

CRAFT: Camera-Radar 3D Object Detection with Spatio-Contextual Fusion Transformer

Youngseok Kim¹, Sanmin Kim¹, Jun Won Choi², Dongsuk Kum¹

¹Korea Advanced Institute of Science and Technology

²Hanyang University

{youngseok.kim, sanmin.kim, dskum}@kaist.ac.kr, junwchoi@hanyang.ac.kr

Abstract

Camera and radar sensors have significant advantages in cost, reliability, and maintenance compared to LiDAR. Existing fusion methods often fuse the outputs of single modalities at the result-level, called the late fusion strategy. This can benefit from using off-the-shelf single sensor detection algorithms, but late fusion cannot fully exploit the complementary properties of sensors, thus having limited performance despite the huge potential of camera-radar fusion. Here we propose a novel proposal-level early fusion approach that effectively exploits both spatial and contextual properties of camera and radar for 3D object detection. Our fusion framework first associates image proposal with radar points in the polar coordinate system to efficiently handle the discrepancy between the coordinate system and spatial properties. Using this as a first stage, following consecutive cross-attention based feature fusion layers adaptively exchange spatio-contextual information between camera and radar, leading to a robust and attentive fusion. Our camera-radar fusion approach achieves the state-of-the-art 41.1% mAP and 52.3% NDS on the nuScenes test set, which is 8.7 and 10.8 points higher than the camera-only baseline, as well as yielding competitive performance on the LiDAR method.

1 Introduction

3D object detection is an essential task for autonomous driving as well as mobile robots. Thanks to the low-cost, high-reliability, and low-maintenance of camera and radar sensors, they are already deployed to a large number of mass production vehicles for active safety systems. Moreover, the rich semantic information of the camera and long-range detection with weather condition robustness of radar are essential attributes that LiDAR cannot provide. Nevertheless, learning-based camera-radar 3D object detection (Lim et al. 2019; Kim, Choi, and Kum 2020; Nabati and Qi 2021) for autonomous driving has not been thoroughly explored.

Existing camera-radar fusion methods often use a late fusion strategy that fuses detection-level outputs of separated camera and radar detection algorithms with heuristic logic. Such methods can have the benefit of exploiting the off-the-shelf detection algorithms that are independently developed by automotive suppliers as modular components. How-

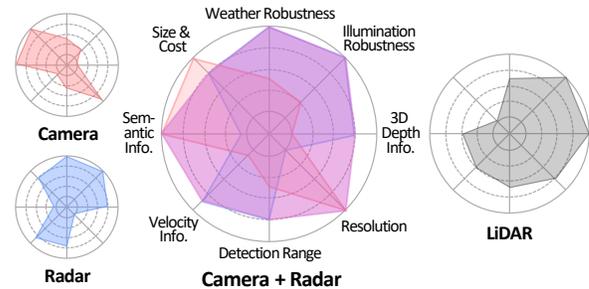


Figure 1: Sensor characteristics of camera, radar, and LiDAR. Camera-radar fusion has high potential considering spatially and contextually complementary properties.

ever, late fusion strategies that rely on heuristics and post-processing techniques suffer from performance-reliability trade-offs, especially when two sensor predictions disagree; thus, late fusion cannot exploit the full potential of each sensor. In contrast, a learning-based early fusion strategy that fuses the intermediate information of each sensor at an early stage has much higher potential, but it requires an in-depth understanding of each sensor’s characteristics to find the “optimal” way to fuse.

However, it is non-trivial to develop early fusion for camera and radar due to the unique characteristics of each sensor, as illustrated in Fig. 1. Camera-LiDAR early fusion methods, which are relatively well-studied, fuse point-wise features (Vora et al. 2020; Huang et al. 2020) by associating image pixel with projected LiDAR point or fuse image and LiDAR proposal features (Ku et al. 2018) by extracting a region of interests (RoIs) with RoIPool (Ren et al. 2015). However, adapting camera-LiDAR fusion methods for camera-radar is unsuitable since those strategies are established by accurate LiDAR measurement ($\pm 2cm$), but radar has low accuracy and measurement ambiguities. Due to the nature of the radar mechanism, the radar has high resolution and accuracy in the radial direction ($0.4m$ and $\pm 0.1m$), which is measured by Fast Fourier Transformation FFT. Meanwhile, the azimuthal measurement obtained by digital beamforming using multiple receive antennas (Johnson and Dudgeon 1992) is inaccurate (4.5° and $\pm 1^\circ$, approximately $4m$ and $\pm 1m$ at $50m$ distance). Measurement

ambiguities are the cases when radar points are occasionally missed on objects (false negatives due to the low radar cross-section) or shown in the background (false positives due to radar multi-path problem), thus these have to be considered. The camera, meanwhile, has very complementary spatial characteristics to radar (Ma et al. 2021; Hung et al. 2022). Dense camera pixels provide accurate azimuthal resolution and accuracy, but the depth information is not provided due to the perspective projection.

Motivated by the aforementioned challenges of camera and radar, camera-radar fusion has to be able to 1) robustly operate even radar points are not reflected from the object and 2) effectively complement spatial (range/azimuth) and contextual (semantic/Doppler) information. To this end, we present Camera RADar Fusion Transformer (CRAFT), a robust and effective fusion framework for 3D object detection. We propose the soft association strategy between image proposal and radar points and cross-attention based feature fusion method, which can effectively exploit spatial characteristics and robustly handle radar measurement ambiguity. Specifically, we first extract camera and radar features by modality-specific feature extractors (*i.e.*, DLA-34 (Yu et al. 2018), PointNet++ (Qi et al. 2017a)) and predict 3D objects with a lightweight camera 3D object detector. Given image proposals and radar points, we associate image proposal with radar points by Soft Polar Association, which queries points within the uncertainty-aware adaptive ellipsoidal threshold in the polar coordinate. To deal with the wrongly associated radar points (background clutters), we then attentively fuse the image proposal feature and radar point features by consecutive cross-attention based encoder layers. Our Spatio-Contextual Fusion Transformer can effectively exchange spatial and contextual information and adaptively determine where and what information to use for fusion. Finally, the class-specific decoder head predicts fusion score and offsets in the polar coordinate to refine the image proposal. Since some objects do not have valid radar points, fused output with a low fusion score is discarded, and the image proposal is used as the final prediction.

