ORA3D: Overlap Region Aware Multi-view 3D Object Detection

Wonseok Roh¹ paulroh@korea.ac.kr Gyusam Chang¹ gsjang95@korea.ac.kr Seokha Moon³ shmoon96@korea.ac.kr Giljoo Nam² namgiljoo@gmail.com Chanyoung Kim¹ kochanha@gmail.com Younghyun Kim⁴ yhkim84@hyundai.com Sangpil Kim1* spk7@korea.ac.kr Jinkyu Kim^{3*} jinkyukim@korea.ac.kr

- ¹ Department of Artificial Intelligence, Korea University, Seoul, Republic of Korea
- ² School of Computing, KAIST, Daejeon, Republic of Korea
- ³ Department of Computer Science and Engineering, Korea University, Seoul, Republic of Korea
- ⁴ Autonomous Driving Center, Hyundai Motor Company R&D Division, Seoul, Republic of Korea

Abstract

Current multi-view 3D object detection methods often fail to detect objects in the overlap region properly, and the networks' understanding of the scene is often limited to that of a monocular detection network. Moreover, objects in the overlap region are often largely occluded or suffer from deformation due to camera distortion, causing a domain shift. To mitigate this issue, we propose using the following two main modules: (1) Stereo Disparity Estimation for Weak Depth Supervision and (2) Adversarial Overlap Region Discriminator. The former utilizes the traditional stereo disparity estimation method to obtain reliable disparity information from the overlap region. Given the disparity estimates as supervision, we propose regularizing the network to fully utilize the geometric potential of binocular images and improve the overall detection accuracy accordingly. Further, the latter module minimizes the representational gap between non-overlap and overlapping regions. We demonstrate the effectiveness of the proposed method with the nuScenes large-scale multi-view 3D object detection data. Our experiments show that our proposed method outperforms current state-of-the-art models, i.e., DETR3D and BEVDet.

1 Introduction

Object detection in 3D space plays a crucial role in various real-world applications, including autonomous driving systems. Existing 3D object detection methods [**5**, **2**, **21**], **26**, **27**] based

© 2022. The copyright of this document resides with its authors. It may be distributed unchanged freely in print or electronic forms.

^{*}Corresponding Authors



2

Figure 1: Examples where DETR3D [13] fails to properly detect objects in the overlapped regions (see dotted line), resulting in performance degradation due to multiple false positives. The first row represents the bird's eye view of the two scenes in the second row (see there are more false positives in the overlapped regions than in others). The second row shows overlapped areas between multi-view images used as input to the network.

on point clouds from LiDAR sensors often yield reliable results, but these methods suffer from a large budget to establish LiDAR sensors per vehicle. Further, camera-based object detection methods [116, 129, 101, 131] using monocular images are economical but, their performance is suboptimal due to insufficient depth cues. Stereo vision-based object detection methods [11, 129, 129, 121] might be an alternative option as they outperform monocular detection approaches with accurate depth estimation. Still, constraints in setting surround-view stereo vision systems need to be resolved. Recently, multi-view (and surround-view) camera systems have become an alternative balanced option as they can resolve some of the weaknesses of monocular and stereo vision systems for the 3D object detection task, potentially replacing LiDAR sensors.

Existing camera-only 3D object detection methods [29, 51, 51, 53] have mainly focused on predicting accurate depth to improve performance. Although estimating precise depth significantly impacts accuracy, it remains challenging. Primarily, there is a problem that depth may not be represented adequately on the pixel (e.g., difficulty in dealing with distant objects on pixels and depth compounding error properly). A landmark work in the camera-only multi-view 3D object detection task is DETR3D [53]. It introduces a promising multi-view detection pipeline that processes six images concurrently in an end-to-end manner, predicting all objects around simultaneously and implicitly utilizing rich information in the overlapping regions. Even though DETR3D performs reasonably well, we found that the network (without explicit guidance) does not totally use the geometric potentials of multi-view camera systems. Specifically, the network's understanding of the scene could be limited to that of a monocular detection network, resulting in multiple false positives in the overlapped regions, as shown in Figure 1. Thus, how to deal with this issue to boost detection accuracy remains a crucial problem.

As reported by Chen *et al.* [**□**], disparity supervision, which fully pilots the network by exploiting the strong association of binocular images, substantially improves detection performance. Inspired by this observation, we propose to use stereo disparity estimation techniques on the overlap region, which is between all adjacent camera pairs in the surround-view setting. Although this region is relatively small, it serves as a geometric link between two images. Consequently, we apply outputs from the traditional stereo disparity estimation model as weak depth supervision to improve the detection accuracy over the overlap region. We

empirically found that this supervision significantly improves the overall detection accuracy.

