

CoIn3D: Revisiting Configuration-Invariant Multi-Camera 3D Object Detection

Zhaonian Kuang^{1,2*} Rui Ding¹ Haotian Wang² Xinhu Zheng^{2†} Meng Yang^{1†} Gang Hua³

¹ National Key Laboratory of Human-Machine Hybrid Augmented Intelligence,
Institute of Artificial Intelligence and Robotics, Xi'an Jiaotong University

² Intelligent Transportation Thrust of the Systems Hub, HKUST(GZ) ³ Amazon Alexa AI

Abstract

Multi-camera 3D object detection (MC3D) has attracted increasing attention with the growing deployment of multi-sensor physical agents, such as robots and autonomous vehicles. However, MC3D models still struggle to generalize to unseen platforms with new multi-camera configurations. Current solutions simply employ a meta-camera for unified representation but lack comprehensive consideration for the devil lies in spatial prior discrepancies across source and target configurations, including different intrinsics, extrinsics, and array layouts. To address this, we propose **CoIn3D**, a generalizable MC3D framework that enables strong transferability from source configurations to unseen target ones. CoIn3D explicitly incorporates all identified spatial priors into both feature embedding and image observation through **spatial-aware feature modulation (SFM)** and **camera-aware data augmentation (CDA)**, respectively. SFM enriches feature space by integrating four spatial representations, such as focal length, ground depth, ground gradient, and Plücker coordinate. CDA improves observation diversity under various configurations via a training-free dynamic novel-view image synthesis scheme. Extensive experiments demonstrate that CoIn3D achieves strong cross-configuration performance on landmark datasets such as NuScenes, Waymo, and Lyft, under three dominant MC3D paradigms represented by BEVDepth, BEVFormer, and PETR.

1. Introduction

Driven by the demand for reliable target object localization, multi-camera 3D object detection (MC3D) [9, 18, 21–23, 38, 42] has gained increasing attention in recent years. MC3D has been widely adopted in physical agents such as

Codes are available [here](#).

* The author was funded while visiting HKUST(GZ).

† Corresponding Authors.

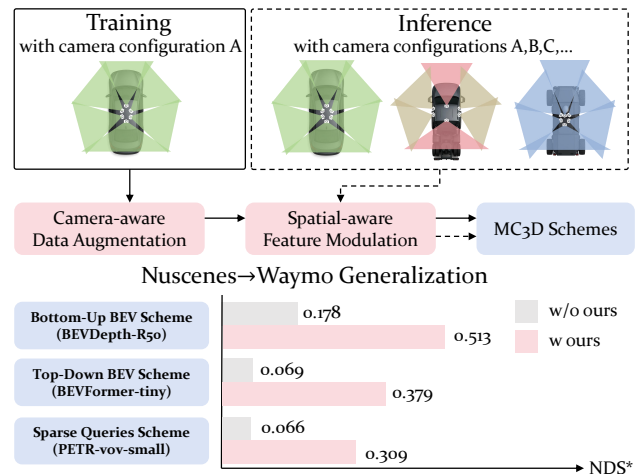


Figure 1. CoIn3D effectively enables model transferability from source configuration A to unseen target configurations B, C, ..., covering variations in intrinsics, extrinsics, and array layouts. Our framework can be applied to three dominant MC3D paradigms, represented by BEVDepth [18], BEVFormer [21], and PETR [23].

autonomous vehicles [3, 17] and robots [4]. Compared with explicit LiDAR-based counterparts [8, 12, 16, 19, 24, 45], vision-based methods rely on multi-camera configurations, including intrinsics, extrinsics, and array layouts, to implicitly perceive spatial structures from image observations.

Although current MC3D methods employ camera configurations to help connect image observations and spatial perception, their focus primarily lies in seeking vision invariance for unified image representation rather than achieving camera invariance for cross-configuration applications. In this paper, we present that the latter also plays a crucial role in MC3D model generalization.

Specifically, the configuration gap makes MC3D models perform poorly when tested on unseen configurations, leading to inflexibility and high costs in real-world deployment. A platform with a new configuration requires recollecting and reannotating data to retrain a specialized model.

To tackle the generalization problem caused by the con-

figuration gap, some previous works [20, 43, 47] directly warp images into meta-camera. Specifically, they resize and crop the image to align focal length and field-of-view (FoV), and then warp image based on spherical [43] or cylindrical [20] 3D space assumption to align extrinsics and arrays. However, warping input image causes resolution reduction and 3D scene structure distortion.

Some other works [2, 26, 37] treat all cameras as possessing a virtual meta-focal and rescale the predicted depth during training and testing. PD-BEV [26] and UDGA-BEV [2] further constrain feature and depth consistency for better generalization. In addition, [15] introduces object bottom ground depth to address camera height gaps. However, these methods do not fully and explicitly consider camera configurations, and they cannot be applied to all MC3D paradigms due to their depth-based design.

This paper comprehensively revisits the impact of configurations on MC3D generalization. We identify that the devil lies in the spatial prior discrepancies across source and target configurations. First, regarding intrinsics, different focal lengths cause object pixel size ambiguity and different FoVs lead to different scene perspective. Second, regarding extrinsic, different camera mounting positions and orientations alter ground plane geometries and observed scene geometries. Both of these two factors affect model’s spatial understanding. Lastly, different array layouts cause different camera numbers and overlap region, which affect the pattern of multi-camera correlation and feature fusion.