To summarize, our contributions are as follows:

- We investigate the characteristics of camera-radar and propose a proposal-level early fusion framework that can mitigate the coordinate system discrepancy and measurement ambiguities.
- We propose Soft Polar Association and Spatio-Contextual Fusion Transformer that effectively exchanges complementary information between camera and radar features.
- We achieve the state-of-the-art 41.1% mAP and 52.3% NDS on nuScenes test set, which significantly boosts the camera-only baseline with a marginal additional computation cost.

2 Related Work

2.1 Camera and Radar for 3D Object Detection

With an advance in 2D object detection (Ren et al. 2015; Tian et al. 2019), image view approaches extend

a 2D detector with additional 3D regression branches. CenterNet (Zhou, Wang, and Krähenbühl 2019) and FCOS3D (Wang et al. 2021b) directly regress depth to the object and 3D size from its center without using an anchor, which is later improved in PGD (Wang et al. 2021a) by utilizing geometric prior. Image view approaches can well exploit GPU-friendly operations and therefore runs fast but suffer from inaccurate depth estimation, which is the naturally ill-posed problem.

Another approaches are to transform image features in the perspective view into a top-down view and predict 3D bounding boxes from the BEV features using LiDAR detection heads (Lang et al. 2019; Yin, Zhou, and Krahenbühl 2021). CaDDN (Reading et al. 2021) and LSS (Phillion and Fidler 2020) generate the BEV features using the depth distribution, and PETR (Liu et al. 2022) and BEVFormer (Li et al. 2022) generate the BEV features by using pre-defined grid-shaped BEV queries to image features. Such BEV methods can utilize more depth information during training, thus showing better localization performance, but operates slow.

Radar data can have various representations such as 2-D FFT (Lin et al. 2018), Range-Azimuth-Doppler Tensor (Major et al. 2019; Kim, Kim, and Kum 2020); and a radar point cloud (Caesar et al. 2020; Meyer and Kuschik 2019a) is an affordable representation for autonomous driving applications. Although the radar point cloud is similar to LiDAR, naively adopting LiDAR methods to radar is inappropriate. (Ulrich et al. 2022) adapt PointPillars with KP-Conv (Thomas et al. 2019) and Radar-PointGNN (Svenningsson, Fioranelli, and Yarovoy 2021) adapt GNN to radar, but their performances are much inferior to LiDAR methods. Although the high potential of radar, radars have not yet been thoroughly investigated in autonomous driving.

2.2 Camera-LiDAR for 3D Object Detection

Camera-LiDAR fusion has gained significant interest for 3D detection, and existing approaches can be roughly classified into result-, proposal- and point-level fusion. Thanks to the rich intermediate feature representation of the learning-based method, a line of work (Ku et al. 2018; Yoo et al. 2020) fuses RoI proposal features to complement two modalities' information. Recent works (Vora et al. 2020; Bai et al. 2022) further propose point-level fusion that can fuse more fine-grained features at an earlier stage. However, methods using LiDAR are established on the strong assumption that LiDAR measurement and camera-LiDAR calibration is accurate, which is not valid under a radar setting.

2.3 Camera-Radar Fusion

Only a few methods attempt to detect 3D objects using camera radar fusion in autonomous driving. (Lim et al. 2019) transforms the camera image into BEV using inverse perspective mapping, assuming a planar road scene such as a highway to mitigate the coordinate discrepancy between camera and radar. (Meyer and Kuschik 2019b; Kim, Kim, and Kum 2020) adapt AVOD (Ku et al. 2018) to fuse front view camera image with point- and tensor-like radar data,

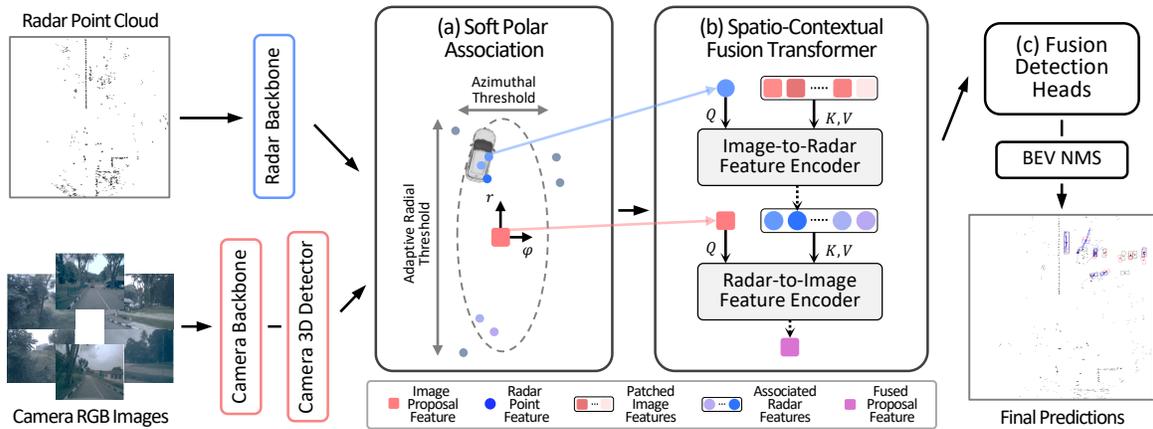


Figure 2: Overall architecture of CRAFT. Modality-specific backbones extract features, and camera 3D detector predicts 3D image proposals. (a) Given the image proposals, we associate radar points around the image proposals using adaptive thresholds in polar coordinates. (b) Then, consecutive cross-attention layers adaptively complement spatial and contextual information of camera and radar features. (c) Finally, fusion detection heads predict fusion score and offset to refine the image proposal and to output final predictions.

respectively. GRIF Net (Kim, Choi, and Kum 2020) further proposes a gating mechanism to adaptively fuse camera and radar RoI features handling the radar sparsity. Center-Fusion (Nabati and Qi 2021) lifts image proposals into a 3D frustum and associates a single closest radar point inside RoI to fuse. Thus, existing camera-radar 3D detectors have not thoroughly considered the spatial properties of camera and radar.

