point cloud and conduct depth completion within the band. We down-sample the dense point cloud by randomly selecting 200 LIDAR points as input.

Methods	RMSE (m)	REL	δ_1 (%)	δ_2 (%)	δ_3 (%)
Ma et al. [9]	0.42	0.040	98.3	99.6	99.8
Fu et al. [4]	0.35	0.032	98.5	99.6	99.8
Proposed	0.28	0.023	98.9	99.7	100.0

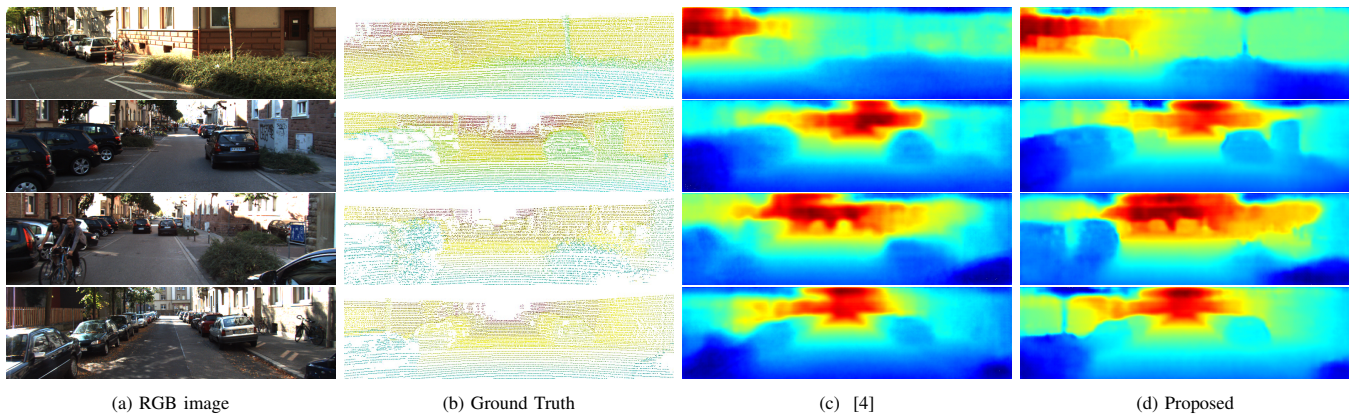


Fig. 6: Visualization result on KITTI dataset on official validation sets. The images from left to right are: RGB image, ground truth dense depth (brightness, contrast changed for visual enhancement), predicted dense depth image from previous work [4] and prediction dense depth image by proposed method. We are using 200 depth points in this experiment setting.

V. CONCLUSION

In this paper, we propose an inductive fusion architecture which fuses planar LIDAR with RGB images to complete the depth map of the surrounding environment. The proposed inductive fusion block optimally combines the dense context features from RGB images with the sparse depth features from planar LIDAR. Our method outperforms conventional methods on the NYUDepthV2 dataset and KITTI odometry dataset. We also achieve promising results on the simulated planar LIDAR dataset. Further work will apply the dense depth image generated by the proposed architecture to perception tasks such as object detection and vehicle heading estimation. We will also apply the proposed method to real planar LIDAR depth completion tasks.

VI. ACKNOWLEDGMENTS

This work was supported by the CMU Argo AI Center for Autonomous Vehicle Research.