To address this, we propose CoIn3D, a generalizable MC3D framework that enables effective transferability from source configurations to unseen target configurations. CoIn3D explicitly incorporates identified spatial priors into both feature embedding and image observation through spatial-aware feature modulation (SFM) and camera-aware data augmentation (CDA), respectively.

Our SFM enriches the feature space by explicitly integrating four spatial representations. Specifically, we introduce four over-complete pixel-level spatial embeddings that formulate camera configurations. Among them, the inverse focal map formulates focal length, the ground depth and gradient maps jointly formulate ground geometry, and the Plücker raymap provides a holistic representation of camera configurations.

We first obtain focal-invariant features by multiplying the inverse focal map with image features. Then, we concatenate the ground depth, gradient, and Plücker raymap into a mixture prior map. The mixture prior map is encoded into a spatial feature embedding and added to the focal-invariant feature to obtain spatial-embedded feature. We further concatenate four raw prior maps into the spatial-embedded feature to obtain spatial-aware feature. This spatial-aware feature can be used by all MC3D schemes.

Our CDA introduces a cost-efficient, training-free novel-view image synthesis scheme based on 3D Gaussian splat-

ting (3DGS) to further enhance the generalization of MC3D. Specifically, we reconstruct ego-centric texture point clouds for each annotated data frame and transform them into a 3D Gaussian representation using predefined parameters. By randomly sampling cameras with diverse configurations, these Gaussians can be dynamically rendered into novel-view images for training.

Our framework can be widely applied to all dominant MC3D paradigms, including bottom-up BEV, top-down BEV, and sparse queries schemes, for which we choose three representative base models: BEVDepth [18], BEVFormer [21], and PETR [23], respectively.

Our framework is evaluated across three real-world landmark datasets with different camera configurations, including Nuscenes [1], Waymo [32], and Lyft [6]. Experiments show that our scheme achieves significant gains on all paradigms under different cross-dataset settings as shown in Fig. 1. In addition, our scheme achieves state-of-the-art (SOTA) performance in all settings based on BEVDepth.

The main contributions can be summarized as follows:

- We revisit the influence of multi-camera configurations, and identify that the devil lies in spatial prior discrepancies across source and target configurations.
- We propose SFM to enrich feature by explicitly integrating four spatial representations, including focal length, ground depth, ground gradient, and Plücker coordinate.
- We propose CDA, a cost-efficient, training-free novel-view image synthesis scheme based on 3D Gaussian splatting, to dynamically render augmented training images with diverse configurations.
- Our framework consistently achieves significant gains on all MC3D paradigms under different cross-dataset settings and achieves SOTA based on BEVDepth.

2. Related Work

2.1. Multi-Camera 3D Object Detection

Current MC3D models fall into three paradigms.

Bottom-up BEV schemes [9, 10, 18, 28, 31, 44] splat the image features into BEV feature, and then objects are detected on the BEV feature. Among them, BEVDepth [18] and BEVDet4D [9] are two representative schemes, which consider LiDAR depth supervision and temporal fusion.

Top-down BEV schemes, represented by BEVFormer [21, 42], update BEV query features by applying spatial cross-attention with image features and temporal self-attention with adjacent BEV features. Then, the detection results will be decoded on BEV features.

Sparse queries schemes [22, 23, 27, 36, 38] encode 3D scenes into sparse representations. Among them, PETR [23] injects 3D position embeddings into image features and StreamPETR [36] explores long sequence modeling.

These MC3D paradigms have their own advantages. While BEV feature-based schemes can provide dense scene

