

BIDIRECTIONAL PROPAGATION FOR CROSS-MODAL 3D OBJECT DETECTION

Yifan Zhang¹, Qijian Zhang¹, Junhui Hou¹, Yixuan Yuan², and Guoliang Xing²

¹City University of Hong Kong, ²The Chinese University of Hong Kong

{yzhang3362-c, qijizhang3-c}@my.cityu.edu.hk; jh.hou@cityu.edu.hk; yxyuan@ee.cuhk.edu.hk; glxing@ie.cuhk.edu.hk

ABSTRACT

Recent works have revealed the superiority of feature-level fusion for cross-modal 3D object detection, where fine-grained feature propagation from 2D image pixels to 3D LiDAR points has been widely adopted for performance improvement. Still, the potential of heterogeneous feature propagation between 2D and 3D domains has not been fully explored. In this paper, in contrast to existing pixel-to-point feature propagation, we investigate an opposite point-to-pixel direction, allowing point-wise features to flow inversely into the 2D image branch. Thus, when jointly optimizing the 2D and 3D streams, the gradients back-propagated from the 2D image branch can boost the representation ability of the 3D backbone network working on LiDAR point clouds. Then, combining pixel-to-point and point-to-pixel information flow mechanisms, we construct a bidirectional feature propagation framework, dubbed BiProDet. In addition to the architectural design, we also propose normalized local coordinates map estimation, a new 2D auxiliary task for the training of the 2D image branch, which facilitates learning local spatial-aware features from the image modality and implicitly enhances the overall 3D detection performance. Extensive experiments and ablation studies validate the effectiveness of our method. Notably, we rank 1st on the highly competitive KITTI benchmark on the cyclist class by the time of submission. The source code is available at <https://github.com/Eaphan/BiProDet>.

1 INTRODUCTION

In recent years, 3D object detection has received increasing academic and industrial attention driven by the thriving of autonomous driving scenarios, where the two dominating data modalities 3D point clouds and 2D RGB images demonstrate complementary properties in many aspects. 3D point clouds scanned by LiDAR encode accurate structure and depth cues in the geometric domain but typically suffer from severe sparsity, incompleteness, and non-uniformity. The 2D RGB images captured by optical cameras convey rich semantic features in the visual domain and facilitate the application of well-developed image learning architectures (He et al., 2016) but pose challenges in reliably modeling spatial structures. More importantly, due to the essential difficulties in processing *irregular* and *unordered* point clouds, the actual learning ability of neural network backbones working on 3D LiDAR data is still relatively *insufficient*, thus limiting the performance of LiDAR-based detection frameworks. Practically, despite the efforts made by previous studies (Wang et al., 2019; Vora et al., 2020) devoted to cross-modal learning, how to reasonably mitigate the big modality gap between images and point clouds and fuse the corresponding heterogeneous feature representations still remains an open and non-trivial problem.

Depending on specific fusion strategies between image and point cloud modalities, existing cross-modal 3D object detection pipelines can be categorized into: 1) *proposal-level*; 2) *result-level*; 3) *pseudo-LiDAR-based*; and 4) *point-level* processing paradigms, as illustrated in Fig. 1. Concretely, proposal-level fusion methods (Chen et al., 2017; Ku et al., 2018) perform feature fusion for both modalities over each of the proposals/anchors, while result-level fusion methods (Pang et al., 2020) directly manipulate the output 2D and 3D object candidates by leveraging geometric and semantic consistencies. However, the former two fusion paradigms do not make full use of fine-grained correspondence between 3D LiDAR points and 2D images, and thus show inferior detection performance.

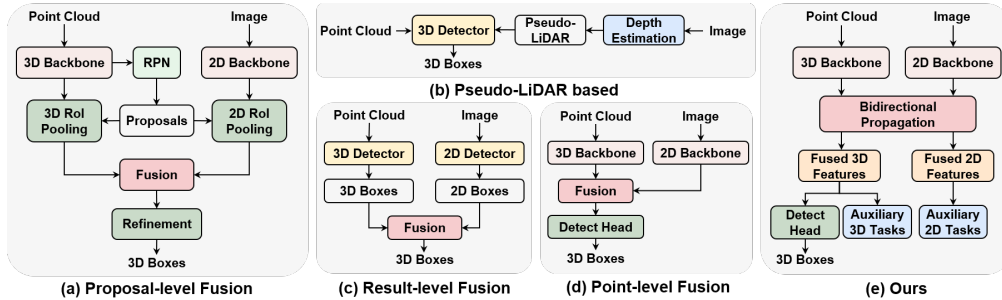


Figure 1: Comparison on the schematic paradigms of different cross-modal 3D object detectors.

For another alternative fusion paradigm based on pseudo-LiDAR representations (You et al., 2020), auxiliary learning modules specialized for stereo depth estimation are integrated to synthesize dense pseudo-LiDAR point clouds as additional complements to the original real-scanned sparse LiDAR data. Such a strategy indeed brings performance boosts but it becomes more technically cumbersome and demanding, since ground-truth depth maps are required as supervision signals. Moreover, as investigated previously (Qian et al., 2020), it is non-trivial to harmonize the joint end-to-end optimization of the object detector and the depth estimator. To make use of point-pixel correspondences between images and point clouds, there has been another family of point-level fusion approaches (Huang et al., 2020; Vora et al., 2020) characterized by using projection from the 3D space of point cloud to 2D image planes. That is, for each 3D point, its geometric feature extracted from the 3D LiDAR learning branch is fused with the semantic feature (retrieved from the 2D image learning branch) of the corresponding 2D pixel. Benefiting from more fine-grained feature manipulation, point-level fusion typically demonstrates more significant performance improvement and thus is considered as a more promising cross-modal fusion paradigm.

Still, based on our observations, despite the superiority of point-level fusion, previous studies have not fully exploited its great potential in terms of specific architectural design and technical workflow. Fundamentally, the performance boost achieved by previous methods can be mainly credited to the fact that additional discriminative 2D semantic information is incorporated into the 3D detection pipeline. However, there is no evidence that the actual representation capability of the 3D LiDAR backbone network, which is also one of the most critical influencing factors, is strengthened. Besides, it is also noticed that the choice of 2D auxiliary tasks used for training the image learning branch also makes a difference, which essentially determines the consequent 2D feature representations to be propagated. Previous methods directly resort to classic semantic understanding tasks (e.g., image segmentation) for the learning of visual patterns, which may not be the optimal task that harmonizes with the 3D LiDAR learning branch for geometric modeling.

Given the above issues, in this paper we propose a bidirectional feature propagation architecture, namely BiProDet, for cross-modal 3D object detection. Foremost, our core component lies in a novel point-to-pixel feature propagation mechanism allowing 3D geometric features of LiDAR point clouds to flow into the 2D image learning branch. Our BiProDet *differentiates* from previous cross-modal 3D detection pipelines that solely pay attention to unidirectional 2D-to-3D (pixel-to-point) feature propagation while neglecting the potential effects of the opposite 3D-to-2D (point-to-pixel) direction. Functionally, different from existing pixel-to-point propagation, the proposed point-to-pixel propagation turns out to *be capable of enhancing the representation ability of 3D LiDAR backbone networks*. In fact, it is worth emphasizing that this is an interesting phenomenon, and we reason that this is because *the gradients back-propagated from the auxiliary 2D image branch built upon mature 2D convolutional neural networks (CNNs) with the impressive learning ability can effectively strengthen the 3D LiDAR branch*. Taking a step forward, we integrate pixel-to-point and point-to-pixel propagation directions into a unified interactive bidirectional feature propagation framework such that 2D and 3D learning branches are inter-strengthened, eventually contributing to more significant performance gains. Note that the previous work also develops a bidirectional model for joint optical and scene flow estimation (Liu et al., 2022). Still, we design a concise yet effective bidirectional propagation strategy without bells and whistles. Unlike Liu et al. (2022) detach the gradients in bidirectional fusion modules, we demonstrate that gradients back-propagated from the auxiliary 2D image branch could enhance the 3D backbone. In addition, we also particularly customize a new 2D auxiliary task called normalized local coordinate (NLC) map estimation, which facilitates learning local spatial-aware feature representations from the 2D image. Such a task also

implicitly helps to capture more discriminative geometric information from the main 3D LiDAR point cloud learning branch, thus improving the overall detection performance, especially for hard cases of distant and/or highly-occluded objects.

We conduct extensive experiments on the prevailing KITTI (Geiger et al., 2012) benchmark dataset. Compared with state-of-the-art single-modal and cross-modal 3D object detectors, our framework achieves remarkable performance improvement. Notably, our method ranks 1st on the cyclist class of KITTI 3D detection benchmark¹. Conclusively, our main contributions are three-fold:

- In contrary to existing pixel-to-point cross-modal fusion strategies, we explore an opposite point-to-pixel feature propagation paradigm, which turns out to be capable of enhancing the representation capability of 3D backbone networks working on LiDAR point clouds.
- Combining the existing pixel-to-point and the newly proposed point-to-pixel feature propagation schemes, we further construct the powerful bidirectional feature propagation for cross-modal 3D objection detection.
- We reveal the importance of specific 2D auxiliary tasks used for training the 2D image learning branch which has not received enough attention in previous works, and introduce 2D NLC map estimation to facilitate learning spatial-aware features and implicitly boost the overall 3D detection performance.

2 RELATED WORK

LiDAR-based 3D Object Detection could be roughly divided into two categories. 1) *Voxel-based 3D detectors* typically voxelize the point clouds into grid-structure forms of a fixed size (Zhou & Tuzel, 2018). Yan et al. (2018) introduced a more efficient sparse convolution to accelerate training and inference. He et al. (2020a) employed auxiliary tasks including center estimation and foreground segmentation to guide the network to learn the intra-object relationship. 2) *Point-based 3D detectors* consume the raw 3D point clouds directly and generate predictions based on (downsampled) points. Shi et al. (2019) applied a point-based feature extractor and generated high-quality proposals on foreground points. 3DSSD (Yang et al., 2020) adopted a new sampling strategy named F-FPS as a supplement to D-FPS to preserve enough interior points of foreground instances. It also built a one-stage anchor-free 3D object detector based on feasible representative points. Zhang et al. (2022c) and Chen et al. (2022) introduced semantic-aware down-sampling strategies to preserve foreground points as much as possible. Shi & Rajkumar (2020b) proposed to encode the point cloud by a fixed-radius near-neighbor graph. Besides, PV-RCNN (Shi et al., 2020a) adopted a voxel-based RPN and leverages point-wise features extracted from the voxel set abstraction module to refine the proposals.

Camera-based 3D Object Detection. Early works (Mousavian et al., 2017; Manhardt et al., 2019) designed monocular 3D detectors by referring to 2D detectors (Ren et al., 2015; Tian et al., 2019) and utilizing 2D-3D geometric constraints. Another way is to convert images to pseudo-lidar representations via monocular depth estimation and then resort to LiDAR-based methods (Wang et al., 2019). Chen et al. (2021) built dense 2D-3D correspondence by predicting object coordinate map and adopted uncertainty-driven PnP to estimate object location. Recently, multi-view detectors like BEVStereo (Li et al., 2022) have also achieved promising performance, getting closer to LiDAR-based methods.