This is only part of a story. Unlike the human vision system that quickly identifies an object across overlapped cameras, we empirically observe that DNN has a strong inductive bias toward identifying objects individually in each single-view image. This often results in failing to utilize additional information of the same instance appearing in other different-view images. Moreover, we observe a domain shift effect between the overlap region (i.e., the region far from the camera center) and the non-overlap region (i.e., the region near the camera center) due to camera lens distortion. Thus, we further propose to train an adversarial overlap region discriminator, which minimizes the domain gap between objects in the non-overlap regions vs. overlap regions. We validate from experiments that such adversarial training makes the overlap region performance more robust.

We start from the state-of-the-art multi-view 3D object detection model called DETR3D [3] as shown in Figure 2. Built upon DETR3D, we introduce the following two main modules: (i) Stereo Disparity Estimation for Weak Supervision and (ii) Adversarial Overlap Region Discriminator. We evaluate the effectiveness of our proposed method using the nuScenes [1] dataset, which is a widely used large-scale multi-view 3D object detection benchmark. Through comprehensive experiments, we verify that our proposed model generally outperforms the state-of-the-art approaches in the camera-only 3D object detection task. Our main contributions are summarized as follows:

- We report that existing works often neglect properly dealing with objects in the overlap region, which limits fully using the geometric potentials of multi-view camera systems, causing performance degradation.
- We propose to use outputs from the traditional stereo disparity estimation model on the overlap region and apply them as weak supervision to improve the detection accuracy over the overlap region. We empirically find that this supervision significantly improves the overall detection accuracy.
- We introduce an overlap region discriminator that adversarially learns to minimize the covariate shift between objects from non-overlap regions vs. from overlap regions.

2 Related Work

3D Object Detection. As the study of inferring objects in 3D plane begins, Mono3D conducts 3D object detection task using multiple information such as RGB images, object instruction segmentation, context, and location prior information. The model proposed by Roddick, T. et al. [2] utilizes Bird-Eye-View(BEV) to compensate for performance degradation due to incorrect depth feature extracted from solely RGB images. MonoPSR presents a method to significantly reduce 3D search space using 2D object detection and exploit point clouds to recover local shape and scale. FCOS3D [23] and RTM3D [16] predicts 3D bounding box more accurately through features guided by geometric information. In addition, PGD [5] shows that geometric interactions between objects enhance the reliability of depth. Yet, the methods mentioned above perform detection tasks independently for each image. Thus, to expand for multiple cameras, each frame processing is performed before integrating the outputs in the post-processing stage. To address this problem, DETR3D introduces a 3D object detection method that simultaneously processes multi-view images. DETR3D predicts 3D bounding boxes via backward geometric projection and operates setto-set prediction without post-processing. Also, ImVoxelNet [22] implements a surrounding view 3D object detector similar to the LiDAR system by optimizing the chronic computation 4



Figure 2: An overview of our proposed architecture. Built upon DETR3D [53], our model takes multi-view camera inputs and outputs a set of 3D bounding boxes for objects in the scene. Our model consists of two main modules: (1) Stereo Matching Network for Weak Depth Supervision, where our depth estimation head is trained to predict a dense depth map of the overlap region. The ground-truth depth map is obtained by a traditional stereo disparity estimation algorithm. (2) Adversarial Overlap Region Discriminator, which minimizes the gap between non-overlap regions vs. overlap regions, improving the overall detection performance.

of multiple images. Although they use multi-view images, they do not take advantage of the features of multiple camera settings. Hence, we focus on overlaying clues between images, which are multi-view characteristics, to improve detection performance.

In addition to the approaches mentioned above, recently, another solution of transforming the image features into BEV representation and applying it to 3D object detection has been widely studied. Lift-Splat-Shoot (LSS) [22] introduces a view transform method that infers depth distribution and projects multiple image features into BEV representation. BEVDet [11] and BEVDet4D [11], which extend LSS, demonstrate that applying the BEV features to 3D object detection is practical. Following BEVDet, BEVDepth [12] constructs better BEV features with trustworthy depth prediction. BEVFormer [12] employs spatiotemporal cues by interacting with spatial and temporal space via predefined grid-shaped BEV queries. Note that our proposed method could potentially be applied to BEV-based approaches in an adhoc manner, as we focus on regularizing networks to deal with objects in the overlap region properly. However, we leave it as future work and will first focus on improving DETR3D-based approaches.