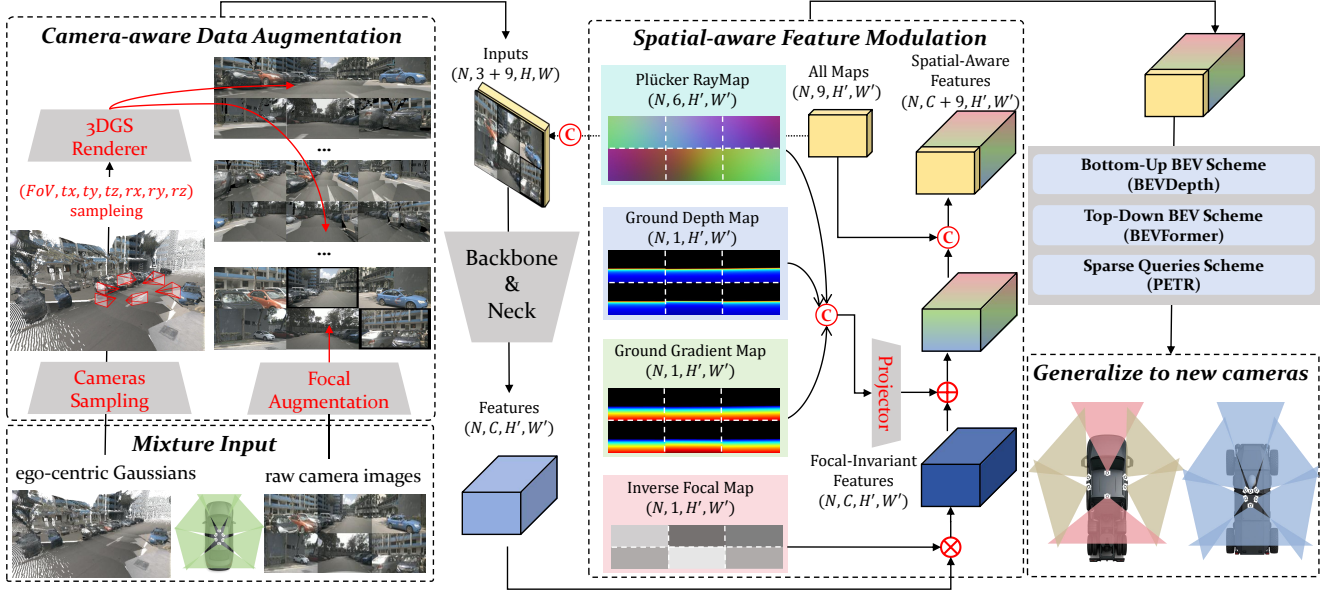


Figure 2. Illustration of our CoIn3D framework for generalizable MC3D across multi-camera configurations. During training, we apply the camera-aware data augmentation (CDA) to generate N images with randomly sampled camera configurations, followed by spatial-aware feature modulation (SFM). SFM modulates activations using an inverse focal map to obtain focal-invariant features, then projects prior maps (ground depth, gradient map, and Plücker raymap) to create spatial embeddings, which are added to the focal-invariant features. These maps are concatenated with image input and features to provide raw priors. Finally, the spatial-aware features can be easily integrated into MC3D for downstream tasks. During inference, we use raw images and apply spatial-aware modulation to generalize to new camera configurations. Our framework is applicable to dominant MC3D paradigms, including bottom-up BEV, top-down BEV, and sparse-queries.

representation for holistic 3D understanding, sparse queries offer lower computational costs. In this paper, we seek a unified framework to improve the generalization of all these dominant paradigms across camera configurations.

2.2. Camera Configuration Generalization

To tackle the generalization problem caused by the configuration gap, on one hand, some works try to adjust the image input [20, 43, 47], and on the other hand, some works try to adjust the predicted target [2, 15, 26, 37].

For image input adjustment, [20, 43, 47] resize and crop the image to align focal length and FoV. To adjust the extrinsic gap, [43] and [20] warp image based on spherical and cylindrical space assumption, respectively. However, resizing image will cause image shape misalignment and cannot take advantage of batch acceleration. Resizing long-focal images to short focal lengths will result in a loss of resolution. Cropping image will loss texture information. Warping images based on spherical or cylindrical space assumption will destroy the 3D scene structure.

For prediction adjustment, DG-BEV, PD-BEV and UDGA-BEV [2, 26, 37] treat all cameras as having an identical virtual focal length in both the training and testing data, and rescale the virtual depth to real depth according to the ratio between virtual and real focal lengths. PD-BEV further constrains feature consistency between image features and BEV feature, while UDGA-BEV further con-

strains depth and photometric consistency between different cameras to enhance the generalization of models. [15] introduces the ground depth of the object bottom center to help model overcome camera height gap. However, these methods do not fully and explicitly consider camera configurations. Furthermore, these methods cannot apply to all MC3D paradigms due to their depth-based design.

In this paper, we comprehensively revisit the influence of multi-camera configurations on MC3D and then propose CoIn3D to uniformly mitigate the camera configuration gap. Our scheme can be applied to all MC3D paradigms.

2.3. 3D Gaussian splatting

Novel view synthesis [14] can augment images under different configurations, showing the potential to achieve configuration-invariant MC3D. 3DGS [13] has a fast rendering speed and makes dynamic augmentation possible. However, current 3DGS schemes [5, 11, 25, 39, 41] have high training costs. Unlike previous works, we propose a cost-efficient, training-free ego-centric Gaussian construction pipeline tailored for MC3D data augmentation.

3. Revisit cameras configuration in MC3D

3.1. MC3D task

MC3D takes N surrounding camera images $I = \{i_1, i_2, \dots, i_N\} \in \mathbb{R}^{N \times 3 \times H \times W}$ as input. Each camera has

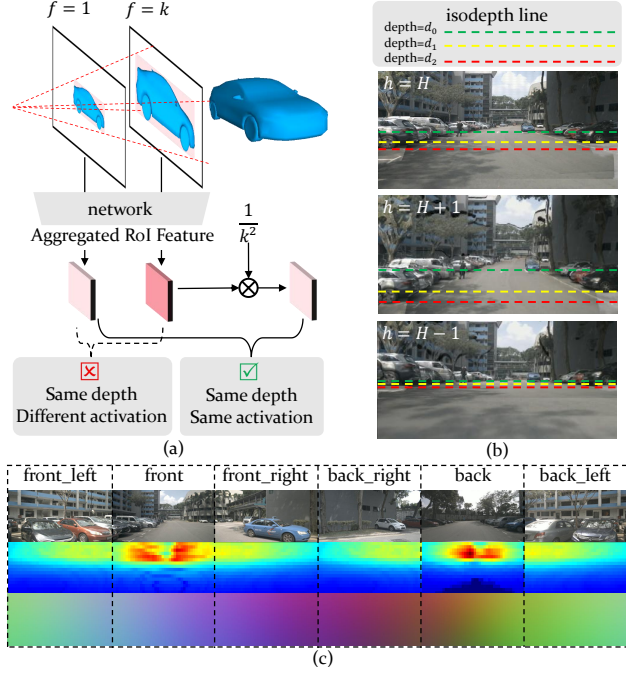


Figure 3. Illustration of the spatial discrepancies under different camera configurations: (a) focal-ambiguity for a same object; (b) ground depth and depth increasing rate under different camera heights; (c) the scene structure (1st row), depth distribution (2nd row), and Plücker raymap (3rd row) for surround-view cameras.

an intrinsic matrix $\mathbf{K} \in \mathbb{R}^{3 \times 3}$ and an extrinsic matrix $\mathbf{T} = [\mathbf{R}; \mathbf{t} | \mathbf{0}, 1] \in \mathbb{R}^{4 \times 4}$, which can transform points from the camera system to the ego system. Taking all these as input, MC3D aims to localize and categorize objects in 3D ego space. MC3D needs to predict position (x, y, z) , size (l, w, h) , orientation θ , and category of each object.