Some literature exploits camera and radar for depth completion tasks considering radar sparsity, inaccurate measurement, and ambiguity. (Lin, Dai, and Van Gool 2020) proposes a radar noise filtering module using a space-increasing discretization (SID) threshold to filter outlier noises, then uses filtered radar points to improve depth prediction. RC-PDA (Long et al. 2021b) further proposes pixel depth association to find the one-to-many mapping between radar point and image pixels that filter and densify radar depth map. In our work, we explore a robust fusion method using an attention mechanism to handle these limitations of radar effectively.

3 Method

In the camera-radar fusion-based 3D detection task, we take surrounding images and radar point clouds with corresponding intrinsics and extrinsics as input. Given camera feature maps from the backbone, camera detection heads first generate object proposals with a 3D bounding box. A *Soft Polar Association (SPA)* module then associates radar points with object proposals using adaptive thresholds in the polar coordinate to handle the spatial characteristics of sensors effectively. Further, a *Spatio-Contextual Fusion Transformer (SCFT)* fuses camera and radar features to complement spatial and contextual information of a single modality to another. Finally, *Fusion Detection Heads* decode the fused object features to refine initial object proposals. An overall framework is provided in Fig. 2.

3.1 Backbones and Camera 3D Object Detector

Given N camera images $I \in \mathbb{R}^{N \times H \times W \times 3}$, the 3D detector aims to classify and localize objects with a set of 3D bounding boxes. Specifically, we feed multi-camera images to the backbone network (e.g., DLA-34 (Yu et al. 2018)) and obtain downsampled image features $F \in \mathbb{R}^{N \times h \times w \times C}$ of multiple camera views. Taking camera view features F as input, the convolutional detection heads predict the projected 3D center (u, v) (called keypoint), depth d , and other attributes in the image plane. We additionally predict the variance of depth σ following (Kendall and Gal 2017) to represent the uncertainty of depth regression. Keypoints are transformed into the 3D camera coordinate system with camera intrinsics (x', y', f_u, f_v) as follows:

$$z_c = d, x_c = \frac{(u - x') \times z}{f_u}, y_c = \frac{(v - y') \times z}{f_v}, \quad (1)$$

then 3D bounding boxes in each camera coordinate system are transformed into the vehicle coordinate system. The bounding box $b = (x, y, z, \sigma, w, l, h, \theta, v)$ consists of the 3D center location (x, y, z) with depth variance σ , 3D dimension (w, l, h) , yaw orientation θ , and velocity v . In the keypoint detection setting (Zhou, Wang, and Krähenbühl 2019), the feature at the keypoint $f_{n,u,v} \in \mathbb{R}^C$ implicitly represents the object. Thus, we denote a set of M image proposals $\{i_m\}_{m=1}^M$, where $i_m = [b_m, f_m^i]$ indicates the bounding box and its features. Note that proposing a camera 3D detector is out of the scope of this paper, and we expect that other research progress (Wang et al. 2021b; Park et al. 2021) could further improve our results.

For a set of K radar points $\{p_k = [v_k, f_k]\}_{k=1}^K$ with 3D location $v_i \in \mathbb{R}^3$ and properties $f_i \in \mathbb{R}^F$ (e.g., radar cross-section (RCS), ego-motion compensated radial velocity), we extract the high-dimensional features $f_k^r \in \mathbb{R}^C$ using a point feature extractor (e.g., PointNet++ (Qi et al. 2017b)). As the

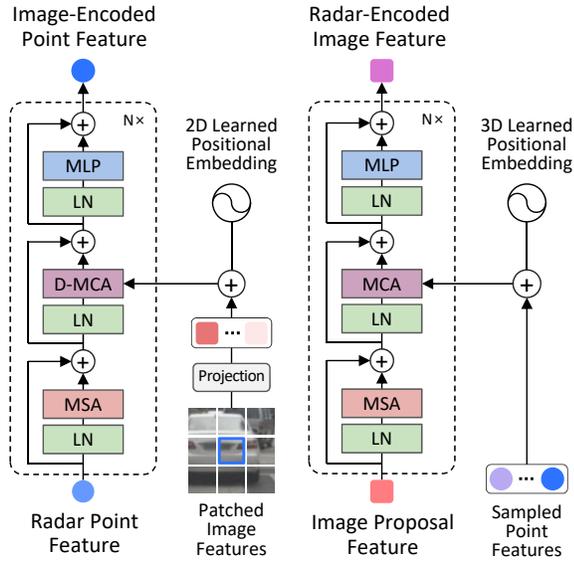


Figure 3: Architectures of the Image-to-Radar (left) and Radar-to-Image (right) Feature Encoder. Radar and image features exchange spatial and contextual information by consecutive feature encoders. We use RGB image for patched image features only for visualization.

number of radar points K is considerably small, we discard subsampling strategies such as farthest point sampling (FPS) (Shi, Wang, and Li 2019; Yang et al. 2020). The detailed architectures for the camera 3D detector and radar backbone are provided in Appendix A.

3.2 Soft Polar Association

Specifically, we calculate eight corners from the image proposal bounding box b and transform corners and radar points in Cartesian coordinate to polar coordinate $O^{(P)} = \{(r, \phi, z)_{m,j}\}_{m,j=1}^{M,8}, v^{(P)} = \{(r, \phi, z)_k\}_{k=1}^K$. Inspired by Ball Query (Qi et al. 2017b), we find a subset of radar points $\{p^m = [v_i, f_i]\}_{i=1}^{K'}$ for each object proposal within a certain distance to the image proposal but with adaptive thresholds.

$$\phi_{m,l} < \phi_i < \phi_{m,r} \quad (2)$$

$$r_{m,f} - (\gamma + \sigma \frac{r_c}{\delta}) < r_i < r_{m,b} + (\gamma + \sigma \frac{r_c}{\delta}) \quad (3)$$

Eq. 2 is the azimuthal threshold for radar points in p^m where $\phi_{m,l}$ and $\phi_{m,r}$ denote the angle of the most left and right corner points in polar coordinate. The radial threshold is represented in Eq. 3, where $r_{m,f}, r_{m,b}$ refer to radial distances of most front and back corner points of the image proposal, γ is a minimum range, δ is a hyper-parameter to modulate the radius size, and σ is the depth variance of the image proposal. $r_c = \frac{r_{m,f} + r_{m,b}}{2}$ denotes the radial distance of the center point of the image proposal. In this way, the proposed association using an adaptive threshold can maximize the chance of using informative foreground radar points while excluding most clutter points in the azimuthal direction.