Cross-Modal 3D Object Detection can be roughly divided into four categories: proposal-level, result-level, point-level, and pseudo-lidar-based approaches. *Proposal-level* fusion methods (Chen et al., 2017; Ku et al., 2018; Liang et al., 2019) adopted a two-stage framework to fuse image features and point cloud features corresponding to the same anchor or proposal. For *result-level* fusion, Pang et al. (2020) exploited the geometric and semantic consistencies of 2D detection results and 3D detection results to improve single-modal detectors. Qi et al. (2018) generated 2D region proposals using a CNN and lifted them to frustum proposals, where a PointNet-based network is utilized to estimate 3D box. Obviously, both proposal-level and decision-level fusion strategies are coarse and do not make full use of correspondence between LiDAR points and images. *Pseudo-LiDAR* based methods (You et al., 2020; Wu et al., 2022; Liang et al., 2019) employed depth estimation/completion and converted the estimated depth map to pseudo-LiDAR points to complement raw sparse point clouds. However, such methods require extra depth map annotations that are high-cost. Regarding *point-level* fusion methods, Liang et al. (2018) proposed the continuous fusion layer to retrieve corresponding image features of nearest 3D points for each grid in BEV feature map. Especially, Vora

¹www.cvlibs.net/datasets/kitti/eval_object.php?obj_benchmark=3d

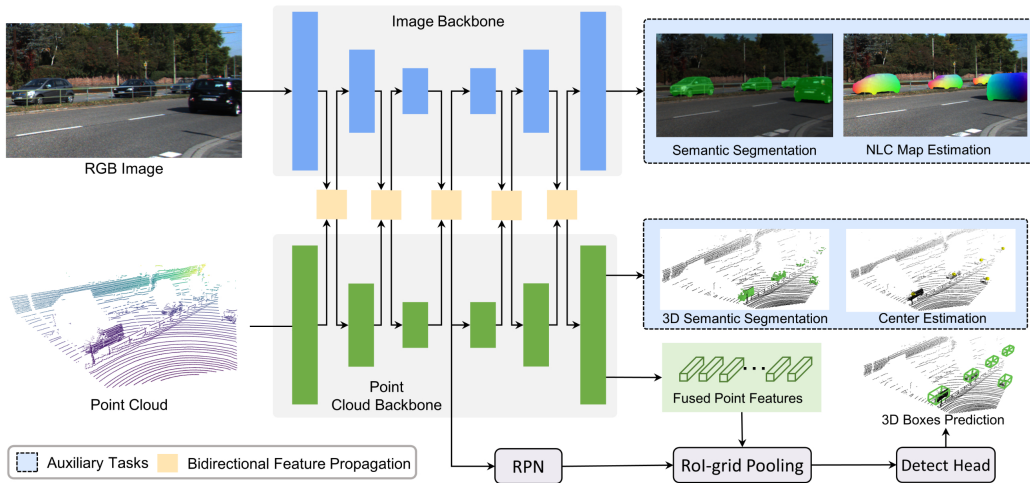


Figure 2: Flowchart of the proposed BiProDet for cross-modal 3D object detection. It contains two branches separately extracting features from input 2D RGB images and 3D point clouds. BiProDet is mainly featured with (1) the multi-stage bidirectional propagation manner for deep interaction between two modalities and (2) the supervision task of NLC map estimation for promoting the image branch to exploit local spatial features.

et al. (2020) and Huang et al. (2020) retrieved the semantic scores or features by projecting points to image plane. Liu et al. (2021) proposed EPNNet++, where an LI-Fusion layer is used to enable more interaction between the two modalities to obtain more comprehensive features. And Wang et al. (2021) decorated point clouds with middle CNN features extracted from 2D detection models. However, neither semantic segmentation nor 2D detection tasks can enforce the network to learn 3D-spatial-aware image features. By contrast, we introduce NLC Map estimation as an auxiliary task to promote the image branch to learn local spatial information to supplement the sparse point cloud. Besides, Bai et al. (2022) proposed TransFusion, a transformer-based detector that performs fine-grained fusion with attention mechanisms on the BEV level. Phillion & Fidler (2020) and Liang et al. (2022) estimated the depth of multi-view images and transformed the camera feature maps into the BEV space before fusing them with LiDAR BEV features.

3 PROPOSED METHOD

Overview. Architecturally, the overall processing pipeline of cross-modal 3D object detection consists of an image branch and a point cloud branch, learning feature representations from 2D RGB images and 3D LiDAR point clouds, respectively. Existing methods implement cross-modal learning only by incorporating pixel-wise semantic cues from the 2D image branch into the 3D point cloud branch for feature fusion. Here, motivated by the observation that point-based networks typically show insufficient learning ability due to the essential difficulties in processing irregular and unordered point cloud data modality (Zhang et al., 2022a), we additionally seek to boost the expressive power of the point cloud backbone network with the assistance of the 2D image branch, whose potential is ignored in previous studies. It is also worth reminding that our ultimate goal is not to design a more powerful point cloud backbone network structure but to elegantly make full use of the available resources in this cross-modal application scenario.

As shown in Fig. 2, in addition to pixel-to-point feature propagation explored by mainstream point-level fusion methods (Liang et al., 2018; Huang et al., 2020), we propose point-to-pixel propagation allowing features to flow inversely from the point cloud branch to the image branch, based on which we can achieve bidirectional feature propagation implemented in a multi-stage fashion. In this way, not only the image features propagated via the pixel-to-point module can provide additional semantic information, but also the gradients backpropagated from the training objectives of the image branch can boost the representation ability of the point cloud backbone. Besides, we also employ auxiliary tasks to train the pipeline, aiming to enforce the network to learn rich semantic and spatial representations. In particular, we propose NLC map estimation to promote the image branch to learn spatial-aware features, providing a necessary complement to the sparse spatial representation

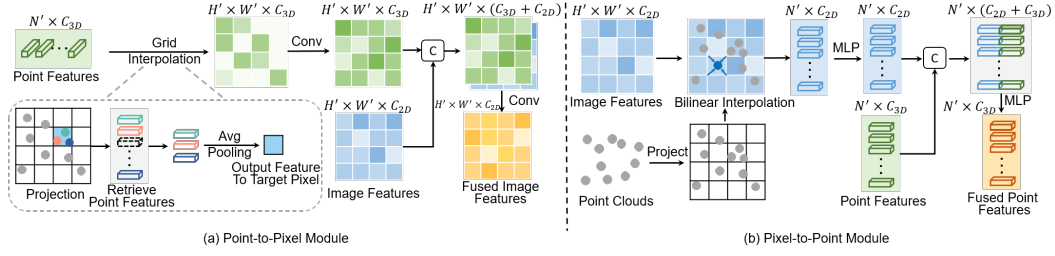


Figure 3: Illustration of the proposed bidirectional feature propagation.

extracted from point clouds, especially for distant or highly occluded cases. We refer the readers to Appendix A.2 for detailed network structures of the image and point cloud backbones.

3.1 BIDIRECTIONAL FEATURE PROPAGATION

As illustrated in Fig. 3, the proposed bidirectional feature propagation consists of a pixel-to-point module and a point-to-pixel module, which bridges the two learning branches that operate on RGB images and LiDAR point clouds. Functionally, the point-to-pixel module applies grid-level interpolation on point-wise features to produce a 2D feature map, while the pixel-to-point module retrieves the corresponding 2D image features by projecting 3D points to the 2D image plane.

In general, we partition the overall 2D and 3D branches into the same number of stages, such that we can perform bidirectional feature propagation at each stage between the corresponding layers of the image and point cloud learning networks. Without loss of generality, below we only focus on a certain stage, where we can acquire a 2D feature map $\mathbf{F} \in \mathbb{R}^{C_{2D} \times H' \times W'}$ with C_{2D} channels and dimensions $H' \times W'$ from the image branch and a set of C_{3D} -dimensional embedding vectors $\mathbf{G} = \{\mathbf{g}_i \in \mathbb{R}^{C_{3D}}\}_{i=1}^{N'}$ of the corresponding N' 3D points.

Point-to-Pixel Module. We design this module to propagate geometric information from the 3D point cloud branch into the 2D image branch. Formally, we expect to construct a 2D feature map $\mathbf{F}^{P2I} \in \mathbb{R}^{C_{3D} \times H' \times W'}$ by querying from the acquired point-wise embeddings \mathbf{G} . More specifically, for each pixel position (r, c) over the targeted 2D grid of \mathbf{F}^{P2I} of dimensions $H' \times W'$, we collect the corresponding geometric features of points that are projected inside the range of this pixel and then perform feature aggregation by avg-pooling $\text{AVG}(\cdot)$:

$$\mathbf{F}_{r,c}^{P2I} = \text{AVG}(\{\mathbf{g}_j | j = 1, \dots, n\}), \quad (1)$$

where n counts the number of points that can be projected inside the range of the pixel located at the r^{th} row and the c^{th} column. Particularly, for empty pixel positions where there is no point projected inside, we simply define its feature value as $\mathbf{0}$.

After that, instead of directly feeding the “interpolated” 2D feature map \mathbf{F}^{P2I} into the subsequent 2D convolutional stage, we introduce a few additional refinement procedures for feature smoothness as described below:

$$\mathbf{F}^{fuse} = \text{CONV}(\text{CAT}[\text{CONV}(\mathbf{F}^{P2I}), \mathbf{F}]), \quad (2)$$

where $\text{CAT}[\cdot, \cdot]$ and $\text{CONV}(\cdot)$ stand for feature channel concatenation and the convolution operation, respectively, and the resulting \mathbf{F}^{fuse} is further fed into the subsequent layers of the 2D image branch.

Pixel-to-Point Module. From an opposite direction, this module propagates semantic information from the 2D image branch into the 3D point cloud branch. We start by projecting the corresponding N' 3D points onto the 2D image plane, obtaining their planar coordinates denoted as $\mathbf{X} = \{\mathbf{x}_i \in \mathbb{R}^2\}_{i=1}^{N'}$. Considering that these 2D coordinates are continuously and irregularly distributed over the projection plane, we apply bilinear interpolation to compute exact image features at projected positions, deducing a set of point-wise embeddings $\mathbf{F}^{I2P} = \{\mathbf{F}(\mathbf{x}_i) \in \mathbb{R}^{C_{2D}} | i = 1, \dots, N'\}$.

Similar to our practice in the point-to-pixel module, the interpolated 3D point-wise features \mathbf{F}^{I2P} need to be refined by the following procedures:

$$\mathbf{G}^{fuse} = \text{MLP}(\text{CAT}[\text{MLP}(\mathbf{F}^{I2P}; \theta_1), \mathbf{G}]; \theta_2), \quad (3)$$

where $\text{MLP}(\cdot; \theta)$ denotes shared multi-layer perceptrons parameterized by learnable weights θ , and the resulting G^{fuse} further passes through the subsequent layers of the 3D point cloud branch.

3.2 AUXILIARY TASKS FOR TRAINING

The semantic and spatial-aware information is essential for determining the category and boundary of objects. To promote the network to uncover such information of input data, we also leverage auxiliary tasks when training the whole pipeline. For the point cloud branch, following SA-SSD (He et al., 2020b), we introduce 3D semantic segmentation and center estimation to learn structure-aware features. For the image branch, we particularly propose NLC map estimation, in addition to 2D semantic segmentation.

Normalized Local Coordinate (NLC) System.

We define the NLC system as taking the center of an object as the origin, aligning the x -axis towards the head direction of its ground-truth (GT) bounding box. Fig. 4 (a) shows an example of the NLC system for a typical car. With the geometry relationship between the input RGB image and point cloud, we can project NLCs to the 2D image plane to construct a 2D NLC map with three channels corresponding to three spatial dimensions, as illustrated in Fig. 4 (b). We refer the reader to Appendix A.1 for the detailed process.




