

MLOD: Awareness of Extrinsic Perturbation in Multi-LiDAR 3D Object Detection for Autonomous Driving

Jianhao Jiao*, Peng Yun*, Lei Tai, Ming Liu

Abstract—Extrinsic perturbation always exists in multiple sensors. In this paper, we focus on the extrinsic uncertainty in multi-LiDAR systems for 3D object detection. We first analyze the influence of extrinsic perturbation on geometric tasks with two basic examples. To minimize the detrimental effect of extrinsic perturbation, we propagate an uncertainty prior on each point of input point clouds, and use this information to boost an approach for 3D geometric tasks. Then we extend our findings to propose a multi-LiDAR 3D object detector called *MLOD*. *MLOD* is a two-stage network where the multi-LiDAR information is fused through various schemes in stage one, and the extrinsic perturbation is handled in stage two. We conduct extensive experiments on a real-world dataset, and demonstrate both the accuracy and robustness improvement of *MLOD*. The code, data and supplementary materials are available at: <https://ram-lab.com/file/site/mlod>.

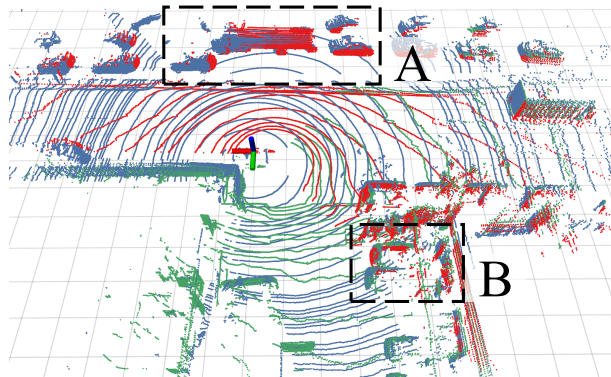
I. INTRODUCTION

3D object detection is a fundamental module in robotic systems. As the front-end of a system, it enables vehicles to recognize key objects such as cars and pedestrians, which is indispensable for high-level decision making. Compared with camera-based detectors, the LiDAR-based detectors perform better in several challenging scenarios because of their activeness as sensors and distance measurement ability for surroundings. However, LiDARs commonly suffer from data sparsity and a limited vertical field of view [1]. For instance, LiDARs' points distribute loosely, which induces a mass of empty regions between two nearby scans. In this paper, we consider a multi-LiDAR system, a setup with excellent potential to solve 3D object detection. Compared with single-LiDAR setups, multi-LiDAR systems enable a vehicle to maximize its perceptual awareness of environments and obtain sufficient measurements. The decreasing price of LiDARs also makes such systems accessible to many modern self-driving cars [1]–[3].

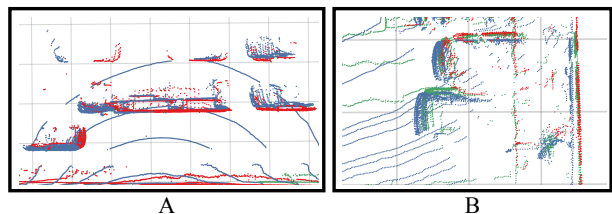
However, two challenges affect the development of multi-LiDAR object detection. One difficulty is multi-LiDAR fusion. As described in [4], current data fusion methods all have pros and cons. It requires us to perform extensive efforts on possible fusion schemes in real tests for an excellent detection algorithm. Another issue is that extrinsic perturbation

This work was supported by the National Natural Science Foundation of China, under grant No. U1713211, Collaborative Research Fund by Research Grants Council Hong Kong, under Project C4063-18G, and the Research Grant Council of Hong Kong SAR Government, China, under Project No. 11210017, awarded to Prof. Ming Liu.

* Equal contribution (alphabet order). Jianhao Jiao, Peng Yun, Ming Liu are with the Robotics Institute, Intelligent Autonomous Driving Center, The Hong Kong University of Science and Technology, Hong Kong SAR, China. {jjiao, pyun, eelium}@ust.hk. Lei Tai is with IAS BU, Huawei Technologies. {tailei2}@huawei.com.



(a) Multiple point clouds are merged with little extrinsic perturbation.



(b) The details in the region A and B of the point cloud.

Fig. 1. (a) A point cloud example from the LYFT dataset [3]. It is merged by transforming point clouds perceived by top (blue), front-left (green), and front-right (red) LiDARs into the base frame. We simulate perturbation on the rotation of 1° in yaw. (b) The details in the region A and B with a zooming effect. Massive noisy points are induced by extrinsic perturbation, while they should be sharp or flat in perturbation-free situations.

is inevitable for sensor fusion during long-term operation because of factors such as vibration, temperature drift, and calibration error [5]. Especially, wide baseline stereo cameras or vehicle-mounted multi-LiDAR systems suffer even more extrinsic deviations than the normal one. Even though online calibration has been proposed to handle this issue [6], life-long calibration is always challenging [7]. These methods typically require environmental or motion constraints with full observability. Otherwise, the resulting extrinsics may become suboptimal or unreliable. As identified in both Fig. 1 and Section IV-B, extrinsic perturbation is detrimental to the measurement accuracy even with a small change. But this fact is often neglected by the research community.

To tackle the challenges, we propose a **Multi-LiDAR Object Detector (*MLOD*)** to predict states of objects on point clouds. We design a two-stage procedure to estimate 3D bounding boxes, and demonstrate two innovations in *MLOD*. First, we explore three fusion schemes to exploit multi-view point clouds in a stage-1 network for generating object proposals. These schemes perform LiDAR fusion at the input, feature, and output phases, respectively. Second,

we develop a stage-2 network to handle extrinsic perturbation and refine the proposals. The experimental results demonstrate that *MLOD* outperforms single-LiDAR detectors up to 9.7AP in perturbation-free cases. When considering extrinsic perturbation, *MLOD* consistently improves the performance of its stage-1 counterpart. To the best of our knowledge, this is the first work to systematically study the multi-LiDAR object detection with consideration of extrinsic perturbation. Overall, our work has the following contributions:

- 1) We analyze the influence of extrinsic perturbation on multi-LiDAR-based geometric tasks, and also demonstrate that the usage of input uncertainty prior improves the robustness of an approach against such effect.
- 2) We propose a unified and effective two-stage approach for multi-LiDAR 3D object detection, where the multi-LiDAR information is fused in the first stage, and extrinsic perturbation is handled in the second stage.
- 3) We exhaustively evaluate the proposed approach in terms of accuracy and robustness under extrinsic perturbation on a real-world self-driving dataset.

II. RELATED WORK

We briefly review methods on LiDAR-based object detection and uncertainty estimation of deep neural networks.