Stereo 3D Object Detection. The stereo 3D object detection task that utilizes binocular information is similar to the human system. Inspired by depth estimation models [2, 12, 13], DSGN [2] proposes an end-to-end model that simultaneously uses plane-sweep volume and 3D geometric volume to predict 3D bounding boxes. PLUMEnet [52] directly con-

structs a pseudo-LiDAR feature volume (PLUME) in 3D space. Reliable depth features extracted from the stereo view yield outstanding performance, whereas the 3D cost volume formed in the stereo network directs a lot of computation complexity. YOLOStereo3D [11], a single-stage 3D detection network, effectively deals with the trade-off between computational complexity and depth accuracy with light-weight cost volume. Additionally, MobileStereoNet [23] introduces a way to leverage MobileNets [2] to reduce the computation cost of deep networks without sacrificing accuracy. From these observations, our method utilizes the disparity information as 3D geometric cues in multi-view settings to enhance the accuracy of 3D object detection.

ORA3D 3

In the following sections, we present a novel multi-view 3D object detection model that leverages rich information from a camera-only multi-view vision system. Our model is built upon the state-of-the-art DETR3D [1] model, and we propose to use the following two main modules: Stereo Disparity Estimation for Weak Supervision (Section 3.1) and Adversarial Overlap Region Discriminator (Section 3.2).

Stereo Disparity Estimation for Weak Depth Supervision 3.1

Existing work suggests that 3D object detectors from stereo vision can take advantage of estimating accurate depth for objects from binocular images. Such depth information is helpful for camera-only object detectors, which often lack reliable depth information. Our surroundview camera setting differs from the conventional stereo vision task - only a tiny portion of the overlap region (i.e., less than 20%) is available. In this work, we advocate for leveraging such overlap regions to supervise networks learning depth cues, potentially providing better 3D detection performance.

Learning Depth Cue by Multi-view Stereo Matching. When considering the multi-view camera system, adjacent cameras have a strong association. We regard this association comes from overlap regions and can be extended to geometric guides. To interactively supervise the network, we train the Stereo Disparity Estimation head, which reconstructs a dense disparity map with overlap region pairs of neighboring cameras.

We follow the recent work by



Figure 3: Following Liu et al. [19], our Stereo Disparity Estimation head is co-trained to compute the disparity map from two overlapped images.

Liu et al. [1] to implement the disparity estimation head. As illustrated in Figure 3, our stereo network extracts features of an image pair with a standard visual encoder. Our disparity estimation head outputs a cost volume through multi-scale stereo matching. To obtain the target disparity map, we use the output from the conventional stereo matching algorithm $[\mathbf{N}]$, which performs pixel-wise mutual information-based matching.

Further, we use a binary mask to consider losses from overlap regions, blocking gradients from non-overlap regions for training our Stereo Disparity Estimation module. Given the known internal camera parameters and external parameters, we first map the point (x_s, y_s) in the source image coordinate to the point (X,Y,Z) in the world coordinate, i.e., $(x_s,y_s)^\top \mapsto (X,Y,Z)^\top$. Then, we project the point (X,Y,Z) back to the point (x_t,y_t) in the target (or neighboring) camera coordinate frame, i.e., $(X,Y,Z)^\top \mapsto (x_t,y_t)^\top$. We provide more details in the supplemental material.

Stereo Disparity Estimation Loss. We use the following stereo focal loss [**L**] \mathcal{L}_d to optimize our Stereo Disparity Estimation head:

$$\mathcal{L}_{d} = \frac{1}{|\mathcal{I}_{o}|} \sum_{x \in \mathcal{I}_{o}} \sum_{d=0}^{D-1} ((1 - P_{x}(d)))^{-\alpha} (-P_{x}(d) \log \hat{P}_{x}(d))$$
(1)

for pixels *x* in the overlap region \mathcal{I}_o . Note that $d \in \{0, 1, ..., D-1\}$ represents a discretized disparity, and $\alpha \in \mathcal{R}$ is the focus weight that is set to 1.0. $P_x(d)$ represents the target probabilistic distribution given a disparity *d* for a pixel *x* defined as follows: $P_x(d) = \text{Softmax}(-2|d - d^{gt}|)$ where d^{gt} is the target (or ground-truth) disparity map. Similarly, $\hat{P}_p(d)$ is the predicted probabilistic distribution of a given disparity *d*.