3.3 Spatio-Contextual Fusion Transformer

In our camera-radar fusion framework, the primary role of radar is to refine the location of the image proposal if there are radar returns from the object. The Spatio-Contextual Fusion Transformer (SCFT) aims to exchange the spatial and contextual information between camera and radar, but determining which image pixels corresponding to the radar point is a difficult problem. Thus we design fusion modules with cross-attention layers to make the fusion network learn where and what information should be taken from image and radar.

Image-to-Radar Feature Encoder For spatio-contextual fusion, Image-to-Radar Feature Encoder (I2R FE) first provides the semantic information from image to radar. Each radar point is projected to the image plane, and then a patch is defined around the projected location. Instead of setting a fixed size patch, we design the size of the patch to be a function of radar distance so that the radar points can attend to a wider region when the object is closer, which takes more pixels in the image. The adaptive size is similar to the inverse function of space-increasing discretization (SID) (Fu et al. 2018):

$$\tau(d) = \lfloor W \exp(-d/\beta + \alpha) \rfloor, \quad (4)$$

where we heuristically set $W = 3.5$, $\alpha = 2$, and $\beta = 55$. Then, $\tau \times \tau$ size patched image feature map is resized to the fixed size $w \times w$ using bilinear interpolation.

Inspired by Deformable DETR (Zhu et al. 2021), we adopt a deformable multi-head cross-attention (D-MCA) module to attend to a small set of key sampling points (patched image features) around a reference point (projected radar pixel). Given a radar query feature $f_k^r \in \mathbb{R}^C$, I2R FE adaptively extracts features from the patched image features $f_k^i \in \mathbb{R}^{w \times w \times C}$ even radar point is not accurately projected into the image. However, the attention network desires proper supervision to learn which image feature is informative for the fusion effectively. Accordingly, we design an auxiliary task for each image-encoded radar feature to predict the probability that the radar point is inside the 3D bounding box.

Radar-to-Image Feature Encoder Following Radar-to-Image Feature Encoder (R2I FE) provides the spatial information of radar points to image proposal. Unlike I2R FE operates in a 2D camera coordinate, R2I FE takes inputs in a polar coordinate since a regression target of 3D object detection is to predict the location in 3D space. Following Transformer methods designed for 3D point cloud (Pan et al. 2021; Misra, Girdhar, and Joulin 2021), we use a cross-attention (Vaswani et al. 2017) with 'pre-norm' techniques (Xiong et al. 2020). Specifically, R2I FE takes the image proposal feature query $f_m^i \in \mathbb{R}^C$ and a set of radar point features $f_k^r \in \mathbb{R}^{K' \times C}$ as input, then produces a radar-encoded image proposal feature that is later used to refine the proposal. We additionally add a batch of zeros to the key-value sequences (`add_zero_attn` in PyTorch `nn.MultiHeadAttention`) so that attention can be assigned to it when none of the associated radar points are reflected from the object.

Method	Input	Backbone	NDS \uparrow	mAP \uparrow	mATE \downarrow	mASE \downarrow	mAOE \downarrow	mAVE \downarrow	mAAE \downarrow	FPS
PointPillars (2019)	L	-	45.3	30.5	0.517	0.290	0.500	0.316	0.368	61
CenterPoint (2021)	L	-	67.3	60.3	0.262	0.239	0.361	0.288	0.136	30
Radar-PointGNN (2021)	R	-	3.4	0.5	1.024	0.859	0.897	1.020	0.931	-
KPConvPillars (2022)	R	-	13.9	4.9	0.823	0.428	0.607	2.081	1.000	-
CenterNet (2019)	C	HGLS	40.0	33.8	0.658	0.255	0.629	1.629	0.142	-
FCOS3D \dagger (2021b)	C	R101	42.8	35.8	0.690	0.249	0.452	1.434	0.124	1.7
PGD \dagger (2021a)	C	R101	44.8	38.6	0.626	0.245	0.451	1.509	0.127	1.4
PETR (2022)	C	R101*	45.5	39.1	0.647	0.251	0.433	0.933	0.143	1.7
BEVFormer-S (2022)	C	R101*	46.2	40.9	0.650	0.261	0.439	0.925	0.147	-
CenterFusion \dagger (2021)	C+R	DLA34	44.9	32.6	0.631	0.261	0.516	0.614	0.115	-
CRAFT \dagger	C+R	DLA34	52.3	41.1	0.467	0.268	0.456	0.519	0.114	4.1

Table 1: State-of-the-art comparison on nuScenes `test` set. ‘L’, ‘C’, and ‘R’ represent LiDAR, camera, and radar, respectively. * are initialized with FCOS3D (Wang et al. 2021b) checkpoint and \dagger use image flipping test time augmentation.

3.4 Detection Heads and Training Objectives

The fusion detection head decodes the above radar-encoded image proposal feature to refine the image proposal localization and other attributes in polar coordinate. Specifically, category-specific regression heads on the top of shared MLP layers predict (a) fusion score, (b) location offsets, (c) center-ness, and (d) velocity.

(a) Fusion score: Due to radar sparsity, some image proposals are associated with only clutter radar points, which can degrade performance. Thus we predict the probability of radar points associated with SPA containing at least one radar point from the object.

(b) Location offsets: Given the image proposal center, we predict offsets in polar coordinate instead of Cartesian. We empirically find that this simple and easy to implement technique effectively reduces the significant localization error disagreement in polar coordinates.

(c) Center-ness: Following (Yang et al. 2020), we assign higher center-ness scores to predictions closer to ground-truth centers.

(d) Velocity: To mitigate the absence of azimuthal velocity in radar Doppler measurement (Long et al. 2021a), we predict the speed and then transform it to velocity using object orientation.

To train the model, we assign the same ground-truth used for training the keypoint-based camera 3D object detector (Law and Deng 2018) to image proposal; thus it can be directly matched to the ground truth without an additional matching algorithm (Carion et al. 2020). Our final loss function can be formulated as a weighted sum of aforementioned classification and regression losses. The full loss function is detailed in Appendix B.