3.2. Intrinsic revisiting

Different cameras has different focal length and field-of-view (FoV). For focal length, as shown in Fig. 3(a), the same object under different focal plane has different sizes, which hinders the model’s consistent depth understanding. For FoV, different FoVs corresponds to different scene perspective geometries. For example, given a camera height H_{cam} , the initial ground depth z_i on image bottom is determined by the FoV angle α : $z_i = \frac{H_{cam}}{\tan(\frac{\alpha}{2})}$. This prior can help model understand scene structure.

3.3. Extrinsic revisiting

Different camera mounting positions and orientations results in varying extrinsics and spatial priors.

First, ground plane can provide a strong prior for MC3D. We assume the ground plane is flat, so we can use N non-collinear points ($N \geq 3$) $\mathbf{p}_e \in \mathbb{R}^{N \times 4}$ on the ground to define a plane in ego space. The ground plane equation $Ax + By + Cz + D = 0$ in each camera space can be

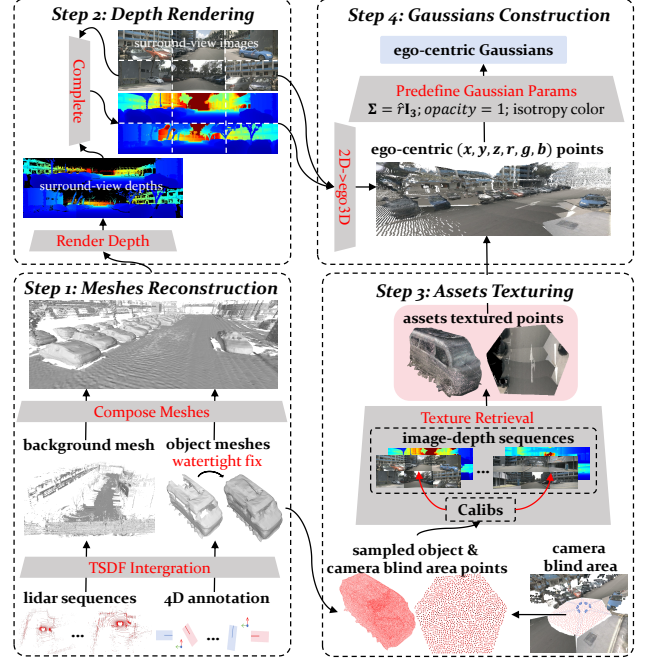


Figure 4. Illustration of the training-free ego-centric Gaussians construction pipeline. We transform reconstructed texture point clouds into Gaussian representations using predefined parameters.

solved by Least Squares from points $\mathbf{p}_c = \mathbf{T}^{-1}\mathbf{p}_e$. For an image with intrinsic parameters f_u, f_v, c_u, c_v , if we assume that a pixel (u, v) with depth $z(u, v)$ is on the ground, then the depth can be solved by:

$$z(u, v) = -\frac{D}{AX + BY + C} \quad (1)$$

where $X = \frac{u-c_u}{f_u}$ and $Y = \frac{v-c_v}{f_v}$.

Different camera mounting heights also lead to varying perspective effects. As shown in Fig.3(b), if we observe the ground in the image from near to far, we will find that, higher camera mounting height has slower ground depth increase rate, while lower height causes a faster increase. Training on one specific camera height will cause overfitting to this rate change behavior.

Finally, different camera orientations result in different observed scene geometry. As shown in Fig. 3(c), each camera captures different depth distribution and scene structures, which should be considered as priors.

3.4. Array revisiting

Different platforms will adopt camera arrays with different numbers of cameras and cameras layouts, resulting in different surround-view structures. As shown in Fig. 3(c), each surround-view image can be seen as a fragment of the omnidirectional image of the scene. Therefore, the scene structure is continuous across different images. Moreover, different layouts result in different overlap regions. These

two factors should be considered by MC3D during multi-camera correlation and feature fusion.

4. Methodology

We identify that discrepancies in intrinsics, extrinsics, and array layouts are the key challenge for cross-configuration generalization. To address this, we propose CoIn3D, a framework for generalizing MC3D across configurations.

The framework is shown in Fig. 2. During training, taking raw images and ego-centric Gaussians as input, we first apply camera-aware data augmentation (CDA) to dynamically generate images with diverse camera configurations. After applying CDA, the augmented images $\mathbf{I}_{im} \in \mathbb{R}^{3 \times H \times W}$ are fed into the backbone and neck to extract features $\mathbf{F}_{im} \in \mathbb{R}^{C \times H' \times W'}$. We then apply spatial-aware feature modulation (SFM) to enrich the image feature by embedding four spatial prior maps formulating configurations. Finally, the modulated spatial-aware feature can be integrated into any MC3D paradigm for downstream tasks. During inference, the model only takes raw images as input and applies SFM to generalize to new camera configurations.

4.1. Spatial-aware feature modulation

In this section, we introduce how we modulate the raw image features to obtain spatial-aware features.