A. 3D Object Detection on Point Clouds

The LiDAR-based object detectors are generally categorized into grid-based [8]–[10] and point-based methods [11]–[13]. 3D-FCN [8] implemented 3D volumetric CNN on voxelized point clouds. But it commonly suffers a high computation cost due to dense convolution on sparse point clouds. VoxelNet [9] was proposed as an end-to-end network to learn features. However, operation on grids is inefficient since LiDAR’s points are sparse. To cope with this drawback, SECOND [10] utilized the spatially sparse convolution [14] to replace the 3D dense convolution layers. In this paper, we adopt SECOND as our basic proposal generator and propose three schemes for multi LiDAR fusion in the first stage. Qi et al. proposed PointNet [11] to directly learn features from raw data with a symmetric function. And their follow-up work presented a set-abstraction block to capture local features [12]. PointRCNN [13] exploits PointNet to learn point-wise features and segments foreground for autonomous driving. In this paper, we extend PointNet to design the stage-2 network of *MLOD* with an awareness of extrinsic perturbation.

Researchers also explored fusion methods that further exploit images to improve the LiDAR-based approaches. For example, MV3D [4] combined features from multiple views, including a bird-eye view and front view, to conduct classification and regression. Several methods [15], [16] separate the process of detection into two stages. They first projected region proposals generated by the camera-based detectors into 3D space, and then used PointNets to segment and to classify objects on point clouds. But these methods are aimed at data fusion from cameras and LiDARs and assume sensors to be well-calibrated. In contrast, our work focuses

TABLE I
NOMENCLATURE

Notation	Explanation
$\{\}^b, \{\}^{l^i}$	Frame of the base and the i^{th} LiDAR.
\mathcal{P}^{l^i}	Raw point cloud captured by the i^{th} LiDAR.
\mathcal{B}	Set of estimated 3D bounding boxes.
\mathbf{T}	Transformation matrix in the <i>Lie group</i> $SE(3)$.
\mathbf{R}	Rotation matrix in the <i>Lie group</i> $SO(3)$.
\mathbf{t}	Translation vector in \mathbb{R}^3 .
Θ	Measurement and extrinsic uncertainty prior.
Ξ	Associated covariance of each point.
α	Scaling parameter of Θ .

on multi-LiDAR fusion, and tries to minimize the negative effect of extrinsic perturbation for 3D object detection.

B. Uncertainty Estimation in Object Detection

The problem of uncertainty estimation is essential to the reliability of an algorithm, and has attracted much attention in recent years. As the two main types of uncertainties in deep neural networks: aleatoric and epistemic uncertainty, they are explained in [17], [18]. In [18], the authors also demonstrated the benefits of modeling uncertainties for vision tasks. As an extension, Feng et al. [19] proposed an approach to capture uncertainties for 3D object detection. Generally, the data noise is modeled as a unique Gaussian variable in these works. However, the sources of data uncertainties in multi-LiDAR systems are much more complicated. In this paper, we take both the measurement noise and extrinsic perturbation into account, and analyze their effect on multi-LiDAR-based object detection. Furthermore, we propagate the Gaussian uncertainty prior of both the extrinsics and measurement to model the input data uncertainty. This additional cue is utilized to improve the robustness of *MLOD* against extrinsic perturbation.

III. NOTATIONS AND PROBLEM STATEMENT

The nomenclature is shown in Tab. I. The transformation from $\{\}^b$ to $\{\}^{l^i}$ is denoted by $\mathbf{T}_{l^i}^b$. Our perception system consists of one primary LiDAR which is denoted by l^1 and multiple auxiliary LiDARs. The primary LiDAR defines the base frame and the auxiliary LiDAR provides an additional field of view (FOV) and measurements to alleviate the occlusion problem and the sparsity drawback of the primary LiDAR. Extrinsics describe the relative transformation from the base frame to frames of auxiliary LiDARs. With the extrinsics, all measurements or features from different LiDARs are transformed into the base frame for data fusion. LiDARs are assumed to be synchronized that multiple point clouds are perceived at the same time.

In this paper, we focus on 3D object detection with a multi-LiDAR system. The extrinsic perturbation is also considered, which indicates the small but unexpected change on transformation from the base frame to other frames over time. Our goal is to estimate a series of 3D bounding boxes covering objects of predefined classes. Each bounding box $\mathbf{b} \in \mathcal{B}$ is parameterized as $[c, x, y, z, w, l, h, \gamma]$, c is the class

of a bounding box, $[x, y, z]$ denotes a box's bottom center, $[w, h, l]$ represent the sizes along the x -, y -, and z - axes respectively, as well as γ for the rotation of the 3D bounding box along the z - axis in the range of $(0, \pi]$.

IV. PROPAGATING EXTRINSIC UNCERTAINTY ON POINTS

We begin with providing preliminaries about the uncertainty representation and propagation. We then use an example to demonstrate the negative effect on points caused by extrinsic perturbation. Finally, we also conduct a plane fitting experiment to show that the extrinsic covariance prior can be utilized to make a fitting approach robust.

A. Preliminaries

We employ the method in [20] to represent the uncertainty. For convenience, rotation and translation are used to indicate a transformation. We first define a random variable for \mathbb{R}^3 with small perturbation according to

$$\mathbf{t} = \boldsymbol{\rho} + \bar{\mathbf{t}}, \quad \boldsymbol{\rho} \sim \mathcal{N}(\mathbf{0}, \mathbf{P}), \quad (1)$$

where $\bar{\mathbf{t}}$ is a noise-free translation and $\boldsymbol{\rho} \in \mathbb{R}^3$ is a zero-mean Gaussian variable with covariance \mathbf{P} . We can also define a rotation for $SO(3)$ as¹

$$\mathbf{R} = \exp(\phi^\wedge) \bar{\mathbf{R}}, \quad \phi \sim \mathcal{N}(\mathbf{0}, \Phi), \quad (2)$$

where $\bar{\mathbf{R}}$ is a noise-free rotation and $\phi \in \mathbb{R}^3$ is a small zero-mean Gaussian variable with the covariance Φ . With (2), we can represent a noisy transformation by storing the mean as $[\bar{\mathbf{t}}, \bar{\mathbf{R}}]$ and using $[\boldsymbol{\rho}, \phi]$ for perturbation on the vector space. Similarly, a point with the perturbation in \mathbb{R}^3 is written as

$$\mathbf{p} = \boldsymbol{\zeta} + \bar{\mathbf{p}}, \quad \boldsymbol{\zeta} \sim \mathcal{N}(\mathbf{0}, \mathbf{Z}), \quad (3)$$

where $\boldsymbol{\zeta}$ is zero-mean Gaussian with the covariance \mathbf{Z} . With the above representations, we can pass the Gaussian representation of a point through a noisy rotation and translation to produce a mean and covariance for its new measurement. Here, we can use $[\bar{\mathbf{t}}_{l_i}^b, \bar{\mathbf{R}}_{l_i}^b]$ to indicate the ground-truth extrinsics of a multi-LiDAR system, and $[\boldsymbol{\rho}, \phi, \boldsymbol{\zeta}]$ to indicate both the extrinsic and measurement perturbation. By transforming $\mathbf{p} \in \mathcal{P}^i$ into $\{\}^b$, we have