Figure 4: Illustration of the (a) NLC system and (b) 2D NLC map.

NLC Map Estimation. Intuitively, vision-based semantic information is useful for network to distinguish the foreground from the background. However, the performance bottleneck of existing detectors lies in the limited localization accuracy in distant or highly occluded cases, where the spatial structure is incomplete due to the sparsity of point clouds. To this end, we further take NLC map as supervision to learn the relative position of each pixel inside the object from the image. We expect such an auxiliary task can drive the image branch to learn local spatial-aware features, which serve as a complement to the sparse spatial representation extracted from point clouds. Besides, this task could have an effect on augmenting the representation ability of the point cloud branch.

Remark. The proposed NLC map estimation shares the identical objective with pseudo-LiDAR-based methods, i.e., enhancing the spatial representation limited by the sparsity of point clouds. However, compared with the manner of learning a pseudo-LiDAR representation via depth completion, our NLC map estimation has the following advantages: **1)** the *local* NLC is easier to learn than the *global* depth owing to its scale-invariant property at different distances; **2)** our NLC map estimation can be naturally incorporated and trained with the detection pipeline end-to-end, while it is non-trivial to optimize the depth estimator and LiDAR-based detector simultaneously in pseudo-LiDAR based methods though it can be implemented in a somewhat complex way (Qian et al., 2020); and **3)** the pseudo-LiDAR representations require ground-true dense depth images as supervision, which may not be available in reality, but our method does not require such information.

Semantic Segmentation. Leveraging semantics of 3D point clouds has demonstrated effectiveness in point-based detectors, owing to the explicit preservation of foreground points at downsampling (Zhang et al., 2022c; Chen et al., 2022). Therefore, we further introduce the auxiliary semantic segmentation tasks not only in the point cloud branch but also in the image branch, to exploit richer semantic information. Needless to say, additional semantic features extracted from images facilitate the network to distinguish true positive candidate boxes from the false-positive.

3.3 LOSS FUNCTION

The NLC map estimation task is optimized with the loss function defined as

$$L_{NLC} = \frac{1}{N_{pos}} \sum_i^N \|\mathbf{p}_i^{NLC} - \hat{\mathbf{p}}_i^{NLC}\|_H \cdot \mathbf{1}_{\mathbf{p}_i}, \quad (4)$$

where \mathbf{p}_i is the i^{th} LiDAR point, N_{pos} is the number of foreground LiDAR points, $\|\cdot\|_H$ is the Huber-loss, $\mathbf{1}_{\mathbf{p}_i}$ indicates the loss is calculated only with foreground points, and \mathbf{p}_i^{NLC} and $\hat{\mathbf{p}}_i^{NLC}$ are the NLCs of ground-truth and prediction at corresponding pixel for foreground points.

We use the standard cross-entropy loss to optimize both 2D and 3D semantic segmentation, denoted as L_{sem}^{2D} and L_{sem}^{3D} respectively. And the loss of center estimation is computed as

$$L_{ctr} = \frac{1}{N_{pos}} \sum_i^N \|\Delta\mathbf{p}_i - \Delta\hat{\mathbf{p}}_i\|_H \cdot \mathbf{1}_{\mathbf{p}_i}, \quad (5)$$

where $\Delta\mathbf{p}_i$ is the target offsets from points to the corresponding object center, and $\Delta\hat{\mathbf{p}}_i$ denotes the output of center estimation head.

Besides the auxiliary tasks, we follow (Chen et al., 2022) to define the loss of the RPN stage as

$$L_{rpn} = L_c + L_r + L_s + L_d, \quad (6)$$

where L_c is the classification loss, L_r the regression loss, L_s the shifting loss between the predicted shifts and residuals from the sampled points to their corresponding object centers in candidate generate layer (Yang et al., 2020), and L_d the point segmentation loss in SA layers to enable semantic-guided sampling. In addition, we also adopt commonly used proposal refinement loss L_{rcnn} as defined in (Shi et al., 2020a).

In all, the loss for optimizing the overall pipeline in an end-to-end manner is written as

$$L_{total} = L_{rpn} + L_{rcnn} + \lambda_1 L_{NLC} + \lambda_2 L_{sem}^{2D} + \lambda_3 L_{sem}^{3D} + \lambda_4 L_{ctr}, \quad (7)$$

where λ_1 , λ_2 , λ_3 , and λ_4 are hyper-parameters that are empirically set to 1.

4 EXPERIMENTS

4.1 EXPERIMENT SETTINGS

Datasets and Metrics. We conducted experiments on the prevailing KITTI benchmark dataset, which contains two modalities of 3D point clouds and 2D RGB images. Following previous works (Shi et al., 2020a), we divided all training data into two subsets, i.e., 3712 samples for training and the rest 3769 for validation. Performance is evaluated by the Average Precision (AP) metric under IoU thresholds of 0.7, 0.5, and 0.5 for car, pedestrian, and cyclist categories, respectively. We computed APs with 40 sampling recall positions by default, instead of 11. For the 2D auxiliary task of semantic segmentation, we used the instance segmentation annotations as provided in (Qi et al., 2019). Besides, we also conducted experiments on the Waymo Open Dataset (WOD) (Sun et al., 2020), which can be found in Appendix A.4.

Implementation Details. For the image branch, we used ResNet18 (He et al., 2016) as the backbone encoder, followed by a decoder composed of pyramid pooling module (Zhao et al., 2017) and several upsampling layers to give the outputs of semantic segmentation and NLC map estimation. For the point cloud branch, we deployed a point-based network like (Yang et al., 2020) with extra FP layers and further applied semantic-guided farthest point sampling (S-FPS) (Chen et al., 2022) in SA layers. Thus, we implemented bidirectional feature propagation between the top of each SA or FP layer and their corresponding locations at the image branch. Please refer to Appendix A.2 for more details.

4.2 COMPARISON WITH STATE-OF-THE-ART METHODS

We submitted our results to the official website of KITTI. As compared in Table 1, our BiProDet outperforms existing state-of-the-art methods on the KITTI test set by a remarkable margin, i.e., an absolute increase of 2.1% mAP over the second best method EQ-PVRCNN. Notably, by the time of submission, we ranked 1st on the KITTI 3D detection benchmark for the cyclist class. Particularly, it is observed that BiProDet shows consistent and more obvious superiority on “moderate” and “hard” levels, i.e., the distant or highly-occluded objects with sparse points. We ascribe such performance gains to the fact that our proposed bidirectional feature propagation mechanism contributes to more adequate exploitation of complementary information between multi-modalities as well as the effects of 2D auxiliary tasks (as verified in Sec. 4.3). Besides, we can observe from the actual visual results (as presented in Appendix A.7) that our BiProDet is able to produce higher-quality 3D bounding boxes in varying scenes.

Table 1: Performance comparisons on the KITTI test set, where the best and the second best results are highlighted in bold and underlined, respectively.

Method	Modality	3D Car (IoU=0.7)			3D Ped. (IoU=0.5)			3D Cyc. (IoU=0.5)			mAP
		Easy	Mod.	Hard	Easy	Mod.	Hard	Easy	Mod.	Hard	
PointRCNN (Shi et al., 2019)	LiDAR	86.96	75.64	70.70	47.98	39.37	36.01	74.96	58.82	52.53	60.33
PointPillars (Lang et al., 2019)	LiDAR	82.58	74.31	68.99	51.45	41.92	38.89	77.10	58.65	51.92	60.65
TANet (Liu et al., 2020)	LiDAR	84.39	75.94	68.82	53.72	44.34	40.49	75.70	59.44	52.53	61.71
IA-SSD (Zhang et al., 2022c)	LiDAR	88.34	80.13	75.04	46.51	39.03	35.60	78.35	61.94	55.70	62.29
STD (Yang et al., 2019)	LiDAR	87.95	79.71	75.09	53.29	42.47	38.35	78.69	61.59	55.30	63.60
Point-GNN (Shi & Rajkumar, 2020a)	LiDAR	88.33	79.47	72.29	51.92	43.77	40.14	78.60	63.48	57.08	63.90
Part-A ² (Shi et al., 2020b)	LiDAR	87.81	78.49	73.51	53.10	43.35	40.06	79.17	63.52	56.93	63.99
PV-RCNN (Shi et al., 2020a)	LiDAR	<u>90.25</u>	81.43	76.82	52.17	43.29	40.29	78.60	63.71	57.65	64.91
3DSSD (Yang et al., 2020)	LiDAR	88.36	79.57	74.55	54.64	44.27	40.23	82.48	64.10	56.90	65.01
HotSpotNet (Chen et al., 2020)	LiDAR	87.60	78.31	73.34	53.10	45.37	41.47	82.59	65.95	59.00	65.19
PDV (Hu et al., 2022)	LiDAR	90.43	81.86	77.36	47.80	40.56	38.46	83.04	67.81	60.46	65.31
MV3D (Chen et al., 2017)	LiDAR+RGB	74.97	63.63	54.00	-	-	-	-	-	-	-
MMF (Liang et al., 2019)	LiDAR+RGB	88.40	77.43	70.22	-	-	-	-	-	-	-
AVOD-FPN (Ku et al., 2018)	LiDAR+RGB	83.07	71.76	65.73	50.46	42.27	39.04	63.76	50.55	44.93	56.84
F-PointNet (Qi et al., 2018)	LiDAR+RGB	82.19	69.79	60.59	50.53	42.15	38.08	72.27	56.12	49.01	57.86
PointPainting (Vora et al., 2020)	LiDAR+RGB	82.11	71.70	67.08	50.32	40.97	37.84	77.63	63.78	55.89	60.81
F-ConvNet (Wang & Jia, 2019)	LiDAR+RGB	87.36	76.39	66.69	52.16	43.38	38.80	81.98	65.07	56.54	63.15
CAT-Det (Zhang et al., 2022b)	LiDAR+RGB	89.87	81.32	76.68	54.26	45.44	41.94	83.68	68.81	61.45	67.05
EQ-PVRCNN (Yang et al., 2022)	LiDAR+RGB	90.13	<u>82.01</u>	<u>77.53</u>	55.84	47.02	42.94	85.41	69.10	62.30	68.03
BiProDet (Ours)	LiDAR+RGB	89.13	82.97	80.05	<u>55.59</u>	48.77	46.12	86.74	74.32	67.45	70.13

Table 2: Ablative experiments on different feature exploitation and propagation schemes.

Exp.	Method	3D Car (IoU=0.7)			3D Ped. (IoU=0.5)			3D Cyc. (IoU=0.5)			mAP
		Easy	Mod.	Hard	Easy	Mod.	Hard	Easy	Mod.	Hard	
(a)	Single-Modal	91.92	85.22	82.98	68.82	61.47	56.39	91.93	74.56	69.58	75.88
(b)	Input Fusion	92.18	85.68	83.43	71.23	64.78	60.27	91.81	74.52	69.69	77.07
(c)	Point-to-pixel	91.83	85.11	82.91	70.66	63.57	58.19	92.78	76.31	71.77	77.01
(d)	Pixel-to-point	92.07	85.79	82.95	69.46	65.53	59.56	94.39	75.24	72.23	77.47
(e)	Bidirectional	92.63	85.77	83.13	72.68	67.64	62.25	94.39	77.77	71.47	78.64

4.3 ABLATION STUDY

We conducted comprehensive ablation studies to validate the effectiveness and explore the impacts of key modules involved in the overall learning framework.