REFERENCES

- [1] Y. Qian, M. Yang, X. Zhao, C. Wang, and B. Wang, "Oriented spatial transformer network for pedestrian detection using fish-eye camera," *IEEE Transactions on Multimedia*, vol. PP, pp. 1–1, 07 2019.
- [2] C. Fu, P. Hu, C. Dong, C. Mertz, and J. M. Dolan, "Camera-based semantic enhanced vehicle segmentation for planar lidar," in *IEEE Intelligent Transportation Systems Conference (ITSC)*, 11 2018, pp. 3805–3810.
- [3] F. Ma, G. V. Cavalheiro, and S. Karaman, "Self-supervised sparse-to-dense: Self-supervised depth completion from lidar and monocular camera," *arXiv preprint arXiv:1807.00275*, 2018.
- [4] C. Fu, C. Mertz, and J. M. Dolan, "Lidar and monocular camera fusion: On-road depth completion for autonomous driving," in *IEEE Intelligent Transportation Systems Conference (ITSC)*. IEEE, 2019, pp. 273–278.
- [5] D. Held, J. Levinson, and S. Thrun, "Precision tracking with sparse 3d and dense color 2d data," in *International Conference on Robotics and Automation (ICRA)*, May 2013.
- [6] J. Lu, D. Min, R. S. Pahwa, and M. N. Do, "A revisit to MRF-based depth map super-resolution and enhancement," in *ICASSP, IEEE International Conference on Acoustics, Speech and Signal Processing - Proceedings*, 2011, pp. 985–988.
- [7] J. Lu and D. Forsyth, "Sparse depth super resolution," in *Proceedings of the IEEE Computer Society Conference on Computer Vision and Pattern Recognition*, vol. 07-12-June-2015, 2015, pp. 2245–2253.
- [8] J. Uhrig, N. Schneider, L. Schneider, U. Franke, T. Brox, and A. Geiger, "Sparsity invariant cnns," in *International Conference on 3D Vision (3DV)*, 2017.
- [9] F. Ma and S. Karaman, "Sparse-to-dense: Depth prediction from sparse depth samples and a single image," 2018.
- [10] 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," in *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2019.
- [11] Y. Chen, B. Yang, M. Liang, and R. Urtasun, "Learning joint 2d-3d representations for depth completion," in *The IEEE International Conference on Computer Vision (ICCV)*, October 2019.
- [12] J. Maximilian, d. C. Raoul, W. Emilie, P. Xavier, and N. Fawzi, "Sparse and Dense Data with CNNs: Depth Completion and Semantic Segmentation," 2018.
- [13] J. Tang, F.-P. Tian, W. Feng, J. Li, and P. Tan, "Learning guided convolutional network for depth completion," 2019.
- [14] M. Garnelo, D. Rosenbaum, C. Maddison, T. Ramalho, D. Saxton, M. Shanahan, Y. W. Teh, D. Rezende, and S. M. A. Eslami, "Conditional neural processes," in *Proceedings of the 35th International Conference on Machine Learning*, vol. 80. PMLR, 10–15 Jul 2018, pp. 1704–1713.
- [15] K. He, X. Zhang, S. Ren, and J. Sun, "Deep Residual Learning for Image Recognition," *Arxiv.Org*, vol. 7, no. 3, pp. 171–180, 2015.
- [16] A. Roy and S. Todorovic, "Monocular depth estimation using neural regression forest," *2016 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 5506–5514, 2016.
- [17] E. David and F. Rob, "Predicting depth, surface normals and semantic labels with a common multi-scale convolutional architecture," 2014.
- [18] I. Laina, C. Rupprecht, V. Belagiannis, F. Tombari, and N. Navab, "Deeper depth prediction with fully convolutional residual networks," in *3D Vision (3DV), 2016 Fourth International Conference on*. IEEE, 2016, pp. 239–248.
- [19] Y. Liao, L. Huang, Y. Wang, S. Kodagoda, Y. Yu, and Y. Liu, "Parse geometry from a line: Monocular depth estimation with partial laser observation," 05 2017, pp. 5059–5066.
- [20] B.-U. Lee, H.-G. Jeon, S. Im, and I. S. Kweon, "Depth completion with deep geometry and context guidance," in *Proceedings of IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, May 2019.
- [21] A. Eldesokey, M. Felsberg, and F. S. Khan, "Propagating confidences through cnns for sparse data regression," *arXiv preprint arXiv:1805.11913*, 2018.
- [22] M. Mancini, G. Costante, P. Valigi, and T. Ciarfuglia, "Fast robust monocular depth estimation for obstacle detection with fully convolutional networks," 07 2016.
- [23] X. Cheng, P. Wang, and R. Yang, "Depth estimation via affinity learned with convolutional spatial propagation network," in *Proceedings of the European Conference on Computer Vision (ECCV)*, 2018, pp. 103–119.
- [24] D. Kingma and J. Ba, "Adam: A method for stochastic optimization," *International Conference on Learning Representations*, 12 2014.
- [25] A. Paszke, S. Gross, S. Chintala, G. Chanan, E. Yang, Z. DeVito, Z. Lin, A. Desmaison, L. Antiga, and A. Lerer, "Automatic differentiation in pytorch," 2017.