$$\begin{aligned} \mathbf{y} &\triangleq \bar{\mathbf{R}}_{l_i}^b \mathbf{p} + \bar{\mathbf{t}}_{l_i}^b = \exp(\phi^\wedge) \bar{\mathbf{R}}_{l_i}^b (\boldsymbol{\zeta} + \bar{\mathbf{p}}) + (\boldsymbol{\rho} + \bar{\mathbf{t}}_{l_i}^b) \\ &\approx (\mathbf{I} + \phi^\wedge) \bar{\mathbf{R}}_{l_i}^b (\boldsymbol{\zeta} + \bar{\mathbf{p}}) + (\boldsymbol{\rho} + \bar{\mathbf{t}}_{l_i}^b), \end{aligned} \quad (4)$$

where we have kept the first-order approximation of the exponential map. If we multiply out the equation and retain only those terms that are first-order in ϕ or $\boldsymbol{\zeta}$, we have

$$\mathbf{y} \approx \mathbf{h} + \mathbf{H}\boldsymbol{\theta}, \quad (5)$$

where

$$\begin{aligned} \boldsymbol{\theta} &= [\boldsymbol{\rho}, \phi, \boldsymbol{\zeta}] \\ \mathbf{h} &= \bar{\mathbf{R}}_{l_i}^b \bar{\mathbf{p}} + \bar{\mathbf{t}} \\ \mathbf{H} &= [\mathbf{I} \quad -(\bar{\mathbf{R}}_{l_i}^b \bar{\mathbf{p}})^\wedge \quad \bar{\mathbf{R}}_{l_i}^b], \end{aligned} \quad (6)$$

¹The \wedge operator turns a 3×1 vector into the corresponding skew symmetric matrix in the *Lie algebra* $\mathfrak{so}(3)$. The exponential map \exp associates an element of $\mathfrak{so}(3)$ to a rotation in $SO(3)$. The closed-form expression for the matrix exponential can be found in [20].

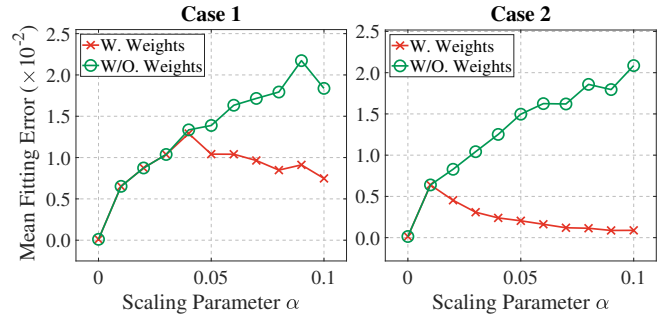


Fig. 2. Mean fitting error $e_{\mathbf{x}} = \|\mathbf{x}_{gt} - \mathbf{x}_{est}\|$ of methods with weights (use uncertainty prior) or without weights on two planar surface cases. The weighted method does better, and the error does not increase along with α .

We embody the uncertainties of extrinsics and measurements into $\boldsymbol{\theta}$ that is subjected to a zero-mean Gaussian with 9×9 covariance $\boldsymbol{\Theta} = \text{diag}(\mathbf{P}, \Phi, \mathbf{Z})$. Linearly transformed by $\boldsymbol{\theta}$, \mathbf{y} is a Gaussian variable with mean and covariance as

$$\begin{aligned} \boldsymbol{\mu} &\triangleq E[\mathbf{y}] = \mathbf{h} \\ \boldsymbol{\Xi} &\triangleq E[(\mathbf{y} - \boldsymbol{\mu})(\mathbf{y} - \boldsymbol{\mu})^\top] = \mathbf{H}\boldsymbol{\Theta}\mathbf{H}^\top. \end{aligned} \quad (7)$$

where we follow [21] to use the trace, i.e., $\text{tr}(\boldsymbol{\Xi})$, as the criterion to quantify the magnitude of a covariance.

B. Uncertainty Propagation Example

In this section, a simple example of uncertainty propagation on a point is presented, where we quantify the data noise caused by both extrinsic and measurement perturbation. The associated covariance is propagated by passing the Gaussian uncertainty from extrinsics and measurements through a transformation. We use the ground-truth extrinsics having rotation with $[10, 10, 10]^\circ$ in roll, pitch, yaw and translation with $[1, 1, 1]m$ along the x -, y -, and z - axis respectively. Let $\mathbf{p} = [10, 10, 10]^\top m$ be a landmark. $\boldsymbol{\Theta}$ is treated as our prior knowledge and considered as a constant matrix²

$$\boldsymbol{\Theta} = \text{diag}^2 \left\{ \frac{1}{20}, \frac{1}{20}, \frac{1}{20}, \frac{1}{10}, \frac{1}{10}, \frac{1}{10}, \frac{1}{50}, \frac{1}{50}, \frac{1}{50} \right\}, \quad (8)$$

where the first six diagonal entries can be multiplied by a scaling parameter α , allowing us to parametrically increase the extrinsic covariance in experiments. According to (7), we have an approximate expression

$$\boldsymbol{\Xi} \approx \alpha \times \begin{bmatrix} 2.36 & -1.24 & -1.20 \\ -1.24 & 2.41 & -1.18 \\ -1.20 & -1.18 & 2.49 \end{bmatrix}, \quad (9)$$

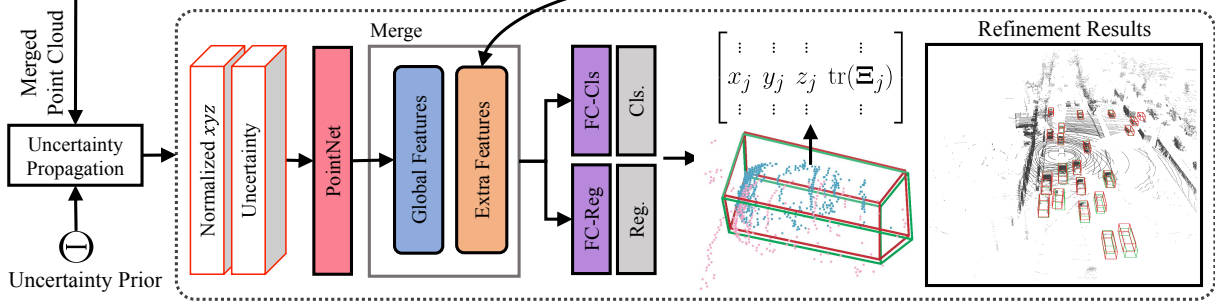
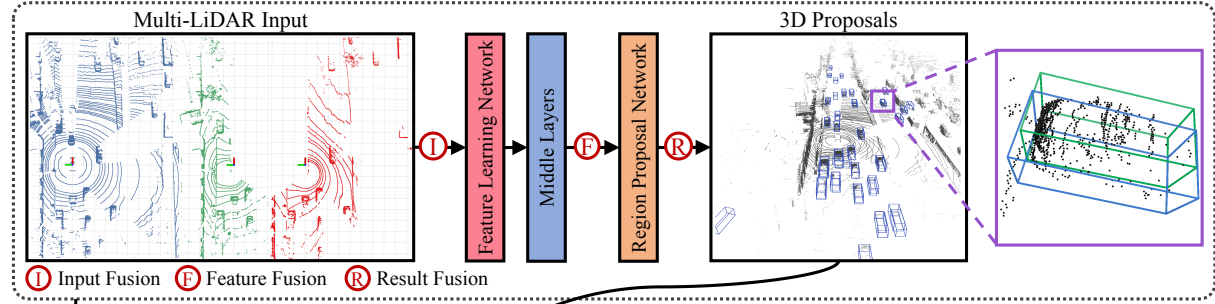
where the new position has a high variance around $0.22m$ on each axis even for a small input covariance (i.e., $\alpha = 0.02$). This accuracy is totally unacceptable for autonomous driving.