3.2 Adversarial Overlap Region Discriminator

In general, the human visual system can easily recognize objects at the edge of the image as well as in the center. Furthermore, it is not difficult to identify the same instances that appear simultaneously in different-view images. However, empirically, we discovered that DNN-based multi-view object detection models have a strong inductive bias toward identifying objects individually in each image. Note that objects in the overlap region are often occluded mainly due to limited Field Of View (FOV) of the camera or suffer from deformation due to camera lens distortion (e.g., pincushion distortion). Finally, we observe a domain shift effect between the overlap region (i.e., region far from camera center) vs. non-overlap region (i.e., region near camera center).

In addition to leveraging stereo disparity information, which gives an implicit bias to process the same object in different-view images together, we propose to use a regularizer to explicitly minimize the representational gap between non-overlap regions vs. overlap regions. Specifically, we constrain the object detection head from learning region-invariant information via an adversarial framework using Gradient Reversal Layer (GRL) [**G**].

Given a query feature \mathbf{q}_i for $i = \{1, 2, ..., |Q|\}$ in the transformer (i.e., object detection decoder), an overlap region discriminator f_d takes such query features. Formally, this discriminator needs to correctly predict its source region. Thus, this classifier f_d is trained using region classification loss \mathcal{L}_r as follows: $\mathcal{L}_r = -\mathbb{E}_{q,y_r \sim \mathbb{D}} [\sum_{r \in \mathcal{R}} y_r \log f_d(q)_r]$ where $\mathbb{E}_{q,y_r \sim \mathbb{D}}$ indicates an expectation over samples (q, y_r) , which are drawn from the (input) data distribution \mathbb{D} . Plus, f_d is trained to classify whether a query feature is from the overlap region $(y_r \text{ is set to } 1)$ or from the non-overlap region $(y_r \text{ is set to } 0)$, so the output dimension from this module is 2 (i.e. $y_r \in \{0, 1\}$). Notably, to reduce inductive bias, we design this ordinary discriminator as a special. We intentionally make a loss negative (i.e., $-\mathcal{L}_r$), preventing the discriminator from rightly distinguishing. Finally, our network is robust to all areas, minimizing bias that boosts performance degradation for overlapped regions.

Loss Function. Ultimately, our model is trained end-to-end by minimizing the following loss function \mathcal{L}_{total} :

$$\mathcal{L}_{\text{total}} = \lambda_{\text{cls}} \mathcal{L}_{\text{cls}} + \lambda_{\text{box}} \mathcal{L}_{\text{box}} + \lambda_d \mathcal{L}_d - \lambda_r \mathcal{L}_r$$
(2)

where λ_{cls} , λ_{box} , λ_d , and λ_r are hyperparameters that are drawn from a grid search to control the strength of λ_{cls} , \mathcal{L}_{box} , \mathcal{L}_d , and \mathcal{L}_r , respectively.

Table 1: Comparing our 3D object detector and the state-of-the-art on the nuScenes [I] dataset. All methods based on camera modality. † : trained with CBGS [I]. * : initialized from pre-trained model on extra data. ‡ : initialized from DD3D checkpoint.

Model	Mono Multi	Backbone	$NDS(\uparrow)$	mAP(↑)	Model	Mono Multi	Backbone	$NDS(\uparrow)$	mAP(†)
CenterNet []	Mono	DLA34	0.328	0.306	FCOS3D [[22]	Mono	ResNet101	0.428	0.358
FCOS3D []	Mono	ResNet101	0.415	0.343	PGD [[22]	Mono	ResNet101	0.448	0.386
PGD []	Mono	ResNet101	0.428	0.369	DD3D* [[22]]	Mono	V2-99	0.477	0.418
DETR3D [53]	Multi	ResNet101	0.425	0.346	DETR3D [‡] [☑]	Multi	V2-99	0.479	0.412
DETR3D [†] [53]	Multi	ResNet101	0.434	0.349	BEVDet [‡] [☑]	Multi	V2-99	0.482	0.422
Ours	Multi	ResNet101	0.445	0.367	Ours‡	Multi	V2-99	0.489	0.423

(a) Validation set

(b) Test set

Table 2: Average precision (AP) for each object class on the nuScenes [I] test set. Higher value is better. *Abbr.* C.V: construction vehicle, T.C: traffic cone.