4 Experiments

Dataset and Metrics We evaluate our method on a large-scale and challenging nuScenes dataset (Caesar et al. 2020), which consists of 700/150/150 scenes for `train/val/test` set. Each 20 second long sequence has 3D bounding box annotations of 10 classes by 2Hz frequency, and it contains six surrounding camera images, one LiDAR point cloud, and five radar point clouds covering 360 degrees. The official evaluation metrics are mean Average Precision (mAP), True Positive metrics (*i.e.*, trans-

lation, scale, orientation, velocity, and attribute error), and the nuScenes detection score (NDS). Particularly, NDS is a weighted sum of mAP and True Positive metrics. The matching thresholds for calculating AP are the center distance of $0.5m$, $1m$, $2m$, and $4m$, instead of IoU.

Implementation Details We implement our fusion network using CenterNet (Zhou, Wang, and Krähenbühl 2019) with DLA34 (Yu et al. 2018) backbone. Note that our camera 3D object detector is trained from scratch on nuScenes without large-scale depth pre-training on DDAD15M proposed in DD3D (Park et al. 2021). For the image backbone, we use 448×800 size image as network input and use a single feature map of the last layer for the image proposal feature. For radar, we accumulate six radar sweeps following GRIF Net (Kim, Choi, and Kum 2020) and set the maximum number of radar points as 2048. To maximize the recall of the camera 3D detector, we set the image proposal threshold to 0.05 and apply Non-Maximum Suppression (NMS) after fusion. More detailed implementation details are in Appendix B.

Training and Inference As a proof of concept that the proposed method is flexible, we pre-train the camera 3D object detector and keep its weights frozen during training the camera-radar fusion network, while other fusion modules are trained in an end-to-end manner. Our model takes six surrounding images and five radar point clouds as a single batch to perform 360-degree detection, and data augmentation is applied to both image and radar. We train our models for 24 epochs with a batch size of 32, cosine annealing scheduler, and 2×10^{-4} learning rate on 4 RTX 3090 GPUs. Inference time is measured on an Intel Core i9 CPU and an RTX 3090 GPU without test time augmentation for fusion. See Appendix C for more detailed experimental settings.

4.1 Comparison with State-of-the-Arts

We compare our method with state-of-the-arts on nuScenes `test` set. To eliminate the effect of multi-sequence input and large-scale depth pre-training and compare with previous methods fairly, we report methods with single frame input trained on nuScenes dataset only. Additional comparisons with other methods and results on `val` set are provided in Appendix D.

Method	Input	Car	Truck	Bus	Trailer	C.V.	Ped.	M.C.	Bicycle	T.C.	Barrier	mAP
PointPillars (2019)	L	79.9	35.7	42.8	26.1	5.5	71.7	39.4	10.6	33.4	52.0	39.7
FCOS3D (2021b)	C	47.9	23.3	31.4	11.2	5.7	41.1	30.5	30.2	55.0	45.5	32.2
CenterNet (2019)	C	48.4	23.1	34.0	13.1	3.5	37.7	24.9	23.4	55.0	45.6	30.6
CenterFusion (2021)	C+R	52.4(+4.0)	26.5(+3.4)	36.2(+2.2)	15.4(+2.3)	5.5(+2.0)	38.9(+1.2)	30.5(+5.6)	22.9(-0.5)	56.3(+1.3)	47.0(+1.4)	33.2(+2.6)
CRAFT-I	C	52.4	25.7	30.0	15.8	5.4	39.3	28.6	29.8	57.5	47.8	33.2
CRAFT	C+R	69.6(+17.2)	37.6(+11.9)	47.3(+17.3)	20.1(+4.3)	10.7(+5.3)	46.2(+6.9)	39.5(+10.9)	31.0(+1.2)	57.1(-0.4)	51.1(+3.3)	41.1(+7.9)

Table 2: Per-class AP comparison on nuScenes val set. ‘C.V.’, ‘M.C.’, and ‘T.C.’ denote construction vehicle, motorcycle, and traffic cone, respectively.

(a) Fusion method				(b) Association method				(c) Coordinate system				
	AP \uparrow		ATE \downarrow		AP \uparrow		Assoc. RC	AP \uparrow		Rad. Err.	Azim. Err.	
	0.5m	mean			0.5m	mean		0.5m	mean			
Img.	19.6	52.4	0.49	No Assoc.	OOM		83.5	Cart.	27.7	58.6	1.73	0.52
FP	34.7	61.8	0.37	RoIPool	29.8	58.7	0.39	Polar	41.3	65.1	1.26	0.25
FC	38.9	63.1	0.34	Ball Query	39.5	64.3	0.33	Imp.	+13.6	+6.5	-0.47	-0.27
FT	41.3	65.1	0.32	SPA	41.3	65.1	0.32					

Table 3: Ablation experiments to validate our fusion framework design choices.

As shown in Table 1, our CRAFT outperforms all competing camera-radar and single-frame camera methods on test set. Importantly, our CRAFT has more than two times faster inference speed compared to other methods. The most performance gain comes from the improved localization and velocity estimation, which improve the recall at strict thresholds (e.g., 0.5m or 1m) and NDS scores, respectively. As our lightweight CenterNet-like camera 3D detector, denote as CRAFT-I, has no bells and whistles as a proof of concept, we leave a more advanced camera 3D detector for camera-radar fusion as a future work.

Table 2 shows the per-class mAP comparison with LiDAR, camera, and camera-radar methods on val set. Although CenterFusion (Nabati and Qi 2021) and ours have a similar camera 3D detector architecture, our CRAFT achieves a remarkable performance boost (+17.2% vs. 4.0% on Car, +17.3% vs. 2.2% on Bus), which shows the importance of an appropriate fusion strategy considering sensor characteristics. The performance improvement is higher in metallic objects (car, truck, bus and motorcycle) than non-metallic objects (pedestrian, bicycle, traffic cone, and barrier) since metallic objects have more valid radar returns and are thus easier to be distinguished from background clutters. Trailer and construction vehicle are metallic but have less performance gain, assuming these classes are commonly surrounded by other trailers or construction materials, which can possibly harm the radar performance. More experiments of robustness against the number of radar points are provided in Appendix D.