C. Plane Fitting Experiment

We denote \mathcal{P}_m a merged point cloud which is obtained by transforming \mathcal{P}^{l_1} and \mathcal{P}^{l_2} into $\{\}^b$ with ground-truth extrinsics. We assume that there is a dominant plane in \mathcal{P}_m , and our task is to estimate the plane coefficients. They can

²The measurement noise is found in LiDAR datasheets. The extrinsic perturbation typically has the variances around $5cm$ in translation and 5° in rotation. This is based on our studies on multi-LiDAR systems [1], [22].

Stage-1: Fusion-Based Proposal Generation (Section V-A)



Stage-2: Box Refinement From Uncertain Points (Section V-C)

Fig. 3. Overview of the proposed detector (*MLOD*) for multi-LiDAR 3D object detection. Red circles: the position where the fusion may be performed; Blue cubes: proposals in stage one; Red cubes: refinements in stage two; Green cubes: ground truths. We used the score and parameter of a proposal from the stage-1 network as the extra features. Global features are extracted by PointNet [11]. They are merged as the input for final classification and regression. For simplicity, we consider three LiDARs as an example; however, *MLOD* is extensible to more.

be obtained by solving a linear system $\mathbf{Ax} = \mathbf{b}$ from \mathcal{P}_m , where row elements of \mathbf{A} are the point coordinates, and \mathbf{b} is set as an identity vector with least-squares methods.

Here, we show how the input uncertainty prior Θ can be utilized to acquire better fitting results. We define \mathbf{W} as a diagonal matrix to weight the linear system. For a point $\mathbf{p}_j \in \mathcal{P}_m$, we propagate its associated covariance Ξ using (7). The corresponding entry (i.e., \mathbf{W}_{jj}) is set as the inverse of $\text{tr}(\Xi)$. After that, the fitting problem is turned into a weighted least-squares regression and the optimal results can be obtained as $\mathbf{x}^* = (\mathbf{A}^T \mathbf{W} \mathbf{A})^{-1} \mathbf{A}^T \mathbf{W} \mathbf{b}$.

An experiment is conducted on a toy example to compare the performance of two fitting methods (i.e., with or without weights). We use the ground-truth extrinsics and the prior covariance in Section IV-B to sample perturbation with 100 trials. We also generate 10000 points according to a planar surface which are subjected to zero-mean Gaussian noise with a standard deviation of $0.02m$ to produce \mathcal{P}_m . \mathcal{P}_m is randomly split into two equal parts to form \mathcal{P}^{l^1} and \mathcal{P}^{l^2} . At each trial, we evaluate the plane fitting results of each method by comparing them to the ground truth as $e_{\mathbf{x}} = \|\mathbf{x}_{gt} - \mathbf{x}_{est}\|$.

The mean fitting error of each method on two different cases over $\alpha \in [0, 0.1]$ is shown in Fig. 2. We see that the weighted least-squares method does better in estimating the plane coefficients, and the mean fitting error does not increase along with α . It shows that the point-wise uncertainty information can be utilized to improve the robustness of algorithms in geometric tasks. But in practice, the value of α should be carefully set by manual or adaptively obtained from an online method. Otherwise, a few correct measurements are discarded during operation.

V. METHODOLOGY

In this section, we extend our findings from the basic geometric tasks to multi-LiDAR-based 3D object detection. The overall structure of the proposed two-stage *MLOD* is illustrated in Fig. 3. We adopt SECOND [10], which is a sparse convolution improvement of VoxelNet [9], to generate 3D proposals in the first stage.³ It consists of three components: a feature learning network (FLN) for feature extraction; middle layers (ML) for feature embedding with sparse convolution; and a region proposal network (RPN) for box prediction and regression. Readers are referred to [9], [10] for more details about the network structures. We first present three species of 3D proposal generation with different multi-LiDAR fusion schemes: *Input Fusion*, *Feature Fusion* and *Result Fusion*. Then, we introduce the architecture of our stage-2 network to tackle the extrinsic uncertainty and refine the proposals.

A. Proposal Generation With Three Fusion Schemes

According to stages at which the information from multiple LiDARs is fused, we propose three general fusion schemes, which we call *Input Fusion*, *Feature Fusion*, and *Result Fusion*. These approaches are developed from SECOND with several modifications.

1) *Input Fusion*: The fusion of point clouds is performed at the input stage. We transform raw point clouds perceived

³We use SECOND V1.5 in our experiments: <https://github.com/traveller59/second.pytorch>. Different from the original SECOND [10], the FLN of SECOND V1.5 is simplified by computing the mean value of points within each voxel, which consumes less memory.

by all LiDARs into the base frame to obtain the fused point cloud, and then feed it to the network as the input.

2) *Feature Fusion*: To enhance the feature interaction, the LiDAR data is also fused in the feature level. The extracted features from the FLN and ML are transformed into the base frame, and then fused by adopting the maximum value as

$$\mathcal{F}_{\text{fused}} = \mathcal{F}^{l^1} \oplus (\mathbf{T}_{l^2}^b \mathcal{F}^{l^2}) \oplus \dots \oplus (\mathbf{T}_{l^I}^b \mathcal{F}^{l^I}), \quad (10)$$

where \mathcal{F}^{l^i} are the extracted features, \oplus denotes the max operator, and I is the number of LiDARs.

3) *Result Fusion*: The result fusion takes the box proposals and the associated points as inputs, and produces a set of boxes with high scores. We transform all box proposals into the base frame, and then filter them as

$$\mathcal{B}_{\text{fused}} = \mathcal{B}^{l^1} \oplus (\mathbf{T}_{l^2}^b \mathcal{B}^{l^2}) \oplus \dots \oplus (\mathbf{T}_{l^I}^b \mathcal{B}^{l^I}), \quad (11)$$

where \oplus denotes the non-maximum suppression (NMS) on 3D intersection-over-Union (IoU). After transformation, each object is associated with several candidate boxes, and these boxes with low confidences are filtered by the NMS.

B. Box Refinement From Uncertain Points

Although the stage-1 fusion-based network generates promising proposals, its capability in handling uncertain data uncertainty is fragile. Therefore, we propose a stage-2 module to improve the robustness of *MLOD*. A straightforward idea is to eliminate the highly uncertain points according to their associated covariances. But this method is sensitive to a pre-set threshold. Inspired by [13], it is more promising that training a neural network to embed features and refine the proposals with an awareness of uncertainty.