Model	Car	Truck	Bus	Trailer	C.V	Ped.	Motor.	Bicycle	T.C	Barrier	mAP
DETR3D [53]	0.603	0.333	0.290	0.358	0.170	0.455	0.413	0.308	0.627	0.565	0.412
Ours	0.609	0.338	0.323	0.347	0.174	0.467	0.420	0.311	0.649	0.589	0.423

4 Experiments

Dataset. We use the nuScenes [I] dataset, a large-scale multi-view object detection benchmark. The nuScenes dataset provides a full 360-degree field of view captured by six different viewing cameras. This comprises 20-second-long 1,000 video sequences, which are fully annotated with 3D bounding boxes for 10 object classes. The dataset covers 28k annotated samples for training, and validation and test contain 6k scenes each.

Evaluation Metrics. We follow the official evaluation protocol of nuScenes [I]. We use a set of True Positive metrics (TP metrics) for each prediction that was matched with a ground-truth box. We employ the following 5 TP metrics: Average Translation Error(ATE), Average Scale Error(ASE), Average Orientation Error(AOE), Average Velocity Error(AVE), and Average Attribute Error(AAE). All TP metrics are also calculated using a 2m center distance threshold during matching, and they are all designed to be positive scalars. We also measure mean average precision (mAP). Lastly, we use the nuScenes Detection Score (NDS) to measures a consolidated scalar metric defined as follows: NDS = $\frac{1}{10}$ [5 mAP + $\sum_{mTP \in TP} (1 - min(1, mTP))$] where TP is five TP metrics. Note that implementation and training details are provided in the supplemental material.

Performance Comparison with SOTA. We compare our proposed ORA3D with existing state-of-the-art methods including CenterNet [1], FCOS3D [2], PGD [3], DD3D [2], DETR3D [3], and BEVDet [1]. The first four approaches take multiple single-view images independently and combine detection outputs for the final output using non-maximum suppression (NMS). The last two approaches (DETR3D and BEVDet) are state-of-the-art multi-view 3D object detection models. As shown in Table 1 (a) and (b), ORA3D generally outperforms other methods (compare the last row vs. others) in terms of (a consolidated metric) NDS and mAP both in validation and test data. Note that FCOS3D uses test-time augmentation and a customized data augmentation strategy, thus requiring more epochs and model ensembles. Though we do not use such an augmentation strategy, ours performs better than these methods. Note that we initialize with DD3D [2] pre-trained model and use the same backbone (V2-99 [1]) for a fair comparison with DETR3D and BEVDet.

Further, we observe in Table 2 that ours show higher mAP scores in all objects except

Table 3: Detection performance comparison with the state-of-the-art approaches for **objects** in the overlap region. nuScenes [II] validation set are used and all use the same backbone.



Figure 4: The computing budget and performance of different paradigms on the nuScenes [II] validation set. Note that DETR3D+ indicates a DETR3D [II] model with CBGS [II].

Trailer. Note that ours is initialized from the same backbone of DETR3D [53]. This confirms that our proposed regularization terms clearly improve the overall detection performance.

Performance Comparison in Overlap Region. We compare our proposed ORA3D again with existing state-of-the-art methods (FCOS3D [23], and DETR3D [53]), but we now focus on overlap regions. As shown in Table 3, our method outperforms FCOS3D and DETR3D in terms of all metrics except mASE and mAAE. This confirms that our overlap region aware approach effectively deals with objects in the overlap region, resulting in a large performance gain. For a fair comparison, all models use a ResNet101-based backbone, and only FCOS3D uses an augmentation strategy.

Qualitative Analysis. Existing CNN-based methods benefit from inductive bias even with relatively small datasets. However, large-scale datasets have become more common, and in these environments, inductive bias easily overfits specific datasets and causes various domain shift issues. In this section, we demonstrate that our proposed methods address these concerns while maximizing the capacity of multi-view camera systems.

Fig. 5 shows the visualized results of 3D bounding boxes predicted by DETR3D (see green boxes) and our proposed method (see pink boxes). The ground-truth bounding boxes (see blue boxes) are overlaid. We project the predicted and ground-truth bounding boxes in the BEV perspective. In general, DETR3D and our proposed method generate reasonable results. However, the difference in the overlap region is apparent between the two methods, where DETR3D exhibits a relatively large number of false positive detections. Especially, our model is more robust for truncated or distorted objects across the entire region. Overall, our simple and effective methods amplify accuracy both qualitatively and quantitatively.