Effectiveness of Bidirectional Propagation. We performed detailed ablation studies on specific multi-modal feature exploitation and interaction strategies. We started by presenting a single-modal baseline (Table 2 (a)) that only preserves the point cloud branch of our BiProDet for both training and inference. Based on our complete bidirectional propagation pipeline (Table 2 (e)), we explored another two variants as shown in (Table 2 (c)) and (Table 2 (d)), solely preserving the point-to-pixel and pixel-to-point feature propagation in our 3D detectors, respectively. Note that in Table 2 (c) the point-to-pixel feature flow was only enabled during training, and we detached the point-to-pixel module as well as the whole image branch during inference. In addition, in Table 2 (b), we replaced the bidirectional feature propagation of our BiProDet with the input fusion strategy as proposed in (Wang et al., 2021) that decorates point clouds with CNN features deduced from the image branch. Empirically, we can draw several aspects of conclusions that strongly support our preceding claims. *First*, the performance of Table 2 (a) turns out to be the worst among all variants, which reveals the superiority of cross-modal learning. *Second*, combining Table 2 (a) and Table 2 (c), the mAP is largely boosted from 75.88% to 77.01%. Considering that, during the inference stage, these two variants have identical forms of input and network structure, the resulting improvement strongly indicates that the representation ability of the 3D LiDAR branch is indeed strengthened. *Third*, comparing Table 2 (c) and Table 2 (d) with Table 2 (e), we can verify the superiority of bidirectional propagation (78.64%) over the unidirectional schemes (77.01% and 77.47%). If we particularly pay attention to Table 2 (d) and Table 2 (e), we can conclude that our newly proposed point-to-pixel feature propagation direction further brings 1.17% mAP increase based on the previously explored pixel-to-point paradigm. *Besides*, by comparing Table 2 (b) (77.07%) and Table 2 (d) (77.47%) whose information propagation directions are both from the 2D image domain to the 3D point cloud domain, we can demonstrate the superiority of feature-level propagation over its input-level counterpart (Wang et al., 2021).

Effectiveness of Image Branch. Comparing Table 3 (d) with Tables 3 (e)-(g), the performance is stably boosted from 75.88% to 78.64%, as the gradual addition of the 2D image branch and two

Table 3: Ablative experiments on key modules of our BiProDet. LiDAR: the point cloud branch trained without auxiliary tasks; Img: the 2D RGB image backbone; 3D Seg: 3D semantic segmentation; 3D Ctr: 3D center estimation; 2D Seg: 2D semantic segmentation; NLC: 2D NLC map estimation.

Exp.	LiDAR	3D Seg	3D Ctr	Img	2D Seg	NLC	Car	Ped.	Cyc.	mAP
(a)	✓						87.40	59.02	79.11	75.18
(b)	✓	✓					86.75	59.90	78.57	75.07
(c)	✓		✓				86.56	61.06	78.58	75.40
(d)	✓	✓	✓				86.71	62.23	78.68	75.88
(e)	✓	✓	✓	✓			86.91	63.23	79.00	76.38
(f)	✓	✓	✓	✓	✓		87.10	64.52	79.94	77.18
(g)	✓	✓	✓	✓	✓	✓	87.18	67.52	81.21	78.64

auxiliary tasks. These results indicate that the additional semantic and geometric features learned from the image modality are indeed effective supplements to the point cloud representation, leading to significantly improved 3D object detection performances.

Effectiveness of Auxiliary Tasks. Comparing Table 3 (a) with Tables 3 (b)-(d), the mAP is boosted by 0.70% when incorporating 3D semantic segmentation and 3D center estimation, and can be further improved by 2.76% after introducing 2D semantic segmentation and NLC map estimation (Tables 3 (d)-(g)). However, only integrating image features without 2D auxiliary tasks (comparing results of Tables 3 (d) and (e)) brings limited improvement of 0.5%. This observation shows that the 2D auxiliary tasks, especially the proposed 2D NLC map estimation, do enforce the image branch to learn complementary information for the detector.

Robustness against Input Corruption. We also conducted extensive experiments to verify the robustness of our BiProDet to sensor perturbation. Specifically, we added Gaussian noises to the reflectance value of points or RGB images. Fig. 5 shows that the mAP value of our cross-modal BiProDet is consistently higher than that of the single-modal baseline and decreases slower with the LiDAR noise level increasing. Particularly, as listed in Table 4, when the variance of the LiDAR noise is set to 0.15, the perturbation affects our cross-modal BiProDet much less than the single-modal detector. Besides, even applying the corruption to both LiDAR input and RGB images, the mAP value of our BiProDet only drops by 2.49%.

Table 4: Performances (mAPs) of the single-modal baseline and our BiProDet on the KITTI val set under input corruptions of simulated LiDAR and image noise sampled from the Gaussian distribution. Note that the image noise is only applicable to multi-modal detectors.

Corruptions Type	Modal	Car	Ped.	Cyc.	mAP
No Corruption	LiDAR	86.71	62.23	78.68	75.87
LiDAR Noise	LiDAR	84.37	49.19	77.27	70.28
No Corruption	LiDAR + RGB	87.18	67.52	81.21	78.64
LiDAR Noise	LiDAR + RGB	86.82	65.61	77.63	76.69
Image Noise	LiDAR + RGB	87.14	66.26	77.97	77.12
LiDAR + Image Noise	LiDAR + RGB	86.60	65.42	76.44	76.15




Figure 5: Comparisons of noise robustness between the single-modal baseline and our BiProDet.

5 CONCLUSION

We have presented a novel cross-modal 3D detector, namely BiProDet, by fully exploiting complementary information from the image domain in two aspects. First, we proposed point-to-pixel feature propagation, which enables gradients backpropagated from the training loss of the image branch to augment the expressive power of the 3D point cloud backbone. Second, we proposed NLC map estimation as an auxiliary task to promote the image branch to learn local spatial-aware representation rather than just semantic features. Our BiProDet achieves state-of-the-art results on the KITTI detection benchmark. And extensive ablative experiments demonstrate the robustness of our BiProDet against sensor noises and generalization to LiDAR signals with fewer beams. The decent results demonstrate the potential of joint training between 3D object detection and more 2D scene understanding tasks. We believe our new perspective will further inspire investigations on multi-task multi-modal learning for scene understanding in autonomous driving.

ACKNOWLEDGEMENT

This work was supported in part by the Hong Kong Research Grants Council under Grants 11219422, 11202320, and 11218121, and in part by Hong Kong Innovation and Technology Fund under Grant MHP/117/21.

REPRODUCIBILITY STATEMENT

We provide the source code and configuration in the supplementary material for the experimental results presented in the main text. We specify the settings of hyper-parameters, the training scheme, and the implementation details of our method in Section 4.1 and Appendix A.2. Besides, we also give a clear explanation of used datasets. We repeat the experiments on the KITTI dataset several times to check the rightness and reproducibility of the implementation.

ETHICS STATEMENT

Proposing effective cross-modal 3D object detectors would benefit a wide range of applications. For example, the incomplete information obtained from the single-modal sensor in autonomous driving vehicles may lead to serious traffic accidents. And the cross-modal detectors are expected to achieve more accurate detection results and be more robust to environment perturbation. Besides, this study was conducted using publicly available datasets and without surveying participants for experiments. Thus we do not consider this study to raise privacy or data security issues. We cited the creators when using existing datasets.

REFERENCES

- Xuyang Bai, Zeyu Hu, Xinge Zhu, Qingqiu Huang, Yilun Chen, Hongbo Fu, and Chiew-Lan Tai. Transfusion: Robust lidar-camera fusion for 3d object detection with transformers. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 1090–1099, 2022.
- Chen Chen, Zhe Chen, Jing Zhang, and Dacheng Tao. Sasa: Semantics-augmented set abstraction for point-based 3d object detection. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 1, pp. 221–229, 2022.
- Hansheng Chen, Yuyao Huang, Wei Tian, Zhong Gao, and Lu Xiong. Monorun: Monocular 3d object detection by reconstruction and uncertainty propagation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 10379–10388, 2021.
- Qi Chen, Lin Sun, Zhixin Wang, Kui Jia, and Alan Yuille. Object as hotspots: An anchor-free 3d object detection approach via firing of hotspots. In *European Conference on Computer Vision*, pp. 68–84. Springer, 2020.
- Xiaozhi Chen, Huimin Ma, Ji Wan, Bo Li, and Tian Xia. Multi-view 3d object detection network for autonomous driving. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 1907–1915, 2017.
- Jiajun Deng, Shaoshuai Shi, Peiwei Li, Wengang Zhou, Yanyong Zhang, and Houqiang Li. Voxel r-cnn: Towards high performance voxel-based 3d object detection. In *Proceedings of the AAAI Conference on Artificial Intelligence*, pp. 1201–1209, 2021.
- Andreas Geiger, Philip Lenz, and Raquel Urtasun. Are we ready for autonomous driving? the kitti vision benchmark suite. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 3354–3361, 2012.
- C. He, H. Zeng, J. Huang, X.-S. Hua, and L. Zhang. Structure aware single-stage 3d object detection from point cloud. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 11870–11879, 2020a.

- Chenhang He, Hui Zeng, Jianqiang Huang, Xian-Sheng Hua, and Lei Zhang. Structure aware single-stage 3d object detection from point cloud. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 11873–11882, 2020b.
- Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 770–778, 2016.
- Jordan SK Hu, Tianshu Kuai, and Steven L Waslander. Point density-aware voxels for lidar 3d object detection. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 8469–8478, 2022.
- Tengteng Huang, Zhe Liu, Xiwu Chen, and Xiang Bai. Epnet: Enhancing point features with image semantics for 3d object detection. In *European Conference on Computer Vision*, pp. 35–52. Springer, 2020.
- D.P. Kingma and J.L. Ba. Adam: A method for stochastic optimization. In *International Conference on Learning Representations*, pp. 1–15, 2015.
- Jason Ku, Melissa Mozifian, Jungwook Lee, Ali Harakeh, and Steven L Waslander. Joint 3d proposal generation and object detection from view aggregation. In *2018 IEEE/RSJ International Conference on Intelligent Robots and Systems*, pp. 1–8. IEEE, 2018.
- A.H. Lang, S. Vora, H. Caesar, L. Zhou, J. Yang, and O. Beijbom. Pointpillars: Fast encoders for object detection from point clouds. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 12689–12697, 2019.
- Yinhao Li, Han Bao, Zheng Ge, Jinrong Yang, Jianjian Sun, and Zeming Li. Bevstereo: Enhancing depth estimation in multi-view 3d object detection with dynamic temporal stereo. *arXiv preprint arXiv:2209.10248*, 2022.
- Ming Liang, Bin Yang, Shenlong Wang, and Raquel Urtasun. Deep continuous fusion for multi-sensor 3d object detection. In *European Conference on Computer Vision*, pp. 641–656, 2018.
- Ming Liang, Bin Yang, Yun Chen, Rui Hu, and Raquel Urtasun. Multi-task multi-sensor fusion for 3d object detection. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 7345–7353, 2019.
- Tingting Liang, Hongwei Xie, Kaicheng Yu, Zhongyu Xia, Zhiwei Lin, Yongtao Wang, Tao Tang, Bing Wang, and Zhi Tang. Bevfusion: A simple and robust lidar-camera fusion framework. *arXiv preprint arXiv:2205.13790*, 2022.
- Haisong Liu, Tao Lu, Yihui Xu, Jia Liu, Wenjie Li, and Lijun Chen. Camliflow: bidirectional camera-lidar fusion for joint optical flow and scene flow estimation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 5791–5801, 2022.
- Zhe Liu, Xin Zhao, Tengpeng Huang, Ruolan Hu, Yu Zhou, and Xiang Bai. Tanet: Robust 3d object detection from point clouds with triple attention. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 34, pp. 11677–11684, 2020.
- Zhe Liu, Bingling Li, Xiwu Chen, Xi Wang, Xiang Bai, et al. Epnet++: Cascade bi-directional fusion for multi-modal 3d object detection. *arXiv preprint arXiv:2112.11088*, 2021.
- Fabian Manhardt, Wadim Kehl, and Adrien Gaidon. Roi-10d: Monocular lifting of 2d detection to 6d pose and metric shape. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 2069–2078, 2019.
- Arsalan Mousavian, Dragomir Anguelov, John Flynn, and Jana Kosecka. 3d bounding box estimation using deep learning and geometry. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 7074–7082, 2017.
- Su Pang, Daniel Morris, and Hayder Radha. Clocs: Camera-lidar object candidates fusion for 3d object detection. In *2020 IEEE/RSJ International Conference on Intelligent Robots and Systems*, pp. 10386–10393. IEEE, 2020.