4.1.1. Inverse focal map

We first address focal ambiguity by assuming that features with different focal lengths should have similar activations. Larger focal lengths can be seen as upsampled versions of smaller ones, leading to different aggregated RoI feature activations for the same object, which causes ambiguity.

To eliminate the ambiguity, we propose using the square of the focal length to normalize the feature activation, which is based on the observation that a k -fold difference in focal length leads to a k^2 -fold difference in pixel size for the same object. Specifically, for a raw image feature \mathbf{F}_{im} with focal length f , we use the inverse of the square focal map $\mathbf{M}_{IF} \in \mathbb{R}^{1 \times H' \times W'}$ to normalize as follows:

$$\begin{aligned} \mathbf{M}_{IF} &= \mathbf{1} * \frac{1}{f^2} \\ \mathbf{F}_{im}^1 &= \mathbf{M}_{IF} \odot \mathbf{F}_{im} \end{aligned} \quad (2)$$

4.1.2. Ground depth and gradient map

After addressing the focal gap, we consider the ground prior under different camera configurations. The ground depth map $\mathbf{M}_{GD} \in \mathbb{R}^{1 \times H' \times W'}$ can be generated by Eq. 1 and provides a straightforward scene spatial prior.

We further propose the ground gradient map. It is based on the observation that, under different camera mounting heights, the ground depth increase rating from near to far will be quite different as shown in Fig. 3 and results in

different perspective effects. To formulate, we introduce ground gradient map $\mathbf{M}_{GG} \in \mathbb{R}^{1 \times H \times W}$, which can be derived from \mathbf{M}_{GD} through cross-row difference: $\mathbf{M}_{GG} = \mathbf{M}_{GD}[:, :, -1, :] - \mathbf{M}_{GD}[:, 1, :, :]$. However, the numerical values of the results are small and vary rapidly, which is unfavorable for network training. To make the numerical value suitable for network, we apply a log-inverse transformation to get the final ground gradient map:

$$\mathbf{M}_{GG} = \log\left(\frac{1}{\mathbf{M}_{GD}[:, :, -1, :] - \mathbf{M}_{GD}[:, 1, :, :]} + 1\right) \quad (3)$$

4.1.3. Plücker raymap

To further enriching the spatial representation, we introduce the Plücker raymap inspired by [29, 46]. The Plücker raymap describes the direction and moment of a cluster of rays emitted from the optical center of camera to each pixel.

Specifically, for a camera with an intrinsic matrix $\mathbf{K} \in \mathbb{R}^{3 \times 3}$ and extrinsic matrix, which describes the transformation from camera to ego $\mathbf{T} = [\mathbf{R}, \mathbf{t}; \mathbf{0}, 1] \in \mathbb{R}^{4 \times 4}$, we can construct the Plücker raymap as follows. For a pixel, denoted its homogeneous coordinates as $\mathbf{p} = [u, v, 1]^T$, its ray direction $\mathbf{d} \in \mathbb{R}^3$ in ego and ray moment related to the camera origin $\mathbf{m} \in \mathbb{R}^3$ can be computed as follows:

$$\begin{aligned} \mathbf{d} &= \mathbf{R}\mathbf{K}^{-1}\mathbf{p} \\ \mathbf{m} &= \mathbf{t} \times \mathbf{d} \end{aligned} \quad (4)$$

where \mathbf{d} together with \mathbf{m} construct the Plücker coordinate $\mathbf{r} = (\mathbf{d}, \mathbf{m}) \in \mathbb{R}^6$ of each pixel in the Plücker raymap $\mathbf{M}_{PR} \in \mathbb{R}^{6 \times H' \times W'}$.

The Plücker raymap has several beneficial properties. First, it formulates camera's FoV, rotation, and translation in the ego system. Second, it continuously formulates the pixel position among multi-camera images which can utilized for camera correlation and feature fusion.

4.1.4. Spatial embedding module

The ground depth map, gradient map and Plücker raymap together provide an over-complete representation that describes camera configurations. To embed these prior maps into the feature, similar to positional embeddings in transformers [33], we first concatenate all these prior maps channel-wise. Then we use a shallow projector to project it to higher-level feature space to obtain spatial embedding. Finally, we add spatial embedding to the image feature. This process can be described as follows:

$$\mathbf{F}_{im}^2 = \mathbf{F}_{im}^1 + projector(cat(\mathbf{M}_{GD}, \mathbf{M}_{GG}, \mathbf{M}_{PR})) \quad (5)$$

In our implementation, the projector has a simple architecture with a 3×3 convolution and a ReLU activation layer. Finally, all prior maps are concatenated with the modulated feature $\mathbf{F}_{im}^3 = cat(\mathbf{M}_{IF}, \mathbf{M}_{GD}, \mathbf{M}_{GG}, \mathbf{M}_{PR}, \mathbf{F}_{im}^2)$ and image input to provide raw prior information. The spatial-aware feature \mathbf{F}_{im}^3 can be used by all kinds of MC3D schemes to decode the 3D detection results.

4.2. Camera-aware data augmentation

Data augmentation is important for enhancing the generalization of MC3D. However, basic augmentations like flipping and photometric distortion cannot augment the camera configuration. To address it, we propose a training-free 3DGS scheme to construct the ego-centric Gaussians for dynamic data augmentation in a cost-efficient manner.