Analysis of Computations. We compare our proposed method with existing 3D object detection methods, including FCOS3D [29], PGD [51], BEVDet [51] and DETR3D [53]. As illustrated in Fig. 4, our method outperforms other methods while requiring only a computational budget comparable to DETR3D. Furthermore, DETR3D uses CBGS [52], a helpful strategy for more balanced data distribution, to achieve higher performance. This strategy,



Figure 5: Qualitative comparison between our method (purple) and DETR3D [53] (green). Detected 3D bounding boxes for objects in the scene are projected into a bird's eye view perspective (left) and the image plane (right). See DETR3D produces more false positives in the overlapped regions than ours.

however, alleviates the problem of data imbalance but requires more than four times the training time. Our method without CBGS demands less training time, but has the most impressive performance in Fig 4. Additionally, although another great work, BEVDet, uses fewer FLOPs by using smaller-sized images, ours shows higher scores (NDS: 0.445, mAP: 0.364). It would be worth exploring as a future work applying our approach to BEV-based models for potential performance improvement.

Ablation Study. We evaluate the variants of our method with DETR3D with and without the following two main loss terms: (i) \mathcal{L}_d : Stereo Disparity Estimation loss, (ii) \mathcal{L}_r : Adversarial Overlap Region Discrimination loss. In Table 4, we quantitatively analyze the importance of our proposed methods for the whole region and Table 4: Ablative analysis of our methods and DETR3D [3] on the nuScenes [3] *mini* validation set for the whole and overlap regions.

Model	Whole	Region	Overlap Region			
	NDS (†)	$\mathrm{mAP}\left(\uparrow\right)$	NDS (\uparrow)	$mAP\left(\uparrow\right)$		
A. DETR3D [0.338	0.247	0.266	0.151		
B. A + \mathcal{L}_d	0.348	0.273	0.276	0.195		
C. A + \mathcal{L}_r	0.339	0.286	0.274	0.215		
D. A + \mathcal{L}_d + \mathcal{L}_r (Ours)	0.353	0.288	0.280	0.207		

the overlap region. Note that we use the nuScenes mini validation set. In Table 4 Whole Region, we observe that adding either \mathcal{L}_d improves most of the metrics, while a similar trend is observed by adding \mathcal{L}_r as well. Indeed, the stereo disparity loss is reflected during training to induce semantic geometric potentials evenly in all areas without distinguishing between overlap and non-overlap regions. In Table 4 Overlap Region, the effect of each component on the overlap region is further confirmed. Most importantly, we discover the remarkable efficiency of \mathcal{L}_r in the target domain. The adversarial loss completely addresses the domain shift effect by deceiving the model into not correctly discriminating which region objects belong to. As a result, we confirm that our components improve overall performance.

Analysis of Domain Shift Effect. We suppose that the multi-view 3D object detection method, such as DETR3D [53], without consideration of lens distortion and limited FOV, is likely to have an inductive bias between central and outer regions, causing a domain shift effect. To find the degree of inductive bias, we visualize the t-SNE [23] results of central (non-overlapped, blue) and edge (overlapped, red) regions' features in Figure 6: t-SNE visualization [23] of non-overlapped (blue) and overlapped (red) regions' features. We discover that our Overlap Region Discriminator suitably overcomes the domain shift effect between two regions. Best viewed in color.



Fig 6. As we clearly see in the first row (DETR3D), the features of each two groups form distinguishable clusters. Our method starts with this observation. The overlap regions in multiview images are clues to improving the overall performance. Ultimately, the features of our method's two groups (second row) are distributed harmoniously, minimizing bias. This distribution indicates that our proposed Overlap Region Discriminator successfully minimizes the representational gap.

5 Conclusion

In this paper, we present a novel pipeline to enhance the performance of 3D object detection using the 3D geometric cues in multi-view settings. We recognize that the overlaying region between images has become a weakness due to issues such as disconnected information and overlapping objects. However, the overlap region is also a novel material with the potential to improve overall performance. We develop approaches to reasonably use this small but highly informative area. The first is a light and robust stereo disparity estimation network for small overlap regions. This module allows the network to fully utilize the geometric potential of binocular images. Next is the adversarial overlap region discriminator, which is adversarially trained to minimize the gap between non-overlap regions and overlap regions. Overall, the experiments demonstrate that our two methods successfully work on detection accuracy.

Acknowledgement

This work is supported by Autonomous Driving Center, Hyundai Motor Company R&D Division. S. Kim, W. Roh, G. Chang, and C. Kim are partially supported by the National Research Foundation of Korea grant (NRF-2022R1F1A1074334) and Institute of Information & communications Technology Planning & Evaluation (IITP) grant funded by the Korea government(MSIT) (No. 2019-0-00079, Artificial Intelligence Graduate School Program(Korea University). S. Moon and J. Kim are partially supported by the ICT Creative Consilience program (IITP-2022-2022-0-01819) and the ITRC(Information Technology Research Center) support program (IITP-2022-RS-2022-00156295).