4.2 Ablation Studies

We conduct a series of ablation studies to validate the design of CRAFT on nuScenes val set. We report the results of the car class for ablation experiments since the performance gains of other classes are less consistent and easier to be fluctuated by other factors.

Fusion Method On top of our CRAFT-I, we conduct additional fusion methods: using only radar points similar to

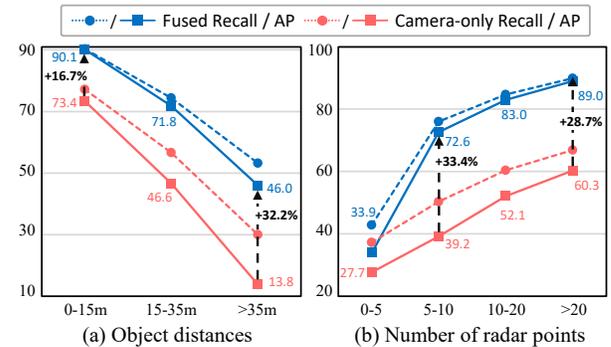


Figure 4: Analysis of different object distances and the number of radar points.

FPointNet (FP) (Qi et al. 2018) and fusing features by AP concatenation (FC) after aggregating radar features by max-pooling. Table 3a shows that our fusion transformer (FT) brings large AP improvement over radar-only (FP, +6.6%) and naive concatenation (FC, +2.4%), especially at the strict 0.5m threshold. It shows that using both image feature and attention-based fusion is beneficial for improving recall and reducing localization error.

Association Method In Table 3b, we ablate the proposed association strategy by replacing it with RoIPool (query points inside image proposal) and Ball Query (query points around image proposal center within a 6m radius) to study how the soft polar association benefits the following fusion method. Association recall denotes the percentage of a proposal containing at least one valid point over associated points. Taking all points to Radar-to-Image Feature Encoder without association suffers from a substantial computational burden, leading to an out-of-memory in our setting. RoIPool often fails to associate valid radar points due to inaccurate image proposal localization, and large size ball query con-

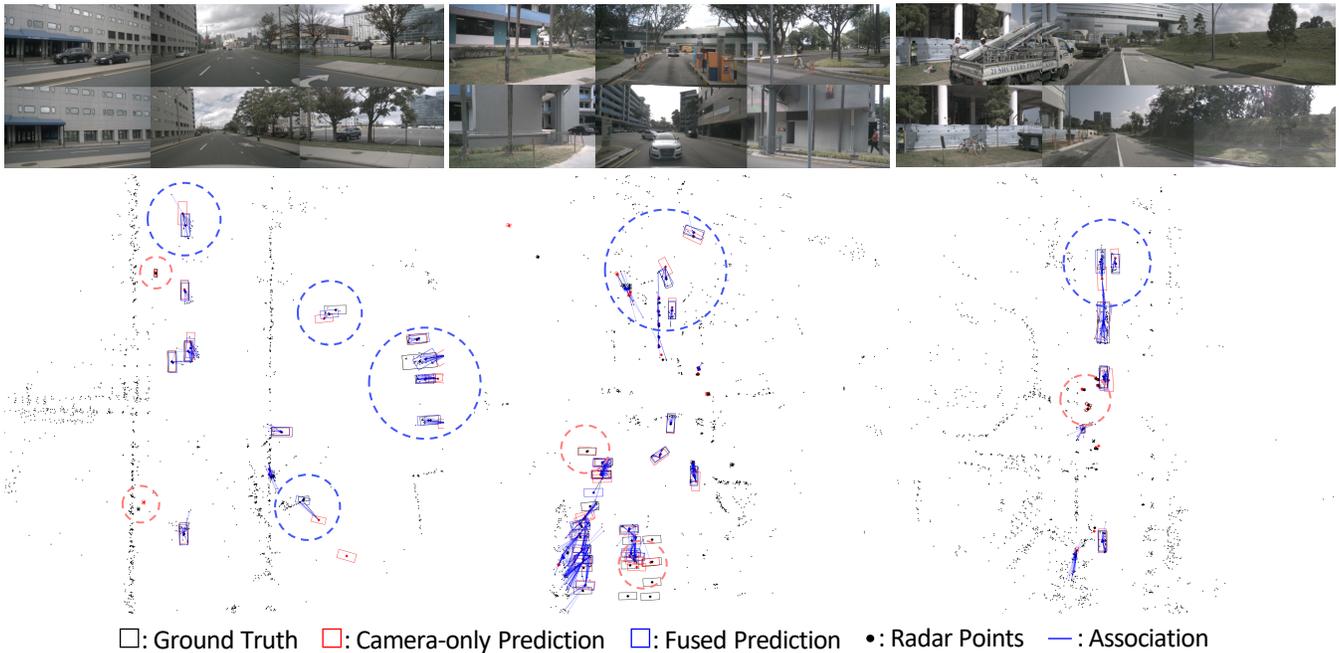


Figure 5: Qualitative results of CRAFT. Blue circles indicate samples that are refined by fusing radar points and have more accurate localization, and red circles indicate samples that are predicted by camera-only since there are no valid radar returns among associated points. Best viewed in color with zoom in.

tains many clutters that can degrade the fusion performance. Our SPA enables the best trade-off between recall and precision for effective camera-radar feature fusion.

Coordinate System Table 3c ablates the coordinate system of point association and regression target. For the Cartesian setting, we associate points using SPA and transform them back to the Cartesian coordinate before feeding them to R2I FE, and also regress offsets in the same coordinate. Using polar coordinate is similar to applying Principal Component Analysis (PCA) since the error variance of image proposals in the azimuthal direction is minimized in polar coordinate. Thus, the simple coordinate transformation makes the network learn spatial information easier and effectively reduces the localization error of image proposal.

4.3 Analysis

We analyze maximum recall and AP performance under different object distances and the number of radar points using a $1m$ matching threshold.

Performance by Distances Our camera-radar fusion achieves significant improvements over the camera-only baseline and brings larger improvements for distant objects, as shown in Fig. 4a. As the depth estimation becomes more inaccurate on distant objects farther than $35m$, our fusion method benefits more from the radar measurement and achieves a significant improvement of 32.2% .