4.2.1. Training-free 3DGS construction

Specifically, the training-free 3DGS construction pipeline can be divided into four steps, as shown in Fig. 4. First, we use 4D annotations to decompose LiDAR sequences into background and objects. Next, we use TSDF integration [34] to reconstruct the background mesh and object meshes. Moreover, we fix the object meshes to watertight surface.

Second, for each annotated timestamp in the sequence, we compose the background mesh and objects meshes according to calibration and annotations. Next, we render the composed meshes to depth [30] for each camera. These mesh-rendered depths have an advantage that their metric is naturally precise and they can ensure surround-view consistency. These merits are well-suited for MC3D data augmentation. Finally, we apply depth completion [35] to fill in holes without meshes and get dense depths.

Third, we reconstruct auxiliary assets, including the texture point model of objects and camera blind area. These assets help complete the unseen parts beyond raw imaging. Specifically, we first sample points from object meshes and camera blind area. Then we warp and retrieve textures through depth matching across images in the sequence.

Fourth, we construct the ego-centric Gaussians for each annotated data frame. We first project RGB-D images to ego space to obtain texture point cloud. Then, we append the texture assets. Finally, we set the covariance matrix of Gaussian as $\Sigma = \hat{r}\mathbf{I}_3$, which means that each Gaussian has no rotation and has a fixed predefined radius. The opacity of each Gaussian is set to 1. We set an isotropy color for each Gaussian, which meaning that we use the raw (r, g, b) color, the center of Gaussian is determined by the (x, y, z) coordinates of the point clouds. In other words, our scheme can be seen as a point-rendering scheme but we transform point cloud into Gaussian representation to take advantage of the fast rendering speed of 3DGS. The ego-centric Gaussians can be rendered at about 450 fps. We provide more details about this pipeline in the supplementary material.

4.2.2. Data augmentation

We will apply our data augmentation to both the ego-centric Gaussians and raw image inputs. These data will be dynamically augmented during training.

Specifically, for ego-centric Gaussians, we randomly sample new camera configurations to render novel-view images. For raw images, we randomly resize them according to the principal point and adjust the focal length.

5. Experiment

We conduct comprehensive experiments across multiple landmark datasets based on different MC3D paradigms.

5.1. Experiment Setup

Datasets. We conduct experiments across three landmark datasets, including Nuscenes [1], Waymo [32], and Lyft [6]. These datasets are collected with quite different camera configurations, making them challenging for evaluating the generalization ability of MC3D across configurations. Our model is trained on one dataset and tested on the others for each experiment. Detailed dataset information is provided in the supplementary material.

Evaluation Metrics. Following previous work [2, 26, 37], we adopt the metric NDS* that aggregates mean Average Precision (mAP), mean Average Translation Error (mATE), mean Average Scale Error (mASE), and mean Average Orientation Error (mAOE):

$$\text{NDS}^* = \frac{1}{6} [3\text{mAP} + \sum_{\text{mTP} \in \text{TP}} (1 - \min(1, \text{mTP}))] \quad (6)$$

Following previous work [2], we consider the unified category “car” for generalization experiment. Moreover, the “car”, “truck”, “construction vehicle”, “bus” and “trailer” in Nuscenes will all be considered as “car”, except for the Lyft \rightarrow Nuscenes evaluation. We only validate results in the range $[-50m, 50m]$ during training and validation.

Implementation Details. We adopt BEVDepth-R50 [18], BEVFormer-tiny [21] and PETR-vov [23] as base models to represent the three MC3D paradigms, respectively. We train the model on 4 RTX 4090 GPUs. The image width will be downsampled to 704 for BEVDepth and 800 for BEVFormer and PETR. The image height will be downsampled by the same ratio as the width. More implementation details are shown in the supplementary material.

5.2. Main Results

5.2.1. Results based on BEVDepth

Tab. 1 verifies the effectiveness of our scheme on bottom-up BEV framework (e.g., BEVDepth-R50). Our scheme significantly enhances the generalization of BEVDepth across datasets with different camera configurations compared to the base model. Moreover, our scheme achieves SOTA performance compared to current methods.

Specifically, directly transfer BEVDepth to new configuration leads to poor performance. Especially in Nuscenes \rightarrow Waymo and Waymo \rightarrow Nuscenes, the mAP drops to nearly 0 because these two datasets have a huge camera configuration gap. With our framework, the generalization performance NDS* gap is bridged from 0.178 to 0.513, 0.296 to 0.534, 0.133 to 0.481, and 0.213 to 0.452 for Nuscenes \rightarrow Waymo, Nuscenes \rightarrow Lyft, Waymo \rightarrow Nuscenes and Lyft \rightarrow Nuscenes settings, respectively.

Table 1. Comparison of our scheme with existing methods based on BEVDepth. “Nuscenes → Waymo” refers to train on Nuscenes and test on Waymo with different configurations, and similarly for the other datasets. “Direct Transfer” refers to directly use BEVDepth to test on target dataset without any adaptation. “Oracle” refers to training and testing on target dataset, serving as an upper bound for performance.