We thus use a deep neural network to deal with this problem. The network takes a series of points of each proposal (a $2m$ margin along x -, y -, and z - axes is expanded) generated by the stage-1 network, and refines the stage-1 proposals. In addition to three-dimensional coordinates, each point is also embedded with its uncertain quantity. As defined in (7), we use the trace, i.e., $\text{tr}(\Xi)$, to quantify the uncertainties. We employ PointNet [11] as the backbone to extract global features. We also encode extra features, including the scores and parameters of each proposal, to concatenate with the global features. At the end of the network, fully connected layers are used to provide classification scores and refinement results with two heads. To reduce the variance of input, we normalize the coordinates of each point of a proposal as

$$\mathbf{p}_n = \mathbf{T}\mathbf{p} \odot \mathbf{s}, \quad (12)$$

where

$$\mathbf{T} = \begin{bmatrix} \cos(\gamma) & \sin(\gamma) & 0 & -x \\ -\sin(\gamma) & \cos(\gamma) & 0 & -y \\ 0 & 0 & 1 & -z \\ 0 & 0 & 0 & 1 \end{bmatrix} \quad (13)$$

$$\mathbf{s} = [1/l \quad 1/w \quad 1/h \quad 1]^\top,$$

where \mathbf{T} and \mathbf{s} are defined in terms of the parameters of a proposal $\mathbf{b} = [c, x, y, z, l, w, h, \gamma]$, \odot is the Hadamard

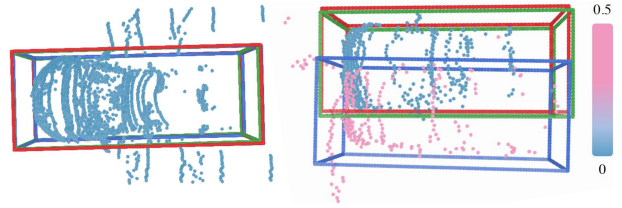


Fig. 4. Illustration of the *MLOD* results on the proposals and merged point cloud given by the stage-1 network. The color of each point represents its uncertainty quantity defined as the trace of the associated covariance. Blue to pink color indicates low to high uncertainty. The estimated and ground-truth bounding boxes are also marked with different colors: blue as the proposals in stage one, red as refinements in stage two, and green as ground truths.

product, and both \mathbf{p}_n and \mathbf{p} are represented in homogeneous coordinates. The output of the classification head is a binary variable indicating the probability of objectiveness. We regress residuals of the bounding box based on the proposal instead of directly regressing the final 3D bounding box. The regression target of the stage-2 network is

$$\mathbf{u}_i \triangleq \left[\frac{x - x_p}{l_p}, \frac{y - y_p}{w_p}, \frac{z - z_p}{h_p}, \frac{l - l_p}{l_p}, \right. \\ \left. \frac{w - w_p}{w_p}, \frac{h - h_p}{h_p}, \sin(\gamma - \gamma_p), \cos(\gamma - \gamma_p) \right], \quad (14)$$

where $\mathbf{u}_i \in \mathbb{R}^8$ are regression outputs for the i^{th} positive proposal respectively. We adopt the classification and regression loss which are defined in [23] as

$$\mathcal{L}_{\text{cls}} = \sum_i f_{\text{cls}}(p_i), \quad \mathcal{L}_{\text{reg}} = \sum_i f_{\text{reg}}(\mathbf{u}_i, \mathbf{u}_i^*), \quad (15)$$

where p_i represents the posterior probability of objectiveness, \mathcal{L}_{cls} is defined as the classification loss, $f_{\text{cls}}(\cdot)$ denotes the focal loss [24], \mathcal{L}_{reg} is defined as the normalized regression loss, and $f_{\text{reg}}(\cdot)$ denotes the smooth-L1 loss. Since we have the observation that the regression outputs \mathbf{u}_i should not be large when the uncertainty is small, we add a regularization term to penalize large \mathbf{u}_i in the low uncertainty cases:

$$\mathcal{L}_{\text{uct}} = \sum_i \exp\left(1 - f\left[\text{MAX}_j \text{tr}(\Xi_j)\right]\right) \cdot \|\mathbf{u}_i'\|_2, \quad (16)$$

where $\text{tr}(\Xi_j)$ is the point-wise uncertainty of the j^{th} point within the i^{th} proposal, MAX is the max operator, \mathbf{u}_i' is the regressed residual without the last cosine term for the i^{th} proposal, and $f(\cdot)$ clamps a value into the range $[10^{-3}, 0.5]$ to stabilize training. Finally, we define the loss function for training the stage-2 network as

$$\mathcal{L} = \mathcal{L}_{\text{reg}} + \eta \mathcal{L}_{\text{cls}} + \lambda \mathcal{L}_{\text{uct}}, \quad (17)$$

where η and λ are hyper-parameters to balance the weight of the classification loss and the uncertainty regularizer. We use $\eta = 2$ and $\lambda = 0.005$ in our experiments.

Fig. 4 visualizes the refinement results of our stage-2 network, which shows that our method produces similar results to the ground truths even if many points are uncertain. More quantitative examples are provided in Section VI-C.

VI. EXPERIMENT

In this section, we evaluate our proposed *MLOD* on LYFT multi-LiDAR dataset [3] in terms of accuracy and robustness under different levels of extrinsic perturbation. In particular, we aim to answer the following questions:

- 1) Can a multi-LiDAR object detector perform better accuracy than single-LiDAR methods?
- 2) Can *MLOD* improve the robustness of a multi-LiDAR object detector under extrinsic perturbation?

A. Implementation Details

1) *Dataset*: The LYFT dataset [3] provides a large amount of data collected in a variety of environments for the task of 3D object detection. Three 40-beam and calibrated Hesai LiDARs are mounted on the top, front-left and front-right position of the vehicle platform. The top LiDAR is set as the primary LiDAR, which is denoted by l^1 . The left and right LiDARs are auxiliary LiDARs, which are denoted by l^2 and l^3 respectively. This setting is according to LiDAR’s field of views (FOV). We select all multi-LiDAR data samples for our experiment. The data contain 4031 samples, 2500 of which are for training as well as validation and 1531 of which are for testing. The testing scenes are different from those in the training and validation sets.

2) *Metric*: Following the KITTI evaluation metrics [25], we compute the average precision on 3D bounding boxes (AP_{3D}) to measure the detection accuracy. We compute the AP_{3D} for 360° around the vehicle instead of evaluating only the 90° front view. According to the distance of objects, we set three different evaluation difficulties: *easy* ($< 20m$), *moderate* ($< 30m$), and *hard* ($< 50m$).

3) *Extrinsic Perturbation Injection*: We inject extrinsic perturbation on the original LYFT dataset to generate another set of data for robustness evaluation. We take Θ defined in (8) and adjust $\alpha \in [0, 0.1]$ with a 0.02 interval to simulate different levels of perturbation. To remove the effects from outliers, we bound the sampled perturbation within the σ position. Although the maximum value of α is small, but the effect of input perturbation on points is obvious. The noisy extrinsics are obtained by adding the sampled perturbation to the ground-truth extrinsics according to (1) and (2).