Performance by Radar Points We demonstrate the performance under a different number of valid radar points re-

turned from the object in Fig. 4b. CRAFT achieves consistent improvements over a various number of radar points, which shows that our fusion method is robust to the radar sparsity. Our method yields robust performance although radar points are not available or only a few ($+6.2\%$) and shows better performance improvement when more radar points are provided ($+28.7\%$).

4.4 Qualitative Results

We show the detection results in a complex scene in Fig. 5. As highlighted with blue circles, CRAFT precisely refines image proposals by adaptively fusing valid radar returns if they are available. More qualitative results are provided in Appendix E.

5 Conclusions

In this paper, we have proposed an effective and robust camera-radar 3D detection framework. With a soft polar association and spatio-contextual fusion transformer, the spatial and contextual information of the camera and radar can be effectively complemented to yield a more accurate prediction given an inaccurate image proposal. CRAFT achieves state-of-the-art performance on nuScenes by outperforming previous camera-radar and camera-only methods and shows the potential of camera-radar fusion. Extensive experiments validate the design choice of our framework and show the advantages over a camera-only baseline. We hope our work will inspire future research on camera-radar fusion for 3D scene understanding.

Acknowledgments

This work was supported by Institute of Information & communications Technology Planning & Evaluation (IITP) grant funded by the Korea government (MSIT) (No.2021-0-00951, Development of Cloud based Autonomous Driving AI learning Software).

References

- Bai, X.; Hu, Z.; Zhu, X.; Huang, Q.; Chen, Y.; Fu, H.; and Tai, C.-L. 2022. Transfusion: Robust lidar-camera fusion for 3d object detection with transformers. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 1090–1099.
- Caesar, H.; Bankiti, V.; Lang, A. H.; Vora, S.; Liong, V. E.; Xu, Q.; Krishnan, A.; Pan, Y.; Baldan, G.; and Beijbom, O. 2020. nuScenes: A multimodal dataset for autonomous driving. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 11621–11631.
- Carion, N.; Massa, F.; Synnaeve, G.; Usunier, N.; Kirillov, A.; and Zagoruyko, S. 2020. End-to-end object detection with transformers. In *Proceedings of the European Conference on Computer Vision (ECCV)*, 213–229.
- Fu, H.; Gong, M.; Wang, C.; Batmanghelich, K.; and Tao, D. 2018. Deep Ordinal Regression Network for Monocular Depth Estimation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2002–2011.
- Huang, T.; Liu, Z.; Chen, X.; and Bai, X. 2020. Epnet: Enhancing point features with image semantics for 3d object detection. In *Proceedings of the European Conference on Computer Vision (ECCV)*, 35–52.
- Hung, W.-C.; Kretschmar, H.; Casser, V.; Hwang, J.-J.; and Angelov, D. 2022. Let-3d-ap: Longitudinal error tolerant 3d average precision for camera-only 3d detection. In *arXiv preprint arXiv:2206.07705*.
- Johnson, D. H.; and Dudgeon, D. E. 1992. *Array signal processing: concepts and techniques*. Simon & Schuster, Inc.
- Kendall, A.; and Gal, Y. 2017. What Uncertainties Do We Need in Bayesian Deep Learning for Computer Vision? In *Advances in Neural Information Processing Systems (NeurIPS)*, 5574–5584.
- Kim, J.; Kim, Y.; and Kum, D. 2020. Low-level sensor fusion network for 3D vehicle detection using radar range-azimuth heatmap and monocular image. In *Proceedings of the Asian Conference on Computer Vision (ACCV)*, 388–402.
- Kim, Y.; Choi, J. W.; and Kum, D. 2020. GRIF Net: Gated region of interest fusion network for robust 3D object detection from radar point cloud and monocular image. In *Proceedings of the IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 10857–10864.
- Ku, J.; Mozifian, M.; Lee, J.; Harakeh, A.; and Waslander, S. L. 2018. Joint 3d proposal generation and object detection from view aggregation. In *Proceedings of the IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 5750–5757.
- Lang, A. H.; Vora, S.; Caesar, H.; Zhou, L.; Yang, J.; and Beijbom, O. 2019. Pointpillars: Fast encoders for object detection from point clouds. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 12697–12705.
- Law, H.; and Deng, J. 2018. Cornernet: Detecting objects as paired keypoints. In *Proceedings of the European Conference on Computer Vision (ECCV)*, 734–750.
- Li, Z.; Wang, W.; Li, H.; Xie, E.; Sima, C.; Lu, T.; Yu, Q.; and Dai, J. 2022. BEVFormer: Learning Bird’s-Eye-View Representation from Multi-Camera Images via Spatiotemporal Transformers. In *Proceedings of the European Conference on Computer Vision (ECCV)*.
- Lim, T.-Y.; Ansari, A.; Major, B.; Fontijne, D.; Hamilton, M.; Gowaikar, R.; and Subramanian, S. 2019. Radar and camera early fusion for vehicle detection in advanced driver assistance systems. In *Advances in Neural Information Processing Systems Workshops (NeurIPSW)*.
- Lin, J.-T.; Dai, D.; and Van Gool, L. 2020. Depth estimation from monocular images and sparse radar data. In *Proceedings of the IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 10233–10240.
- Lin, Y.; Le Kernec, J.; Yang, S.; Fioranelli, F.; Romain, O.; and Zhao, Z. 2018. Human activity classification with radar: Optimization and noise robustness with iterative convolutional neural networks followed with random forests. *IEEE Sensors Journal*, 18(23): 9669–9681.
- Liu, Y.; Wang, T.; Zhang, X.; and Sun, J. 2022. PETR: Position Embedding Transformation for Multi-View 3D Object Detection. In *Proceedings of the European Conference on Computer Vision (ECCV)*, 531–548.
- Long, Y.; Morris, D.; Liu, X.; Castro, M.; Chakravarty, P.; and Narayanan, P. 2021a. Full-Velocity Radar Returns by Radar-Camera Fusion. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, 16198–16207.
- Long, Y.; Morris, D.; Liu, X.; Castro, M.; Chakravarty, P.; and Narayanan, P. 2021b. Radar-camera pixel depth association for depth completion. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 12507–12516.
- Ma, X.; Zhang, Y.; Xu, D.; Zhou, D.; Yi, S.; Li, H.; and Ouyang, W. 2021. Delving into Localization Errors for Monocular 3D Object Detection. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 4721–4730.
- Major, B.; Fontijne, D.; Ansari, A.; Teja Sukhvasi, R.; Gowaikar, R.; Hamilton, M.; Lee, S.; Grzechnik, S.; and Subramanian, S. 2019. Vehicle detection with automotive radar using deep learning on range-azimuth-doppler tensors. In *Proceedings of the IEEE/CVF International Conference on Computer Vision Workshops (ICCVW)*, 924–932.
- Meyer, M.; and Kuschik, G. 2019a. Automotive radar dataset for deep learning based 3d object detection. In *Proceedings of the European Radar Conference (EuRAD)*, 129–132.