- Jonah Philion and Sanja Fidler. Lift, splat, shoot: Encoding images from arbitrary camera rigs by implicitly unprojecting to 3d. In *European Conference on Computer Vision*, pp. 194–210. Springer, 2020.
- Charles R Qi, Wei Liu, Chenxia Wu, Hao Su, and Leonidas J Guibas. Frustum pointnets for 3d object detection from rgb-d data. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 918–927, 2018.
- Lu Qi, Li Jiang, Shu Liu, Xiaoyong Shen, and Jiaya Jia. Amodal instance segmentation with kins dataset. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 3014–3023, 2019.
- Rui Qian, Divyansh Garg, Yan Wang, Yurong You, Serge Belongie, Bharath Hariharan, Mark Campbell, Kilian Q Weinberger, and Wei-Lun Chao. End-to-end pseudo-lidar for image-based 3d object detection. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 5881–5890, 2020.
- Shaoqing Ren, Kaiming He, Ross Girshick, and Jian Sun. Faster r-cnn: Towards real-time object detection with region proposal networks. *Advances in Neural Information Processing Systems*, 28:91–99, 2015.
- S. Shi, X. Wang, and H. Li. Pointcnn: 3d object proposal generation and detection from point cloud. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 770–779, 2019.
- Shaoshuai Shi, Chaoxu Guo, Li Jiang, Zhe Wang, Jianping Shi, Xiaogang Wang, and Hongsheng Li. Pv-rcnn: Point-voxel feature set abstraction for 3d object detection. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 10529–10538, 2020a.
- Shaoshuai Shi, Zhe Wang, Jianping Shi, Xiaogang Wang, and Hongsheng Li. From points to parts: 3d object detection from point cloud with part-aware and part-aggregation network. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 43(8):2647–2664, 2020b.
- W. Shi and R. Rajkumar. Point-gnn: Graph neural network for 3d object detection in a point cloud. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 1708–1716, 2020a.
- Weijing Shi and Raj Rajkumar. Point-gnn: Graph neural network for 3d object detection in a point cloud. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 1711–1719, 2020b.
- Leslie N Smith. Cyclical learning rates for training neural networks. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision*, pp. 464–472. IEEE, 2017.
- P. Sun, H. Kretzschmar, X. Dotiwalla, A. Chouard, V. Patnaik, P. Tsui, J. Guo, Y. Zhou, Y. Chai, B. Caine, V. Vasudevan, W. Han, J. Ngiam, H. Zhao, A. Timofeev, S. Ettinger, M. Krivokon, A. Gao, A. Joshi, Y. Zhang, J. Shlens, Z. Chen, and D. Anguelov. Scalability in perception for autonomous driving: Waymo open dataset. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 2443–2451, 2020.
- Zhi Tian, Chunhua Shen, Hao Chen, and Tong He. Fcos: Fully convolutional one-stage object detection. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pp. 9627–9636, 2019.
- Sourabh Vora, Alex H Lang, Bassam Helou, and Oscar Beijbom. Pointpainting: Sequential fusion for 3d object detection. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 4604–4612, 2020.
- Chunwei Wang, Chao Ma, Ming Zhu, and Xiaokang Yang. Pointaugmenting: Cross-modal augmentation for 3d object detection. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 11794–11803, 2021.

- Yan Wang, Wei-Lun Chao, Divyansh Garg, Bharath Hariharan, Mark Campbell, and Kilian Q Weinberger. Pseudo-lidar from visual depth estimation: Bridging the gap in 3d object detection for autonomous driving. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 8445–8453, 2019.
- Zhixin Wang and Kui Jia. Frustum convnet: Sliding frustums to aggregate local point-wise features for amodal 3d object detection. In *2019 IEEE/RSJ International Conference on Intelligent Robots and Systems*, pp. 1742–1749. IEEE, 2019.
- Xiaopei Wu, Liang Peng, Honghui Yang, Liang Xie, Chenxi Huang, Chengqi Deng, Haifeng Liu, and Deng Cai. Sparse fuse dense: Towards high quality 3d detection with depth completion. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 5418–5427, June 2022.
- Danfei Xu, Dragomir Anguelov, and Ashesh Jain. Pointfusion: Deep sensor fusion for 3d bounding box estimation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 244–253, 2018.
- Yan Yan, Yuxing Mao, and Bo Li. Second: Sparsely embedded convolutional detection. *Sensors*, 18(10):3337, 2018.
- Zetong Yang, Yanan Sun, Shu Liu, Xiaoyong Shen, and Jiaya Jia. Std: Sparse-to-dense 3d object detector for point cloud. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pp. 1951–1960, 2019.
- Zetong Yang, Yanan Sun, Shu Liu, and Jiaya Jia. 3dssd: Point-based 3d single stage object detector. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 11040–11048, 2020.
- Zetong Yang, Li Jiang, Yanan Sun, Bernt Schiele, and Jiaya Jia. A unified query-based paradigm for point cloud understanding. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 8541–8551, 2022.
- Jin Hyeok Yoo, Yecheol Kim, Jisong Kim, and Jun Won Choi. 3d-cvf: Generating joint camera and lidar features using cross-view spatial feature fusion for 3d object detection. In *European Conference on Computer Vision*, pp. 720–736. Springer, 2020.
- Yurong You, Yan Wang, Wei-Lun Chao, Divyansh Garg, Geoff Pleiss, Bharath Hariharan, Mark Campbell, and Kilian Q. Weinberger. Pseudo-lidar++: Accurate depth for 3d object detection in autonomous driving. In *International Conference on Learning Representations*, pp. 1–13, 2020.
- Qijian Zhang, Junhui Hou, and Yue Qian. Multi-view vision-to-geometry knowledge transfer for 3d point cloud shape analysis. *arXiv preprint arXiv:2207.03128*, 2022a.
- Yanan Zhang, Jiaxin Chen, and Di Huang. Cat-det: Contrastively augmented transformer for multi-modal 3d object detection. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 908–917, 2022b.
- Yifan Zhang, Qingyong Hu, Guoquan Xu, Yanxin Ma, Jianwei Wan, and Yulan Guo. Not all points are equal: Learning highly efficient point-based detectors for 3d lidar point clouds. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 18953–18962, 2022c.
- Yifan Zhang, Qijian Zhang, Zhiyu Zhu, Junhui Hou, and Yixuan Yuan. Glenet: Boosting 3d object detectors with generative label uncertainty estimation. *arXiv preprint arXiv:2207.02466*, 2022d.
- Hengshuang Zhao, Jianping Shi, Xiaojuan Qi, Xiaogang Wang, and Jiaya Jia. Pyramid scene parsing network. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 2881–2890, 2017.
- Wu Zheng, Weiliang Tang, Li Jiang, and Chi-Wing Fu. Se-ssd: Self-ensembling single-stage object detector from point cloud. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 14494–14503, 2021.

Y. Zhou and O. Tuzel. Voxnet: End-to-end learning for point cloud based 3d object detection. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 4490–4499, 2018.

A APPENDIX

In this appendix, we provide the details omitted from the manuscript due to space limitation. We organize the appendix as follows.

- Section A.1: Details of the normalized local coordinate (NLC) System.
- Section A.2: Implementation details.
- Section A.3: More quantitative results.
- Section A.4: Experimental results on Waymo dataset.
- Section A.5: More ablation studies.
- Section A.6: Efficiency analysis.
- Section A.7: Visual results of 3D object detection.
- Section A.8: Visual results of 2D semantic segmentation.
- Section A.9: Details on the official KITTI test leaderboard.

A.1 DETAILS OF NORMALIZED LOCAL COORDINATE (NLC) SYSTEM

The aim of the cross-modal 3D detector is to estimate the bounding box parameterized by the object dimension (w, l, h) , the center location (x_c, y_c, z_c) , and the orientation angle θ . The input contains an RGB image $\mathbf{X} \in \mathbb{R}^{3 \times H \times W}$, where H and W represent the height and width of the image respectively, and a 3D LiDAR point cloud $\{\mathbf{p}_i | i = 1, \dots, N\}$, where N is the number of points, and each point $\mathbf{p}_i \in \mathbb{R}^4$ is represented with its 3D location (x_p, y_p, z_p) in the LiDAR coordinate system and the reflectance value ρ .

To build the correspondence between the 2D and 3D modalities, we project the observed points from the 3D coordinate system to the 2D coordinate system on the image plane:

$$d_c \begin{bmatrix} u \\ v \\ 1 \end{bmatrix} = \mathbf{K} [\mathbf{R} \quad \mathbf{T}] \begin{bmatrix} x_p \\ y_p \\ z_p \\ 1 \end{bmatrix}, \quad (8)$$

where u, v, d_c denote the corresponding coordinates and depth on the image plane, $\mathbf{R} \in \mathbb{R}^{3 \times 3}$ and $\mathbf{T} \in \mathbb{R}^{3 \times 1}$ denote the rotation matrix and translation matrix of the LiDAR relative to the camera, and $\mathbf{K} \in \mathbb{R}^{3 \times 3}$ is the camera intrinsic matrix.

Given the point cloud and the bounding box of an object, LiDAR coordinates of foreground points can be transformed into the NLC system by proper translation, rotation, and scaling, i.e.,

$$\begin{bmatrix} x_p^{NLC} \\ y_p^{NLC} \\ z_p^{NLC} \end{bmatrix} = \begin{bmatrix} 1/l & 0 & 0 \\ 0 & 1/w & 0 \\ 0 & 0 & 1/h \end{bmatrix} \cdot \begin{bmatrix} x_p^{LCS} \\ y_p^{LCS} \\ z_p^{LCS} \end{bmatrix} + \begin{bmatrix} 0.5 \\ 0.5 \\ 0.5 \end{bmatrix} = \begin{bmatrix} 1/l & 0 & 0 \\ 0 & 1/w & 0 \\ 0 & 0 & 1/h \end{bmatrix} \cdot \begin{bmatrix} \cos\theta & -\sin\theta & 0 \\ \sin\theta & \cos\theta & 0 \\ 0 & 0 & 1 \end{bmatrix} \cdot \begin{bmatrix} x_p - x_c \\ y_p - y_c \\ z_p - z_c \end{bmatrix} + \begin{bmatrix} 0.5 \\ 0.5 \\ 0.5 \end{bmatrix}, \quad (9)$$

where $(x_p^{NLC}, y_p^{NLC}, z_p^{NLC})$ and $(x_p^{LCS}, y_p^{LCS}, z_p^{LCS})$ denote the coordinates of points in NLC system and local coordinate system, respectively. Suppose that the bounding box is unknown with known global coordinates and NLCs of points, we can build a set of equations about the parameters of box (i.e., the $x_c, y_c, z_c, w, l, h, \theta$). In order to solve the total of 7 parameters, global coordinates and NLCs of at least 7 points are required to construct the equations.