4) *Case Declaration*: We declare all cases of our methods which are tested in the following experiments as

- *LiDAR-Top*, *LiDAR-Left*, *LiDAR-Right*: single-LiDAR object detectors based on SECOND with the input data which are captured by the top LiDAR, front-left LiDAR and front-right LiDAR, respectively.
- *Input Fusion*, *Feature Fusion*, *Result Fusion*: multi-LiDAR object detectors with different fusion schemes, which are introduced in Section V-A.
- *MLOD-I*, *MLOD-F*, *MLOD-R*: *MLOD* with *Input Fusion*, *Feature Fusion*, *Result Fusion* as its stage-1 network respectively, which are described in Section V-B.
- *MLOD-I (OC)*, *MLOD-F (OC)*, *MLOD-R (OC)*: Variants of *MLOD* with an online calibration method to reduce extrinsic perturbation. This method is implemented with a point-to-plane ICP approach [26].

TABLE II

AVERAGE PRECISION ON THE MULTI-LiDAR LYFT TEST SET

Case	AP_{3D} IoU ≥ 0.7			AP_{3D} IoU ≥ 0.5		
	easy	mod.	hard	easy	mod.	hard
<i>LiDAR-Top</i>	62.7	52.9	37.8	89.8	80.1	61.8
<i>LiDAR-Left</i>	41.1	37.0	27.5	62.6	53.9	43.1
<i>LiDAR-Right</i>	40.9	31.5	23.4	53.9	52.4	35.5
<i>Input Fusion</i>	71.9	61.8	45.0	89.4	87.9	61.9
<i>MLOD-I</i>	72.4	62.7	45.6	89.6	88.3	61.8
<i>Feature Fusion</i>	71.1	60.2	44.1	89.6	87.4	61.6
<i>MLOD-F</i>	72.0	61.8	44.9	89.7	88.2	61.6
<i>Result Fusion</i>	71.4	61.6	44.5	89.5	87.7	62.0
<i>MLOD-R</i>	72.1	62.3	45.2	89.6	88.0	61.7

5) *Training of Stage-One Network*: We follow [10] to train the stage-1 network. During training, we conduct dataset sampling as in [10] and an augmentation of random flips on the y -axis as well as translation sampling within $[-0.5, 0.5]m$, $[-0.5, 0.5]m$, $[-0.3, 0.3]m$ on the x -, y -, and z -axes respectively, as well as rotation around the z -axis between $[-5, 5]^\circ$. In each epoch, the augmented data accounts for 30% of the whole training data. In *Result Fusion* and *Feature Fusion*, we directly take the calibrated and merged point clouds from different LiDARs as input. In *Result Fusion*, we calibrate all the proposal generators with temperature scaling [27].

6) *Training of Stage-Two Network*: To train our stage-2 network, we use all cases of the well-trained stage-1 networks to generate proposals. We use the training set to create proposals and inject uncertainties by selecting different α like the above section. In addition, we force the ratio between uncertainty-free samples and uncertain samples to be about 1:1.5 to stabilize the training process. A proposal is considered to be positive if its maximum IoU with ground-truth boxes is above 0.6, and is treated as negative if its maximum 3D IoU is below 0.45. During training, we conduct data augmentation of random flipping, scaling with a scale factor sampled from $[0.95, 1.05]$, translation along each axis between $[-0.1, 0.1]m$ and rotation around the z -axis between $[-3, 3]^\circ$. We also randomly sample 1024 points within the proposals to increase their diversity.

B. Results on the Multi-LiDAR LYFT Dataset

We evaluate both single-LiDAR and multi-LiDAR detectors on the LYFT test set, as reported in Tab. II. In this experiment, we do not consider any extrinsic perturbation between LiDARs. *LiDAR-Top* performs the best among all single-LiDAR detectors. This is because the top LiDAR has a complete 360° horizontal field of view (FOV). The left and right LiDARs, which are partially blocked by the vehicle, only have a 225° horizontal FOV. With sufficient measurements and decreasing occlusion areas, all multi-LiDAR detectors outperform *LiDAR-Top*, and the accuracy improvement gains up to 9.8AP. Tab. II also shows that all *MLOD* variants perform better or comparable to their one-stage counterparts. This proves that our stage-2 network refines the stage-1 proposals in perturbation-free cases.

TABLE III
MEAN AND VARIANCE OF ACCURACY (AP_{3D} IOU ≥ 0.7) UNDER DIFFERENT LEVEL OF THE EXTRINSIC PERTURBATION.

Cases	$\alpha = 0$			$\alpha = 0.02$			$\alpha = 0.04$			$\alpha = 0.1$		
	<i>easy</i>	<i>mod.</i>	<i>hard</i>	<i>easy</i>	<i>mod.</i>	<i>hard</i>	<i>easy</i>	<i>mod.</i>	<i>hard</i>	<i>easy</i>	<i>mod.</i>	<i>hard</i>
<i>Top LiDAR</i>	63.3	53.8	38.2	63.3	53.8	38.2	63.3	53.8	38.2	63.3	53.8	38.2
<i>Input Fusion</i>	71.2	62.3	45.4	64.7 ± 2.5	60.8 ± 0.4	44.4 ± 0.2	63.0 ± 0.5	53.9 ± 2.1	38.0 ± 0.2	60.9 ± 1.0	51.0 ± 0.6	36.1 ± 0.4
<i>MLOD-I</i>	71.2	62.8	46.1	66.3 ± 3.6	61.4 ± 0.4	44.8 ± 0.2	62.8 ± 0.4	56.3 ± 3.5	37.7 ± 0.2	61.6 ± 0.9	51.5 ± 0.6	35.8 ± 0.5
<i>Input Fusion (OC)</i>	-	-	-	66.8 ± 4.0	61.2 ± 0.5	44.5 ± 0.2	63.4 ± 0.4	60.0 ± 0.5	42.8 ± 2.2	62.4 ± 0.6	52.3 ± 0.4	37.2 ± 0.2
<i>MLOD-I (OC)</i>	-	-	-	67.6 ± 4.0	61.5 ± 0.5	44.9 ± 0.4	64.3 ± 2.5	60.4 ± 0.6	43.1 ± 2.5	62.9 ± 0.4	52.9 ± 0.3	36.8 ± 0.3
<i>Feature Fusion</i>	71.2	59.2	42.9	63.9 ± 2.5	59.2 ± 0.9	42.9 ± 0.9	62.9 ± 0.4	58.6 ± 2.5	41.1 ± 3.3	61.9 ± 0.6	51.1 ± 1.0	35.7 ± 0.7
<i>MLOD-F</i>	71.5	61.8	45.2	65.4 ± 3.3	60.9 ± 0.4	44.5 ± 0.2	63.7 ± 0.5	59.6 ± 2.0	42.1 ± 2.8	63.7 ± 0.5	53.9 ± 0.5	37.7 ± 0.4
<i>Feature Fusion (OC)</i>	-	-	-	66.2 ± 3.7	59.5 ± 0.8	43.5 ± 0.9	63.8 ± 2.2	59.6 ± 0.7	43.6 ± 0.7	63.3 ± 2.4	54.3 ± 3.0	37.9 ± 2.0
<i>MLOD-F (OC)</i>	-	-	-	67.7 ± 3.8	61.2 ± 0.4	44.7 ± 0.2	65.3 ± 3.2	60.4 ± 0.4	44.1 ± 0.2	64.9 ± 2.1	58.9 ± 2.9	38.2 ± 0.2
<i>Result Fusion</i>	70.4	62.1	44.0	64.5 ± 2.3	60.1 ± 0.4	43.4 ± 0.2	63.0 ± 0.5	55.5 ± 2.9	38.1 ± 0.2	61.1 ± 1.0	51.5 ± 0.5	36.8 ± 0.3
<i>MLOD-R</i>	70.7	62.4	45.1	65.7 ± 2.8	60.9 ± 0.4	44.2 ± 0.2	63.1 ± 0.5	59.3 ± 0.5	39.4 ± 2.5	62.6 ± 0.9	52.9 ± 0.5	37.0 ± 0.3
<i>Result Fusion (OC)</i>	-	-	-	66.5 ± 3.7	60.4 ± 0.6	43.5 ± 0.2	64.0 ± 2.2	59.3 ± 0.4	42.0 ± 1.9	62.5 ± 0.6	52.5 ± 0.4	37.5 ± 0.2
<i>MLOD-R (OC)</i>	-	-	-	66.8 ± 3.8	61.1 ± 0.5	44.3 ± 0.4	65.1 ± 3.2	59.9 ± 0.6	43.6 ± 0.4	63.3 ± 0.5	53.4 ± 0.3	37.5 ± 0.2