- Meyer, M.; and Kuschik, G. 2019b. Deep learning based 3d object detection for automotive radar and camera. In *Proceedings of the European Radar Conference (EuRAD)*, 133–136.
- Misra, I.; Girdhar, R.; and Joulin, A. 2021. An end-to-end transformer model for 3d object detection. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, 2906–2917.
- Nabati, R.; and Qi, H. 2021. Centerfusion: Center-based radar and camera fusion for 3d object detection. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision (WACV)*, 1527–1536.
- Pan, X.; Xia, Z.; Song, S.; Li, L. E.; and Huang, G. 2021. 3d object detection with pointformer. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 7463–7472.
- Park, D.; Ambrus, R.; Guizilini, V.; Li, J.; and Gaidon, A. 2021. Is Pseudo-Lidar needed for Monocular 3D Object detection? In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, 3142–3152.
- Phillion, J.; and Fidler, S. 2020. Lift, splat, shoot: Encoding images from arbitrary camera rigs by implicitly unprojecting to 3d. In *Proceedings of the European Conference on Computer Vision (ECCV)*, 194–210.
- Qi, C. R.; Liu, W.; Wu, C.; Su, H.; and Guibas, L. J. 2018. Frustum pointnets for 3d object detection from rgb-d data. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 918–927.
- Qi, C. R.; Su, H.; Mo, K.; and Guibas, L. J. 2017a. Pointnet: Deep learning on point sets for 3d classification and segmentation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 652–660.
- Qi, C. R.; Yi, L.; Su, H.; and Guibas, L. J. 2017b. Pointnet++: Deep hierarchical feature learning on point sets in a metric space. In *Advances in Neural Information Processing Systems (NeurIPS)*, 5105–5114.
- Reading, C.; Harakeh, A.; Chae, J.; and Waslander, S. L. 2021. Categorical depth distribution network for monocular 3d object detection. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 8555–8564.
- Ren, S.; He, K.; Girshick, R.; and Sun, J. 2015. Faster R-CNN: Towards Real-Time Object Detection with Region Proposal Networks. In *Advances in Neural Information Processing Systems (NeurIPS)*, 91–99.
- Shi, S.; Wang, X.; and Li, H. 2019. Pointtrnn: 3d object proposal generation and detection from point cloud. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 770–779.
- Svenningsson, P.; Fioranelli, F.; and Yarovoy, A. 2021. Radar-pointgcn: Graph based object recognition for unstructured radar point-cloud data. In *Proceedings of the IEEE Radar Conference (RadarConf)*, 1–6.
- Thomas, H.; Qi, C. R.; Deschaud, J.-E.; Marcotegui, B.; Goulette, F.; and Guibas, L. J. 2019. Kpconv: Flexible and deformable convolution for point clouds. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, 6411–6420.
- Tian, Z.; Shen, C.; Chen, H.; and He, T. 2019. Fcos: Fully convolutional one-stage object detection. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, 9627–9636.
- Ulrich, M.; Braun, S.; Köhler, D.; Niederlöhner, D.; Faion, F.; Gläser, C.; and Blume, H. 2022. Improved Orientation Estimation and Detection with Hybrid Object Detection Networks for Automotive Radar. In *arXiv preprint arXiv:2205.02111*.
- Vaswani, A.; Shazeer, N.; Parmar, N.; Uszkoreit, J.; Jones, L.; Gomez, A. N.; Kaiser, Ł.; and Polosukhin, I. 2017. Attention is all you need. In *Advances in Neural Information Processing Systems (NeurIPS)*, 6000–6010.
- Vora, S.; Lang, A. H.; Helou, B.; and Beijbom, O. 2020. Pointpainting: Sequential fusion for 3d object detection. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 4604–4612.
- Wang, T.; Xinge, Z.; Pang, J.; and Lin, D. 2021a. Probabilistic and geometric depth: Detecting objects in perspective. In *Proceedings of the Conference on Robot Learning (CoRL)*, 1475–1485.
- Wang, T.; Zhu, X.; Pang, J.; and Lin, D. 2021b. Fcos3d: Fully convolutional one-stage monocular 3d object detection. In *Proceedings of the IEEE/CVF International Conference on Computer Vision Workshops (ICCVW)*, 913–922.
- Xiong, R.; Yang, Y.; He, D.; Zheng, K.; Zheng, S.; Xing, C.; Zhang, H.; Lan, Y.; Wang, L.; and Liu, T. 2020. On layer normalization in the transformer architecture. In *Proceedings of the International Conference on Machine Learning (ICML)*, 10524–10533.
- Yang, Z.; Sun, Y.; Liu, S.; and Jia, J. 2020. 3dssd: Point-based 3d single stage object detector. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 11040–11048.
- Yin, T.; Zhou, X.; and Krahenbuhl, P. 2021. Center-based 3d object detection and tracking. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 11784–11793.
- Yoo, J. H.; Kim, Y.; Kim, J.; and Choi, J. W. 2020. 3d-cvf: Generating joint camera and lidar features using cross-view spatial feature fusion for 3d object detection. In *Proceedings of the European Conference on Computer Vision (ECCV)*, 720–736.
- Yu, F.; Wang, D.; Shelhamer, E.; and Darrell, T. 2018. Deep Layer Aggregation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2403–2412.
- Zhou, X.; Wang, D.; and Krähenbühl, P. 2019. Objects as Points. In *arXiv preprint arXiv:1904.07850*.
- Zhu, X.; Su, W.; Lu, L.; Li, B.; Wang, X.; and Dai, J. 2021. Deformable detr: Deformable transformers for end-to-end object detection. In *Proceedings of the International Conference on Learning Representations (ICLR)*.