Discussion. Notably, it is difficult to infer NLCs of points with only LiDAR data for objects far away or with low reflectivity, since the point clouds are sparse. Meanwhile, the contours and appearance of these cases are still visible in RGB images, so we propose to estimate the normalized local coordinates from images. Then, we can retrieve the estimated NLCs of observed points based on the 2D-3D correspondence. Finally, we expect the proposed cross-modal detector to achieve higher detection accuracy with estimated NLCs and known global coordinates.

A.2 IMPLEMENTATION DETAILS

Network Architecture. Fig. 6 illustrates the architectures of the point cloud and image backbone networks. For the encoder of the point cloud branch, we further show the details of multi-scale

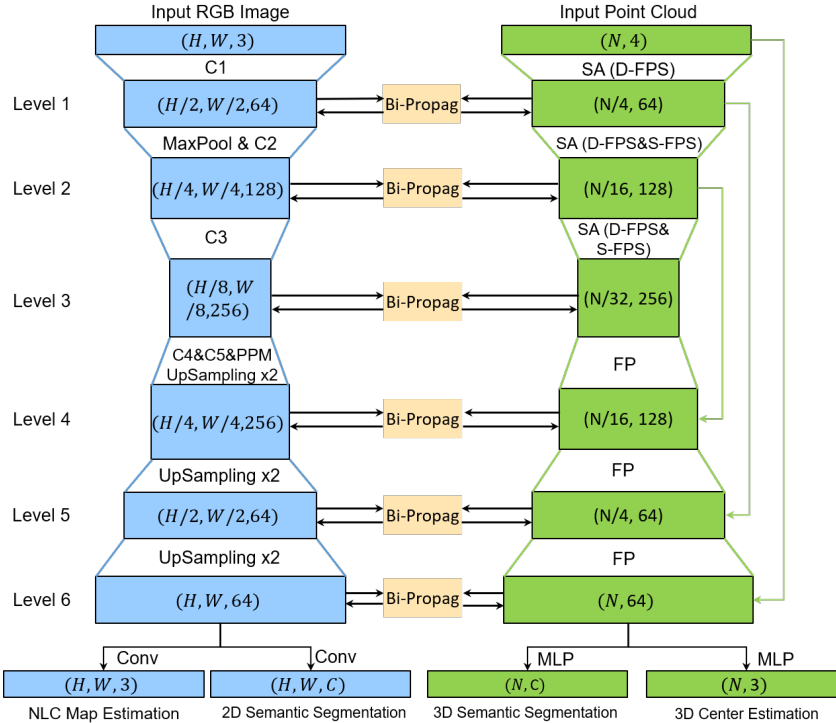


Figure 6: The detailed architecture of 2D and 3D backbones. We adopt ResNet18 as the encoder of the image branch, followed by a decoder with Pyramid Pooling Module (PPM) and several up-sampling blocks. C1, C2, C3, C4, and C5 denote convolutional layers of different stages in ResNet. Extra convolutional layers are deployed after each up-sampling layer. For the point cloud branch, we adopt the PointNet++ structure. SA: set abstraction layer, D-FPS: 3D Euclidean distance-based farthest point sampling, S-FPS: semantic-guided farthest point sampling, FP: feature propagation layer, MLP: shared multi-layer perceptron. Besides, “Bi-Propag” denotes the proposed bidirectional feature propagation between the 2D and 3D backbones.

Table 5: Details of set abstraction layers in the point cloud branch. We report the sampling strategy used in the sampling operation, ball radius of group operation, “nquery” that denotes the number of group points, and dimensions of the unit PointNet layer for multi-scale grouping. The features at different scales are concatenated and dimensionally reduced to the specific output channels.

Layer	Sampling Strategy	Radius	nquery	Feature Dimension	Output Channels
1 st SA	D-FPS	[0.2, 0.4, 0.8]	[32, 32, 64]	[[16, 16, 32], [16, 16, 32], [32, 32, 64]]	64
2 nd SA	D-FPS & S-FPS	[0.4, 0.8, 1.6]	[32, 32, 64]	[[64, 64, 128], [64, 64, 128], [64, 96, 128]]	128
3 rd SA	D-FPS & S-FPS	[1.6, 3.2, 4.8]	[64, 64, 128]	[[128, 128, 256], [128, 196, 256], [128, 256, 256]]	256

grouping (MSG) network in Table 5. Following 3DSSD (Yang et al., 2020), we take key points sampled by the 3rd SA layer to generate vote points and estimate the 3D box. Then we feed these 3D boxes and features output by the last FP layer to the refinement stage. Besides, we adopt density-aware RoI grid pooling (Hu et al., 2022) to encode point density as an additional feature. Note that 3D center estimation (He et al., 2020b) aims to learn the relative position of each foreground point to the object center, while the 3D box is estimated based on sub-sampled points. Thus, the auxiliary task of 3D center estimation differs from 3D box estimation and can facilitate learning structure-aware features of objects.

Training Details. Through the experiments on KITTI dataset, we adopted Adam (Kingma & Ba, 2015) ($\beta_1=0.9$, $\beta_2=0.99$) to optimize our BiProDet. We initialized the learning rate as 0.003 and updated it with the one-cycle policy (Smith, 2017). And we trained the model for a total of 80 epochs in an end-to-end manner. In our experiments, the batch size was set to 8, equally distributed on 4 NVIDIA 3090 GPUs. We kept the input image with the original resolution and padded it to

Table 6: Quantitative comparisons on the KITTI validation set under the evaluation metric of 3D Average Precision (AP) calculated with 11 sampling recall positions. We highlight the best and the second best results in bold and underlined, respectively.

Method	Modality	3D Car (IoU=0.7)			3D Ped. (IoU=0.5)			3D Cyc. (IoU=0.5)			mAP
		Easy	Mod.	Hard	Easy	Mod.	Hard	Easy	Mod.	Hard	
PointPillars (Lang et al., 2019)	LiDAR	86.46	77.28	74.65	57.75	52.29	47.90	80.05	62.68	59.70	66.53
SECOND (Yan et al., 2018)	LiDAR	88.61	78.62	77.22	56.55	52.98	47.73	80.58	67.15	63.10	68.06
3DSSD (Yang et al., 2020)	LiDAR	88.55	78.45	77.30	58.18	54.31	49.56	86.25	70.48	65.32	69.82
PointRCNN (Shi et al., 2019)	LiDAR	88.72	78.61	77.82	62.72	53.85	50.24	86.84	71.62	65.59	70.67
PV-RCNN (Shi et al., 2020a)	LiDAR	89.03	<u>83.24</u>	78.59	63.71	57.37	52.84	86.06	69.48	64.50	71.65
TANet (Liu et al., 2020)	LiDAR	88.21	77.85	75.62	70.80	63.45	58.22	85.98	64.95	60.40	71.72
Part-A ² (Shi et al., 2020b)	LiDAR	89.55	79.40	78.84	65.68	60.05	55.44	85.50	69.90	65.48	72.20
AVOD-FPN (Ku et al., 2018)	LiDAR+RGB	84.41	74.44	68.65	-	58.80	-	-	49.70	-	-
PointFusion (Xu et al., 2018)	LiDAR+RGB	77.92	63.00	53.27	33.36	28.04	23.38	49.34	29.42	26.98	42.75
F-PointNet (Qi et al., 2018)	LiDAR+RGB	83.76	70.92	63.65	70.00	61.32	53.59	77.15	56.49	53.37	65.58
CLOCs (Pang et al., 2020)	LiDAR+RGB	89.49	79.31	77.36	62.88	56.20	50.10	87.57	67.92	63.67	70.50
EPNet (Huang et al., 2020)	LiDAR+RGB	88.76	78.65	78.32	66.74	59.29	54.82	83.88	65.50	62.70	70.96
CAT-Det (Zhang et al., 2022b)	LiDAR+RGB	90.12	81.46	<u>79.15</u>	74.08	<u>66.35</u>	<u>58.92</u>	<u>87.64</u>	<u>72.82</u>	<u>68.20</u>	<u>75.42</u>
BiProDet (Ours)	LiDAR+RGB	<u>89.73</u>	86.40	79.31	<u>71.77</u>	68.49	62.52	89.24	76.91	75.18	77.73

Table 7: Comparison with state-of-the-art methods on the KITTI val set for car 3D detection. All results are reported by the average precision with 0.7 IoU threshold. R11 and R40 denotes AP calculated with 11 and 40 recall sampling recall points, respectively.

Method	Modal	AP _{3D R11} (%)			AP _{3D R40} (%)		
		Easy	Mod.	Hard	Easy	Mod.	Hard
Voxel R-CNN (Deng et al., 2021)	LiDAR	89.41	84.52	78.93	92.38	85.29	82.86
PV-RCNN (Shi et al., 2020a)	LiDAR	89.35	83.69	78.7	92.57	84.83	82.69
SA-SSD (He et al., 2020b)	LiDAR	90.15	79.91	78.78	93.14	84.65	81.86
SE-SSD (Zheng et al., 2021)	LiDAR	90.21	85.71	79.22	93.19	86.12	83.31
GLENet-VR (Zhang et al., 2022d)	LiDAR	89.93	86.46	79.19	93.51	86.10	83.60
MV3D (Chen et al., 2017)	LiDAR+RGB	-	-	-	71.29	62.68	56.56
3D-CVF (Yoo et al., 2020)	LiDAR+RGB	-	-	-	89.67	79.88	78.47
BiProDet (Ours)	LiDAR+RGB	89.72	86.52	79.34	93.03	86.46	83.73

the size of 1248×376 , and down-sampled the input point cloud to 16384 points during training and inference. Following the common practice, we set the detection range of the x , y , and z axis to $[0m, 70.4m]$, $[-40m, 40m]$ and $[-3m, 1m]$, respectively.

Data Augmentation. We applied common data augmentation strategies at global and object levels. The global-level augmentation includes random global flipping, global scaling with a random scaling factor between 0.95 and 1.05, and global rotation around the z -axis with a random angle in the range of $[-\pi/4, \pi/4]$. Each of the three augmentations was performed with a 50% probability for each sample. The object-level augmentation refers to copying objects from other scenes and pasting them to current scene (Yan et al., 2018). In order to perform sampling synchronously on point clouds and images, we utilized the instance masks provided in (Qi et al., 2019). Specifically, we pasted both the point clouds and pixels of sampled objects to the point cloud and images of new scenes, respectively.

A.3 MORE QUANTITATIVE RESULTS

Performance on KITTI Val Set. We also reported the performance of our BiProDet on all three classes of the KITTI validation set in Table 6, where it can be seen that our BiProDet also achieves the highest mAP of 77.73%, which is obviously higher than the second best method CAT-Det.

Performance of Single-Class Detector. Quite a few methods (Deng et al., 2021; Zheng et al., 2021) train models only for car detection. Empirically, the single-class detector performs better in the car class compared with multi-class detectors. Therefore, we also provided performance of BiProDet trained only for the car class, and compared it with several state-of-the-art methods in Table 7.