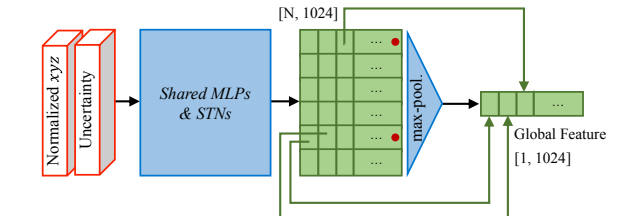
C. Robustness Under Extrinsic Perturbation

In this section, we evaluate the robustness of our proposed methods under extrinsic perturbation. Three levels of perturbation are performed: no perturbation ($\alpha = 0$), moderate perturbation ($\alpha = 0.02$, $\alpha = 0.04$) and high perturbation ($\alpha = 0.1$). An example is displayed in Fig. 6(a) ($\alpha = 0.04$), where massive noisy points (50cm misplacement) appear at 35m away. This is a typical phenomenon of extrinsic perturbation in multi-LiDAR systems. We randomly select 300 data samples from the test set with 10 trials to conduct evaluations at each level.

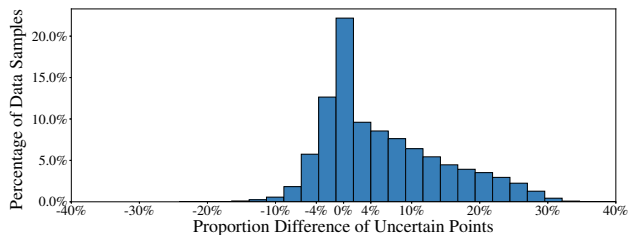
In Tab. III, the detection results in terms of the means and variances as AP_{3D}(IoU ≥ 0.7) are detailed. We see that all variants of *MLOD* (*MLOD-I*, *MLOD-F*, *MLOD-R*) perform better or comparable to their one-stage counterparts. Online calibration baselines (*Input Fusion (OC)*, *Feature Fusion (OC)*, *Result Fusion (OC)*) demonstrate better results than those without online calibration. With the assistance of *MLOD*, their performance is further enhanced. *Feature Fusion* and its variants perform better than the others. We explain that *Feature Fusion* fuses data in a high dimensional embedding space, and each feature is related to an area of the cognitive region. Regarding *LiDAR-Top*, *MLOD*'s variants also perform better or comparable results. A qualitative result is illustrated in Fig. 6, and more results are shown in the supplementary materials.

We also conduct a sensitivity analysis of extrinsic uncertainty with a fixed α as input in the supplementary material. The same conclusion still holds. To further explore the reason behind the performance of *MLOD*, we investigate the characteristics of active points. These points are activated by PointNet before the max-pooling layer (Fig. 5(a)). We compute the difference between the input points and active points on the proportion of uncertain points⁴ for each sample. Fig. 5(b) plots a histogram of the proportional difference

⁴The proportion of the uncertain points is defined as $|\mathcal{P}_u|/|\mathcal{P}|$, where $\mathcal{P}_u \in \mathcal{P}$ is the set of uncertain points with the uncertain quantity > 0.05.



(a) The details of PointNet [11].



(b) Histogram of proportional differences between input uncertain points and active uncertain points.

Fig. 5. (a) Active points (Red dotted) form the global features after applying the max-pooling symmetric function. (b) The proportional differences larger than 0.04 (4%) occupy about 44.0%, and those less than -0.04 (-4%) occupy about 8.5%. Less uncertain points exist in the active points.

overall input samples. It shows that the network deactivates highly uncertain points, and explain why *MLOD* is robust.

VII. CONCLUSION

In this paper, we analyze the extrinsic perturbation effect on multi-LiDAR-based 3D object detection. We propose a two-stage network to both fuse the data from multiple LiDARs and handle extrinsic perturbation after data fusion. We conduct extensive experiments on a real-world dataset and discuss the results in different levels of extrinsic perturbation. In the perturbation-free situation, we show the multi-LiDAR fusion approaches obtain better accuracy than single-LiDAR detectors. Under extrinsic perturbation, *MLOD* performs great robustness with the assistance of the uncertainty prior. A future direction concerns the combination with sensors in various modalities, e.g., LiDAR-camera-radar setups, for developing a more accurate object detector.