Setting	Nuscenes → Waymo					Waymo → Nuscenes				
Method	mAP↑	mATE↓	mASE↓	mAOE↓	NDS*↑	mAP↑	mATE↓	mASE↓	mAOE↓	NDS*↑
Oracle	0.552	0.528	0.148	0.085	0.649	0.475	0.577	0.177	0.147	0.587
Direct Transfer	0.040	1.303	0.265	0.790	0.178	0.032	1.305	0.768	0.532	0.133
CAM-Convs[7]	0.045	1.301	0.253	0.773	0.185	0.038	1.308	0.316	0.506	0.215
Single-DGOD[40]	0.034	1.305	0.262	0.855	0.164	0.014	1.000	1.000	1.000	0.007
DG-BEV[37]	0.297	0.822	0.216	0.372	0.415	0.303	0.689	0.218	0.171	0.472
UDGA-BEV[2]	0.349	0.754	0.289	0.250	0.459	0.326	0.684	0.263	0.168	0.477
Ours	0.381	0.687	0.220	0.155	0.513	0.349	0.727	0.257	0.179	0.481

Setting	Nuscenes → Lyft					Lyft → Nuscenes				
Method	mAP↑	mATE↓	mASE↓	mAOE↓	NDS*↑	mAP↑	mATE↓	mASE↓	mAOE↓	NDS*↑
Oracle	0.602	0.471	0.152	0.078	0.684	0.475	0.577	0.177	0.147	0.587
Direct Transfer	0.112	0.997	0.176	0.389	0.296	0.102	1.143	0.239	0.789	0.213
CAM-Convs	0.145	0.999	0.173	0.368	0.316	0.098	1.198	0.209	1.064	0.181
Single-DGOD	0.159	0.949	0.174	0.358	0.332	0.105	1.166	0.222	0.905	0.198
DG-BEV	0.287	0.771	0.170	0.302	0.437	0.268	0.764	0.205	0.591	0.374
PD-BEV[26]	0.304	0.709	0.169	0.289	0.458	0.263	0.746	0.186	0.790	0.344
UDGA-BEV	0.324	0.709	0.162	0.180	0.487	0.281	0.759	0.183	0.377	0.421
Ours	0.375	0.660	0.161	0.101	0.534	0.303	0.647	0.176	0.377	0.452

Table 2. Generalization experiment of our scheme based on BEVFormer and PETR. Our framework can apply to both top-down BEV and sparse queries schemes, represented by BEVFormer and PETR respectively. Notably, most previous methods do not support these two paradigms. N, L, W, and DT represent Nuscenes, Lyft, Waymo, and Direct Transfer, respectively.

Setting	Method	BEVFormer					PETR				
		mAP↑	mATE↓	mASE↓	mAOE↓	NDS*↑	mAP↑	mATE↓	mASE↓	mAOE↓	NDS*↑
N → L	DT	0.149	1.031	0.755	1.241	0.115	0.013	1.240	0.761	1.484	0.046
	PD-BEV	0.208	/	/	/	0.355	0.032	/	/	/	0.091
	Ours	0.237	0.850	0.191	0.407	0.377	0.332	0.880	0.170	0.210	0.456
N → W	DT	0.054	1.146	0.747	1.606	0.069	0.048	1.191	0.748	1.543	0.066
	Ours	0.249	0.872	0.239	0.365	0.379	0.200	1.141	0.242	0.505	0.309

Further, we compare the performance with current methods: CAM-Convs [7], Single-DGOD [40], DG-BEV [37], PD-BEV [26] and UDGA-BEV [2]. CAM-Convs and Single-DGOD are applied in depth prediction and 2D object detection. DG-BEV, PD-BEV and UDGA-BEV are three current SOTA schemes. We observe that, CAM-Convs and Single-DGOD have little effect because they did not consider the camera configuration problem in MC3D. DG-BEV, PD-BEV, and UDGA-BEV make great progress by considering the focal ambiguity, feature consistency and depth consistency.

In this paper, by fully and explicitly considering camera configuration, we surpass all previous methods and

achieve SOTA performance. Specifically, we surpass UDGA-BEV, the SOTA scheme, by 0.054, 0.047, 0.004 and 0.031NDS* under Nuscenes→Waymo, Nuscenes→Lyft, Waymo→Nuscenes and Lyft→Nuscenes settings.

5.2.2. Results based on BEVFormer and PETR

Tab. 2 verifies that our model-agnostic design enables application on both top-down BEV (e.g., BEVFormer-tiny) and sparse query-based frameworks (e.g., PETR-vov), which were not supported by most previous methods. Notably, PD-BEV did not mention the versions of BEVFormer and PETR used, so we choose the smallest versions for these two base models. Overall, our framework is a unified solution for all MC3D paradigms.

Table 3. Main ablation of our two modules CDA and SFM. CA refers the Camera-Aware SE module designed by BEVDepth.

Module			mAP \uparrow	mATE \downarrow	mAOE \downarrow	NDS* \uparrow
CDA	SFM	CA				
		✓	0.040	1.303	0.790	0.178
✓		✓	0.190	1.021	0.255	0.347
✓	✓	✓	0.381	0.706	0.168	0.504
	✓		0.215	0.887	0.312	0.358
✓			0.041	1.310	0.513	0.224
✓	✓		0.381	0.687	0.155	0.513

5.3. Ablation Studies

In this section, we study the effect of the two modules, SFM and CDA, in our scheme through BEVDepth. All experiments are conducted on Nuscenes \rightarrow Waymo setting because we consider that they have the largest configuration gap.

5.3.1. Main Ablations

Tab. 3 demonstrates the effectiveness of the proposed SFM and CDA. CA represents the camera-aware SE-module designed by BEVDepth, which can embed camera configuration in a simple way. The results show that our SFM can work without CDA, as it explicitly formulates the camera configuration. Applying CDA alone is ineffective due to the discrepancies of configurations that has not been processed. Combining SFM and CDA can promote NDS* from 0.178 to 0.513 and surpass CA and CDA by 0.166. This means that, on one hand, our framework can greatly improve the generalization performance across configurations, on the other hand, SFM provides a more effective way to formulate configuration. In addition, combining CA with our scheme leads to degraded performance because CA disturbs the features we modulate. Therefore, we discard CA.