Performance of NLC Map Estimation. The NLC map estimation task is introduced to guide the image branch to learn local spatial-aware features. We evaluated the predicted NLC of pixels containing at least one projected point with mean absolute error (MAE) for each object:

$$\text{MAE}_q = \frac{1}{N_{obj}} \sum_i^N |q_{p_i}^{NLC} - \hat{q}_{p_i}^{NLC}| \cdot \mathbf{1}_{p_i}, q \in \{x, y, z\}, \quad (10)$$




Figure 7: Quantitative results of NLC map estimation on the KITTI val set.

Table 8: Bidirectional propagation is also effective when the 2D and 3D backbones are asymmetric. Here we adopt a single-stage detector (Chen et al., 2022) whose backbone includes only an encoder, which is asymmetric with the encoder-decoder network in the image branch.

Method	3D Car (IoU=0.7)			3D Ped. (IoU=0.5)			3D Cyc. (IoU=0.5)			mAP
	Easy	Mod.	Hard	Easy	Mod.	Hard	Easy	Mod.	Hard	
SASA	92.17	84.90	82.57	66.75	61.40	56.00	89.91	74.05	69.41	75.24
Ours (SASA)	92.11	85.67	82.99	70.52	63.38	58.16	91.58	74.81	70.22	76.61

where N_{obj} is the number of LiDAR points inside the ground-truth box of the object, $\mathbf{1}_{p_i}$ indicates that the evaluation is only calculated with foreground points inside the box, $q_{p_i}^{NLC}$ and $\hat{q}_{p_i}^{NLC}$ are the normalized local coordinates of the ground-truth and the prediction at the corresponding pixel for the foreground point. Finally, we obtained the mean value of MAE_q for all instances, namely $mMAE_q$. We report the metric over different difficulty levels for three categories on the KITTI val set at Figure 7. We can observe that, for the car class, the $mMAE$ error is only 0.0619, i.e., ± 6.19 cm error per meter. For the challenging pedestrian and cyclist categories, the error becomes larger due to the smaller size and the non-rigid shape.

Generalization to Asymmetric Backbones. As shown in Fig. 6, we originally adopted an encoder-decoder network in the LiDAR branch that is architecturally similar to the image backbone. Nevertheless, it is worth clarifying that our approach is not limited to symmetrical structures and can be generalized to different point-based backbones. Here, we replaced the 3D branch of the original framework with an efficient single-stage detector—SASA (Chen et al., 2022), using a backbone only with the encoder in the LiDAR branch, which is asymmetric with the encoder-decoder structure of the image backbone. Accordingly, the proposed bidirectional propagation is only performed between the 3D backbone and the encoder of the 2D image backbone. The experimental results are shown in Table 8. We can observe that the proposed method works well even when the two backbones are asymmetric, which demonstrates the satisfactory generalization ability of our method for different LiDAR backbones.

A.4 RESULTS ON WAYMO OPEN DATASET

The Waymo Open Dataset (Sun et al., 2020) is a large-scale dataset for 3D object detection. It contains 798 sequences (15836 frames) for training, and 202 sequences (40077 frames) for validation. According to the number of points inside the object and the difficulty of annotation, the objects are further divided into two difficulty levels: LEVEL_1 and LEVEL_2. Following common practice, we adopted the metrics of mean Average Precision (mAP) and mean Average Precision weighted by

Table 9: 3D detection results on the Waymo Open Dataset validation set. “-” denotes that the results are not reported in their papers.

Method	Vehicle L1		Vehicle L2		Pedestrian L1		Pedestrian L2		Cyclist L1		Cyclist L2	
	mAP	mAPH	mAP	mAPH	mAP	mAPH	mAP	mAPH	mAP	mAPH	mAP	mAPH
PointAugmenting	67.41	-	62.7	-	75.42	-	70.55	-	76.29	-	74.41	-
TransFusion	-	-	-	65.14	-	-	-	64.00	-	-	-	67.40
Ours	78.36	77.91	69.45	69.04	76.32	71.67	65.93	61.81	79.64	78.55	76.36	75.26

Table 10: Ablative experiments on the multi-stage manner of bidirectional propagation, where SA and FP denote applying bidirectional propagation at downsampling (in the encoder) and upsampling (in the decoder) stages of the point cloud branch, respectively.

	Stage		3D Car (IoU=0.7)			3D Ped. (IoU=0.5)			3D Cyc. (IoU=0.5)			mAP
	SA	FP	Easy	Mod.	Hard	Easy	Mod.	Hard	Easy	Mod.	Hard	
(a)	-	-	91.92	85.22	82.98	68.82	61.47	56.39	91.93	74.56	69.58	75.88
(b)	✓	-	92.67	85.86	83.40	70.94	65.29	60.35	93.60	75.60	71.04	77.64
(c)	-	✓	92.19	85.44	83.27	69.02	63.41	58.47	92.85	76.39	71.91	76.99
(d)	✓	✓	92.63	85.77	83.13	72.68	67.64	62.25	94.39	77.77	71.47	78.64

Table 11: Effect of the semantic-guided SA layer. Compared with the single-modal baseline, BiProDet can better exploit image semantics and preserve more foreground points during down-sampling.

Method	Single-Modal		BiProDet (Ours)		Improvement	
SA Layer	FG rate	Instance recall	FG rate	Instance recall	FG rate	Instance recall
Level-2	15.87	97.92	20.70	98.23	+4.83	+0.31
Level-3	29.73	97.35	38.03	97.82	+8.29	+0.47

heading accuracy (mAPH), and reported the performance on both LEVEL_1 and LEVEL_2. We set the detection range to $[-75.2\text{m}, 75.2\text{m}]$ for x and y axis, and $[-2\text{m}, 4\text{m}]$ for z axis. Following Wang et al. (2021) and Bai et al. (2022), the training on Waymo dataset consists of two stages to allow flexible augmentations. First, we only trained the LiDAR branch without image inputs and bidirectional propagation for 30 epochs. We enabled the copy-and-paste augmentation in this stage. Then, we trained the whole pipeline for another 6 epochs, during which the copy-and-paste is disabled. Note that the image semantic segmentation head is disabled, since ground-truth segmentation maps are not provided (Sun et al., 2020).

As shown in Table 9, our method achieves substantial improvement compared with previous state-of-the-arts. Particularly, unlike existing approaches including PointAugmenting (Wang et al., 2021) and TransFusion (Bai et al., 2022) where the camera backbone is pre-trained on other datasets and then frozen, we trained the entire pipeline in an end-to-end manner. It can be seen that even without the 2D segmentation auxiliary task, our method still achieves higher accuracy under all scenarios except “Ped L2”, demonstrating its advantage.

A.5 MORE ABLATION STUDIES

Effectiveness of Multi-stage Interaction. As mentioned before, both 2D and 3D backbones adopt an encoder-decoder structure, and we perform bidirectional feature propagation at both downsampling and upsampling stages. Here, we conducted experiments to verify the superiority of the multi-stage interaction over single-stage interaction. As shown in Table 10, only performing the bidirectional feature propagation in the encoder (i.e., Table 10 (b)) or the decoder (i.e., Table 10 (c)) leads to worse performance than that of performing the module in both stages (i.e., Table 10 (d)).

Effect of Semantic-guided Point Sampling. When performing downsampling in the SA layers of the point cloud branch, we adopted S-FPS (Chen et al., 2022) to explicitly preserve as many foreground points as possible. We report the percentage of sampled foreground points and instance recall (i.e., the ratio of instances that have at least one point) in Table 11, where it can be seen that exploiting supplementary semantic features from images leads to substantial improvement of the ratio of sampled foreground points and better instance recall during S-FPS.

Table 12: Comparison between our multi-task training methods and the single-modal 2D semantic segmentation baseline (PSPNet). The results show that point features effectively improve the segmentation performance on pedestrian and cyclist classes.

Method	Car	Pes.	Cyc.	mIoU
PSPNet (Zhao et al., 2017)	77.49	30.45	23.83	43.92
BiProDet (Ours)	78.45	36.15	30.42	48.34

Influence on 2D Semantic Segmentation. We also aimed to demonstrate that the 2D-3D joint learning paradigm benefits not only the 3D object detection task but also the 2D semantic segmentation task. As shown in Table 12, the deep interaction between different modalities yields an improvement of 4.42% mIoU. The point features can naturally complement RGB image features by providing 3D geometry and semantics, which are robust to illumination changes and help distinguish different classes of objects, for 2D visual information. The results suggest the potential of joint training between 3D object detection and more 2D scene understanding tasks in autonomous driving.

Conditional Analysis. To better figure out where the improvement comes from when using additional image features, we compared BiProDet with the single-modal detector on different occlusion levels and distant ranges. The results shown in Table 13 and Table 14 include separate APs for objects belonging to different occlusion levels and APs for moderate class in different distance ranges. For car detection, our BiProDet achieves more accuracy gains for long-distance and highly occluded objects, which suffer from the sparsity of observed LiDAR points. The cyclist and pedestrian are much more difficult categories on account of small sizes, non-rigid structures, and fewer training samples. For these two categories, BiProDet still brings consistent and significant improvements on different levels even in extremely difficult cases.

Table 13: Performance breakdown over different occlusion levels. As defined by the official website of KITTI, occlusion levels 0, 1, and 2 correspond to fully-visible samples, partly-occluded samples, and samples that are difficult to see, respectively.

Class	Car			Pedestrian			Cyclist		
	Level-0	Level-1	Level-2	Level-0	Level-1	Level-2	Level-0	Level-1	Level-2
Single-Modal	91.98	77.18	55.41	67.44	26.76	6.10	91.23	24.66	1.74
BiProDet (Ours)	92.26	77.44	58.39	74.00	35.13	7.99	92.89	30.02	2.53
<i>Improvement</i>	+0.28	+0.26	+2.97	+6.56	+8.38	+1.89	+1.66	+5.36	+0.79

Table 14: Performance breakdown over different distances.

Class	Car			Pedestrian			Cyclist		
	0-20m	20-40m	40m-Inf	0-20m	20-40m	40m-Inf	0-20m	20-40m	40m-Inf
Single-Modal	96.28	85.21	43.91	71.28	38.42	1.63	93.61	61.56	34.48
BiProDet (Ours)	96.36	86.48	49.88	76.56	45.72	2.46	94.12	67.27	39.10
<i>Improvement</i>	+0.07	+1.27	+5.97	+5.28	+7.30	+0.83	+0.51	+5.71	+4.62

Generalization to Sparse LiDAR Signals. We also compared our BiProDet with the single-modal baseline on LiDAR point clouds with various sparsity. In practice, following Pseudo-LiDAR++ (You et al., 2020), we simulated the 32-beam, 16-beam, and 8-beam LiDAR signals by selecting LiDAR points whose elevation angles fall within specific intervals. As shown in Table 15, the proposed BiProDet outperforms the single-modal baseline under all settings. The consistent improvements suggest our method can generalize to sparser signals. Besides, the proposed BiProDet significantly performs better than the baseline in the setting of LiDAR signals with fewer beams, demonstrating the effectiveness of our method in exploiting the supplementary information in the image domain.

A.6 EFFICIENCY ANALYSIS

We also compared the inference speed and number of parameters of the proposed BiProDet with state-of-the-art cross-modal approaches in Table 16. Our BiProDet has about the same number of parameters as CAT-Det (Zhang et al., 2022b), but a much higher inference speed at 9.52 frames per second on a single GeForce RTX 2080 Ti GPU. In general, our BiProDet is inevitably slower than some single-modal detectors, but it achieves a good trade-off between speed and accuracy among cross-modal approaches.