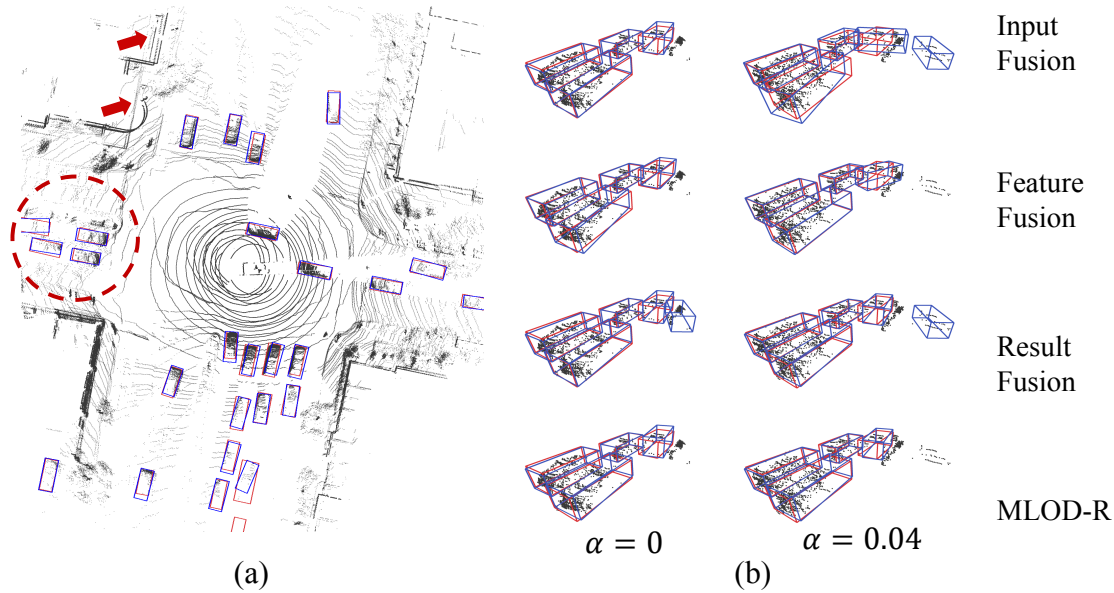


Fig. 6. Visualization of the fusion schemes results (For comparison, here we sample the extrinsic perturbation for all the ($\alpha = 0.04$) cases at the σ position according to Θ). (a) Bird's-eye view of *MLOD*'s detection results when $\alpha = 0.04$. The extrinsic perturbation is observed from where the red arrows indicate. (b) A close-up view of results estimated by different fusion schemes within the red circle. Left to right: $\alpha = 0$, $\alpha = 0.04$, *Input Fusion*, *Feature Fusion* and *Result Fusion* suffer from false positives or inaccurate boxes caused by extrinsic perturbation and inaccurate box location. Compared with its counterpart (*Result Fusion*), *MLOD-R* eliminates the false positives and refines the 3D boxes.

REFERENCES

- [1] J. Jiao, Y. Yu, Q. Liao, H. Ye, R. Fan, and M. Liu, "Automatic calibration of multiple 3d lidars in urban environments," in *2019 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE, 2019, pp. 15–20.
- [2] P. Sun, H. Kretzschmar et al., "Scalability in perception for autonomous driving: Waymo open dataset," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2020, pp. 2446–2454.
- [3] R. Kesten, M. Usman et al., "Lyft level 5 av dataset 2019," [url:https://level5.lyft.com/dataset/](https://level5.lyft.com/dataset/), 2019.
- [4] X. Chen, H. Ma, J. Wan, B. Li, and T. Xia, "Multi-view 3d object detection network for autonomous driving," in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2017, pp. 1907–1915.
- [5] Y. Ling and S. Shen, "High-precision online markerless stereo extrinsic calibration," in *2016 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE, 2016, pp. 1771–1778.
- [6] J. Lin, X. Liu, and F. Zhang, "A decentralized framework for simultaneous calibration, localization and mapping with multiple lidars," *arXiv preprint arXiv:2007.01483*, 2020.
- [7] W. Maddern, G. Pascoe, C. Linegar, and P. Newman, "1 year, 1000 km: The oxford robotcar dataset," *The International Journal of Robotics Research*, vol. 36, no. 1, pp. 3–15, 2017.
- [8] B. Li, "3d fully convolutional network for vehicle detection in point cloud," in *2017 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE, 2017, pp. 1513–1518.
- [9] Y. Zhou and O. Tuzel, "Voxelnet: End-to-end learning for point cloud based 3d object detection," in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2018, pp. 4490–4499.
- [10] Y. Yan, Y. Mao, and B. Li, "Second: Sparsely embedded convolutional detection," *Sensors*, vol. 18, no. 10, p. 3337, 2018.
- [11] C. R. Qi, H. Su, K. Mo, and L. J. Guibas, "Pointnet: Deep learning on point sets for 3d classification and segmentation," in *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2017, pp. 652–660.
- [12] C. R. Qi, L. Yi, H. Su, and L. J. Guibas, "Pointnet++: Deep hierarchical feature learning on point sets in a metric space," in *Advances in neural information processing systems*, 2017, pp. 5099–5108.
- [13] S. Shi, X. Wang, and H. Li, "Pointcnn: 3d object proposal generation and detection from point cloud," in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2019, pp. 770–779.
- [14] B. Graham and L. van der Maaten, "Submanifold sparse convolutional networks," *arXiv preprint arXiv:1706.01307*, 2017.
- [15] C. R. Qi, W. Liu, C. Wu et al., "Frustum pointnets for 3d object detection from rgb-d data," in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2018, pp. 918–927.
- [16] D. Xu, D. Anguelov, and A. Jain, "Pointfusion: Deep sensor fusion for 3d bounding box estimation," in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2018, pp. 244–253.
- [17] A. Kendall and Y. Gal, "What uncertainties do we need in bayesian deep learning for computer vision?" in *Advances in neural information processing systems*, 2017, pp. 5574–5584.
- [18] A. Kendall, V. Badrinarayanan, and R. Cipolla, "Bayesian segnet: Model uncertainty in deep convolutional encoder-decoder architectures for scene understanding," in *British Machine Vision Conference 2017*, BMVC 2017, 2017.
- [19] D. Feng, L. Rosenbaum, F. Timm, and K. Dietmayer, "Leveraging heteroscedastic aleatoric uncertainties for robust real-time lidar 3d object detection," in *2019 IEEE Intelligent Vehicles Symposium (IV)*. IEEE, 2019, pp. 1280–1287.
- [20] T. D. Barfoot and P. T. Furgale, "Associating uncertainty with three-dimensional poses for use in estimation problems," *IEEE Transactions on Robotics*, vol. 30, no. 3, pp. 679–693, 2014.
- [21] Y. Kim and A. Kim, "On the uncertainty propagation: Why uncertainty on lie groups preserves monotonicity?" in *2017 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE, 2017, pp. 3425–3432.
- [22] T. Liu, Q. Liao, L. Gan et al., "Hercules: An autonomous logistic vehicle for contact-less goods transportation during the covid-19 outbreak," *arXiv preprint arXiv:2004.07480*, 2020.
- [23] P. Yun, L. Tai, Y. Wang, C. Liu, and M. Liu, "Focal loss in 3d object detection," *IEEE Robotics and Automation Letters*, vol. 4, no. 2, pp. 1263–1270, 2019.
- [24] T.-Y. Lin, P. Goyal, R. Girshick, K. He, and P. Dollár, "Focal loss for dense object detection," in *Proceedings of the IEEE international conference on computer vision*, 2017, pp. 2980–2988.
- [25] A. Geiger, P. Lenz, C. Stiller, and R. Urtasun, "Vision meets robotics: The kitti dataset," *The International Journal of Robotics Research*, vol. 32, no. 11, pp. 1231–1237, 2013.
- [26] F. Pomerleau et al., "Comparing icp variants on real-world data sets," *Autonomous Robots*, vol. 34, no. 3, pp. 133–148, 2013.
- [27] C. Guo, G. Pleiss, Y. Sun, and K. Q. Weinberger, "On calibration of modern neural networks," in *International Conference on Machine Learning*, 2017, pp. 1321–1330.