References

- [1] Holger Caesar, Varun Bankiti, Alex H Lang, Sourabh Vora, Venice Erin Liong, Qiang Xu, Anush Krishnan, Yu Pan, Giancarlo Baldan, and Oscar Beijbom. nuscenes: A multimodal dataset for autonomous driving. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 11621–11631, 2020.
- [2] Jia-Ren Chang and Yong-Sheng Chen. Pyramid stereo matching network. In Proceedings of the IEEE conference on computer vision and pattern recognition, pages 5410–5418, 2018.
- [3] Xiaozhi Chen, Kaustav Kundu, Ziyu Zhang, Huimin Ma, Sanja Fidler, and Raquel Urtasun. Monocular 3d object detection for autonomous driving. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 2147–2156, 2016.
- [4] Yilun Chen, Shu Liu, Xiaoyong Shen, and Jiaya Jia. Dsgn: Deep stereo geometry network for 3d object detection. In *Proceedings of the IEEE/CVF conference on computer* vision and pattern recognition, pages 12536–12545, 2020.
- [5] Jiajun Deng, Shaoshuai Shi, Peiwei Li, Wengang Zhou, Yanyong Zhang, and Houqiang Li. Voxel r-cnn: Towards high performance voxel-based 3d object detection. arXiv preprint arXiv:2012.15712, 1(2):4, 2020.
- [6] Yaroslav Ganin, Evgeniya Ustinova, Hana Ajakan, Pascal Germain, Hugo Larochelle, François Laviolette, Mario Marchand, and Victor Lempitsky. Domain-adversarial training of neural networks. *The journal of machine learning research*, 17(1):2096–2030, 2016.
- [7] Qingdong He, Zhengning Wang, Hao Zeng, Yi Zeng, Shuaicheng Liu, and Bing Zeng. Svga-net: Sparse voxel-graph attention network for 3d object detection from point clouds. arXiv preprint arXiv:2006.04043, 2020.
- [8] Heiko Hirschmuller. Stereo processing by semiglobal matching and mutual information. *IEEE Transactions on pattern analysis and machine intelligence*, 30(2):328–341, 2007.
- [9] Andrew G Howard, Menglong Zhu, Bo Chen, Dmitry Kalenichenko, Weijun Wang, Tobias Weyand, Marco Andreetto, and Hartwig Adam. Mobilenets: Efficient convolutional neural networks for mobile vision applications. arXiv preprint arXiv:1704.04861, 2017.
- [10] Junjie Huang and Guan Huang. Bevdet4d: Exploit temporal cues in multi-camera 3d object detection. *arXiv preprint arXiv:2203.17054*, 2022.
- [11] Junjie Huang, Guan Huang, Zheng Zhu, Ye Yun, and Dalong Du. Bevdet: Highperformance multi-camera 3d object detection in bird-eye-view. *arXiv preprint arXiv:2112.11790*, 2021.
- [12] Alex Kendall, Hayk Martirosyan, Saumitro Dasgupta, Peter Henry, Ryan Kennedy, Abraham Bachrach, and Adam Bry. End-to-end learning of geometry and context for deep stereo regression. In *Proceedings of the IEEE international conference on computer vision*, pages 66–75, 2017.