A.7 VISUAL RESULTS OF 3D OBJECT DETECTION

In Figure 8, we present the qualitative comparison of detection results between the single-modal baseline and our BiProDet. We can observe that the proposed BiProDet shows better localization capability than the single-modal baseline in challenging cases. Besides, we also show qualitative

Table 15: Comparison with single-modal baselines under LiDAR signals with different beams, where we report $AP_{3D|R40}$ on the KITTI validation set.

LiDAR Beams	Modal	Car	Ped.	Cyc.	mAP
64	LiDAR	86.71	62.23	78.68	75.87
	LiDAR + RGB	87.18	67.52	81.21	78.64
	<i>Improvement</i>	+0.47	+5.29	+2.53	+2.76
32	LiDAR	83.49	57.83	70.82	70.71
	LiDAR + RGB	84.47	62.56	73.93	73.65
	<i>Improvement</i>	+0.98	+4.73	+3.11	+2.94
16	LiDAR	79.80	53.84	60.51	64.71
	LiDAR + RGB	80.78	59.44	67.35	69.19
	<i>Improvement</i>	+0.99	+5.60	+6.84	+4.48
8	LiDAR	64.42	22.57	43.19	43.39
	LiDAR + RGB	67.09	31.03	47.97	48.70
	<i>Improvement</i>	+2.66	+8.46	+4.78	+5.30

Table 16: Comparison of the number of network parameters, inference speed, and detection accuracy of different multi-modal methods on the KITTI test set.

Method	Params (M)	Frames per second	mAP (%)
AVOD-FPN (Ku et al., 2018)	38.07	10.00	56.84
F-PointNet (Qi et al., 2018)	12.45	6.25	57.86
EPNet (Huang et al., 2020)	16.23	5.88	-
CAT-Det (Zhang et al., 2022b)	23.21	3.33	67.05
BiProDet (Ours)	24.98	9.52	70.13

results of BiProDet on the KITTI test split in Figure 9. We can clearly observe that our BiProDet performs well in challenging cases, such as pedestrians and cyclists (with small sizes) and highly-occluded cars.

A.8 VISUAL RESULTS OF 2D SEMANTIC SEGMENTATION

Several examples are shown in Figure 10. For distant cars in the first and the second row as well as the pedestrian in the sixth row, the size of objects is small and PSPNet tends to treat them as background, while our BiProDet is able to correct such errors. In the third row, our BiProDet finds the dim cyclist missed by PSPNet. Our BiProDet also performs better for the highly occluded objects as shown in the fourth and the fifth lines. This observation shows the 3D feature representations extracted from point clouds can boost 2D semantic segmentation, since the image-based method is sensitive to illumination and can hardly handle corner cases with only single-modal inputs.

A.9 DETAILS ON OFFICIAL KITTI TEST LEADERBOARD

We submitted the results of our BiProDet to the official KITTI website, and it ranks **1st** on the 3D object detection benchmark for the cyclist class. Figure 11 shows the screenshot of the leaderboard. Figure 12 illustrates the precision-recall curves along with AP scores on different categories of the KITTI test set. The samples of the KITTI test set are quite different from those of training/validation set in terms of scenes and camera parameters, so the impressive performance of our BiProDet on the test set demonstrates it also achieves good generalization.




Figure 8: Qualitative comparison between single-modal baseline and our multi-modal BiProDet. For each comparison, from top to bottom, we have the image, detection results of single-modal baseline, and detection results of BiProDet. We use red, green, and yellow to denote the **ground-truth**, **true positive** and **false positive** bounding boxes, respectively. We highlight some objects in images with red circles, which are detected by BiProDet but missed by the single-modal method.




Figure 9: Extra qualitative results of BiProDet on the KITTI test set. The predicted bounding boxes of **car**, **pedestrian**, and **cyclist** are visualized in green, cyan, and yellow, respectively. We also show the corresponding projection of boxes on images. Best viewed in color and zoom in for more details.

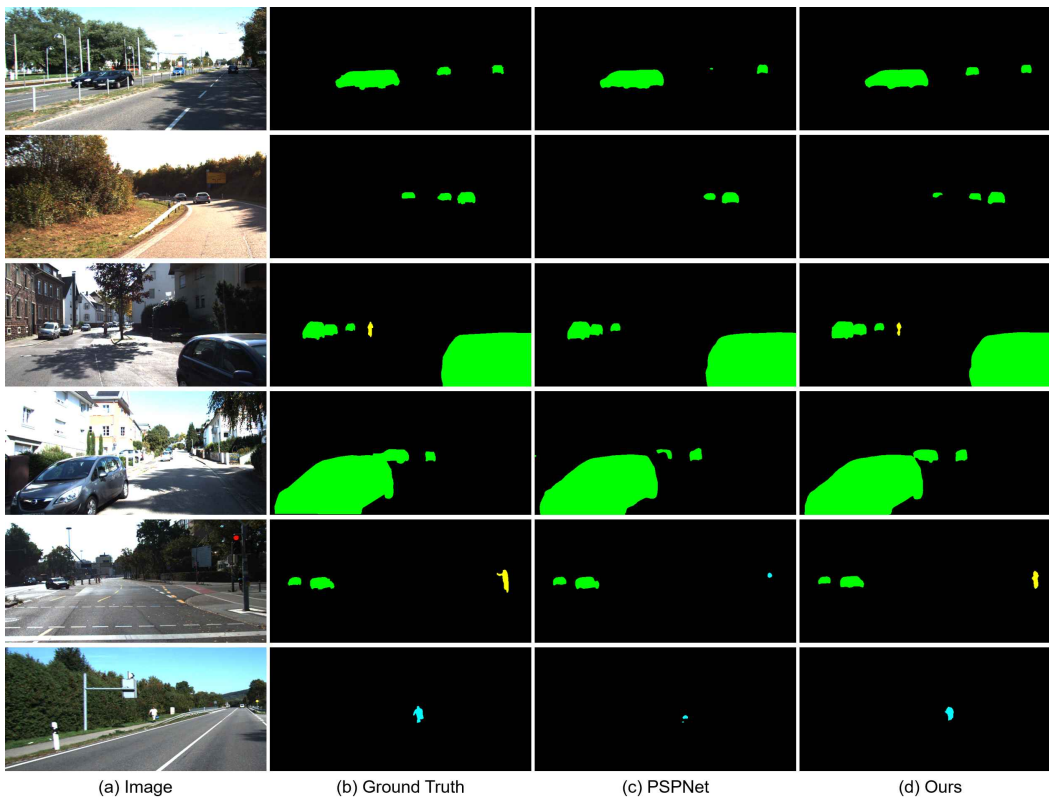
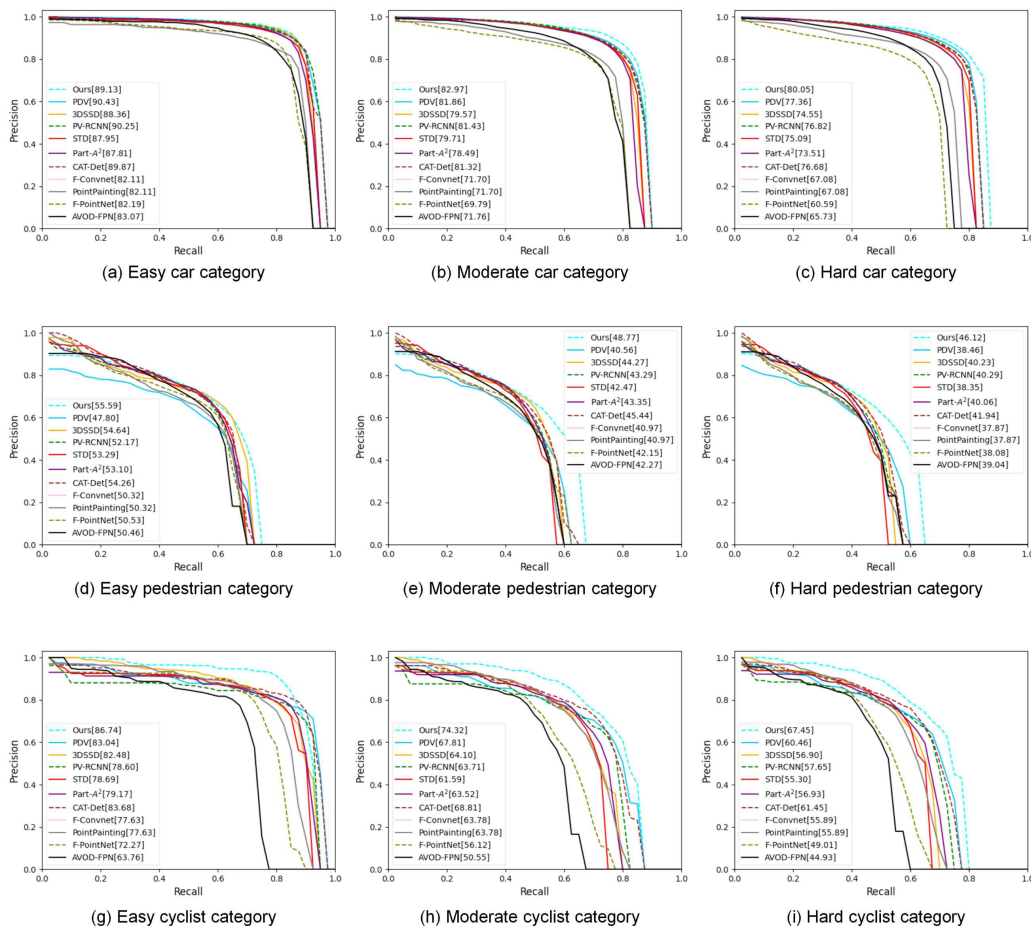


Figure 10: Visual results of 2D semantic segmentation on the KITTI val set. The prediction boxes are shown in green for **car**, cyan for **pedestrian**, and yellow for **cyclist**. Best viewed in color. Compared with PSPNet, our cross-modal BiProDet produces more accurate and detailed results.

	Method	Setting	Code	Moderate	Easy	Hard	Runtime	Environment
1	BiProDet			74.32 %	86.74 %	67.45 %	0.1 s	GPU @ 2.5 Ghz (Python + C/C++)
2	TED			74.12 %	88.82 %	66.84 %	0.1 s	1 core @ 2.5 Ghz (C/C++)
3	CasA++		code	73.79 %	87.76 %	66.84 %	0.1 s	1 core @ 2.5 Ghz (C/C++)
4	CasA		code	73.47 %	87.91 %	66.17 %	0.1 s	1 core @ 2.5 Ghz (C/C++)
5	SGNet			70.40 %	86.75 %	62.73 %	0.09 s	GPU @ 2.5 Ghz (Python)
6	HMFI		code	70.37 %	84.02 %	62.57 %	0.1 s	1 core @ 2.5 Ghz (C/C++)
7	CAD			69.94 %	84.68 %	62.21 %	0.1 s	GPU @ 2.5 Ghz (Python + C/C++)
8	SARFE			69.67 %	84.88 %	62.26 %	0.03 s	1 core @ 2.5 Ghz (C/C++)
9	EQ-PVRCNN		code	69.10 %	85.41 %	62.30 %	0.2 s	GPU @ 2.5 Ghz (Python + C/C++)
10	VoCo			69.00 %	82.74 %	62.46 %	0.1 s	1 core @ 2.5 Ghz (Python + C/C++)

Figure 11: Screenshot of the KITTI 3D object detection benchmark for cyclist class on August 15th, 2022.

Figure 12: Precision-recall curves of different methods on the KITTI 3D object detection test set on Aug. 15th, 2022. We also report APs in different categories for each method.