12 W.ROH ET AL.: OVERLAP REGION AWARE MULTI-VIEW 3D OBJECT DETECTION

- [13] Jason Ku, Alex D Pon, and Steven L Waslander. Monocular 3d object detection leveraging accurate proposals and shape reconstruction. In *Proceedings of the IEEE/CVF* conference on computer vision and pattern recognition, pages 11867–11876, 2019.
- [14] Youngwan Lee and Jongyoul Park. Centermask: Real-time anchor-free instance segmentation. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 13906–13915, 2020.
- [15] Peiliang Li, Xiaozhi Chen, and Shaojie Shen. Stereo r-cnn based 3d object detection for autonomous driving. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 7644–7652, 2019.
- [16] Peixuan Li, Huaici Zhao, Pengfei Liu, and Feidao Cao. Rtm3d: Real-time monocular 3d detection from object keypoints for autonomous driving. In *European Conference* on Computer Vision, pages 644–660. Springer, 2020.
- [17] Yinhao Li, Zheng Ge, Guanyi Yu, Jinrong Yang, Zengran Wang, Yukang Shi, Jianjian Sun, and Zeming Li. Bevdepth: Acquisition of reliable depth for multi-view 3d object detection. arXiv preprint arXiv:2206.10092, 2022.
- [18] Zhiqi Li, Wenhai Wang, Hongyang Li, Enze Xie, Chonghao Sima, Tong Lu, Qiao Yu, and Jifeng Dai. Bevformer: Learning bird's-eye-view representation from multi-camera images via spatiotemporal transformers. *arXiv preprint arXiv:2203.17270*, 2022.
- [19] Yuxuan Liu, Lujia Wang, and Ming Liu. Yolostereo3d: A step back to 2d for efficient stereo 3d detection. In 2021 IEEE International Conference on Robotics and Automation (ICRA), pages 13018–13024. IEEE, 2021.
- [20] Jiageng Mao, Minzhe Niu, Haoyue Bai, Xiaodan Liang, Hang Xu, and Chunjing Xu. Pyramid r-cnn: Towards better performance and adaptability for 3d object detection. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 2723–2732, 2021.
- [21] Dennis Park, Rares Ambrus, Vitor Guizilini, Jie Li, and Adrien Gaidon. Is pseudo-lidar needed for monocular 3d object detection? In *IEEE/CVF International Conference on Computer Vision (ICCV)*, 2021.
- [22] 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*, pages 194–210. Springer, 2020.
- [23] Thomas Roddick, Alex Kendall, and Roberto Cipolla. Orthographic feature transform for monocular 3d object detection. *arXiv preprint arXiv:1811.08188*, 2018.
- [24] Danila Rukhovich, Anna Vorontsova, and Anton Konushin. Imvoxelnet: Image to voxels projection for monocular and multi-view general-purpose 3d object detection. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision*, pages 2397–2406, 2022.
- [25] Faranak Shamsafar, Samuel Woerz, Rafia Rahim, and Andreas Zell. Mobilestereonet: Towards lightweight deep networks for stereo matching. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision*, pages 2417–2426, 2022.

- [26] 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, pages 10529–10538, 2020.
- [27] Shaoshuai Shi, Li Jiang, Jiajun Deng, Zhe Wang, Chaoxu Guo, Jianping Shi, Xiaogang Wang, and Hongsheng Li. Pv-rcnn++: Point-voxel feature set abstraction with local vector representation for 3d object detection. arXiv preprint arXiv:2102.00463, 2021.
- [28] L van der Maaten and G Hinton. Visualizing high-dimensional data using t-629 sne. *Journal of Machine Learning Research*, 9(2579-2605):630, 2008.
- [29] Tai Wang, Xinge Zhu, Jiangmiao Pang, and Dahua Lin. Fcos3d: Fully convolutional one-stage monocular 3d object detection. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 913–922, 2021.
- [30] Tai Wang, ZHU Xinge, Jiangmiao Pang, and Dahua Lin. Probabilistic and geometric depth: Detecting objects in perspective. In *Conference on Robot Learning*, pages 1475– 1485. PMLR, 2022.
- [31] 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, pages 8445–8453, 2019.
- [32] Yan Wang, Bin Yang, Rui Hu, Ming Liang, and Raquel Urtasun. Plumenet: Efficient 3d object detection from stereo images. In 2021 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS), pages 3383–3390. IEEE, 2021.
- [33] Yue Wang, Vitor Campagnolo Guizilini, Tianyuan Zhang, Yilun Wang, Hang Zhao, and Justin Solomon. Detr3d: 3d object detection from multi-view images via 3d-to-2d queries. In *Conference on Robot Learning*, pages 180–191. PMLR, 2022.
- [34] 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 *ICLR*, 2020.
- [35] Jure Zbontar, Yann LeCun, et al. Stereo matching by training a convolutional neural network to compare image patches. *J. Mach. Learn. Res.*, 17(1):2287–2318, 2016.
- [36] Youmin Zhang, Yimin Chen, Xiao Bai, Suihanjin Yu, Kun Yu, Zhiwei Li, and Kuiyuan Yang. Adaptive unimodal cost volume filtering for deep stereo matching. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 34, pages 12926–12934, 2020.
- [37] Xingyi Zhou, Dequan Wang, and Philipp Krähenbühl. Objects as points. *arXiv preprint arXiv:1904.07850*, 2019.
- [38] Benjin Zhu, Zhengkai Jiang, Xiangxin Zhou, Zeming Li, and Gang Yu. Classbalanced grouping and sampling for point cloud 3d object detection. *arXiv preprint arXiv:1908.09492*, 2019.