5.3.2. Ablation on SFM

Tab. 4 demonstrates the effectiveness of all four spatial priors in our scheme based on CDA, including the inverse focal map (IF), ground depth map (GD), ground gradient map (GG), and Plücker raymap (PR).

First, using IF alone improves NDS* by 0.238, showing that inverse focal modulation effectively mitigates focal ambiguity by normalizing feature activation. Second, adding GD brings a 0.036 gain, as the explicit ground depth prior helps the model leverage the crucial ground-plane prior. Third, adding GG also brings an additional 0.008 gain, indicating that over-complete ground-related priors enhance spatial features. Finally, adding PR further brings a 0.007 gain, demonstrating its ability to encode pixel-level spatial priors for camera configuration.

5.3.3. Ablation on CDA

To analyze the effect of each augmentation, we apply focal augmentation (F-Aug) and our novel view synthesis aug-

Table 4. Ablation on our SFM based on CDA. IF, GD, GG, and PR represent the inverse focal map, ground depth map, ground gradient map, and Plücker raymap, respectively.

Spatial Priors				mAP \uparrow	mATE \downarrow	mAOE \downarrow	NDS* \uparrow
IF	GD	GG	PR				
				0.041	1.310	0.513	0.224
✓				0.332	0.811	0.176	0.462
✓	✓			0.365	0.696	0.182	0.498
✓	✓	✓		0.375	0.684	0.177	0.506
✓	✓	✓	✓	0.381	0.687	0.155	0.513

Table 5. Ablation on our CDA based on SFM. F-Aug and NVS-Aug represent focal augmentation and novel view synthesis augmentation, respectively.

Augmentation		mAP \uparrow	mATE \downarrow	mAOE \downarrow	NDS* \uparrow
F-Aug	NVS-Aug				
		0.215	0.887	0.312	0.358
✓		0.314	0.990	0.192	0.418
✓	✓	0.381	0.687	0.155	0.513

mentation (NVS-Aug) based on SFM in Tab. 5. F-Aug improves NDS* by 0.060, indicating that SFM benefits from augmentation. NVS-Aug boosts NDS* by 0.095, showing that the effect of focal augmentation on raw images is limited, while NVS-Aug better enhances generalization by augmenting training data with diverse configurations.

6. Conclusion

In this paper, we revisit how different camera configurations affect the generalization of multi-camera 3D object detection, and we identify that the devil lies in spatial prior discrepancies across source and target configurations. To address this, we propose CoIn3D, a generalizable MC3D framework that enables strong transferability from source configurations to unseen target ones. Extensive experiments across three landmark datasets with different camera configurations show that our framework is effective for dominant MC3D paradigms. We hope that this work can bring insight for industrial applications to facilitate deployment. In the future, we will explore improving the generalization performance of MC3D across different semantic distributions.

Acknowledgments

This work was supported by the New Generation Artificial Intelligence-National Science and Technology Major Project under Grant 2025ZD0124203, NSFC under Grants 62373298 and U24A20252, Guangdong Pearl River Talent Program under Grant 2024D03J0008, and Guangzhou Major Talent Project under Grant 2023ZT10X009.

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 multi-modal dataset for autonomous driving. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 11621–11631, 2020. 2, 6
- [2] Gysam Chang, Jiwon Lee, Donghyun Kim, Jinkyu Kim, Dongwook Lee, Daehyun Ji, Sujin Jang, and Sangpil Kim. Unified domain generalization and adaptation for multi-view 3d object detection. *Advances in Neural Information Processing Systems*, 37:58498–58524, 2024. 2, 3, 6, 7
- [3] Li Chen, Penghao Wu, Kashyap Chitta, Bernhard Jaeger, Andreas Geiger, and Hongyang Li. End-to-end autonomous driving: Challenges and frontiers. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2024. 1
- [4] Weinan Chen, Wenzheng Chi, Sehua Ji, Hanjing Ye, Jie Liu, Yunjie Jia, Jiajie Yu, and Jiyu Cheng. A survey of autonomous robots and multi-robot navigation: Perception, planning and collaboration. *Biomimetic Intelligence and Robotics*, page 100203, 2024. 1
- [5] Ziyu Chen, Jiawei Yang, Jiahui Huang, Riccardo de Lutio, Janick Martinez Esturo, Boris Ivanovic, Or Litany, Zan Gojcic, Sanja Fidler, Marco Pavone, et al. Omnire: Omni urban scene reconstruction. *arXiv preprint arXiv:2408.16760*, 2024. 3
- [6] Christy, Maggie, NikiNikatos, Phil Culliton, Vinay Shet, and Vladimir Iglovikov. Lyft 3d object detection for autonomous vehicles. <https://kaggle.com/competitions/3d-object-detection-for-autonomous-vehicles>, 2019. Kaggle. 2, 6
- [7] Jose M Facil, Benjamin Ummerhofer, Huizhong Zhou, Luis Montesano, Thomas Brox, and Javier Civera. Camconvs: Camera-aware multi-scale convolutions for single-view depth. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 11826–11835, 2019. 7
- [8] Lue Fan, Yuxue Yang, Feng Wang, Naiyan Wang, and Zhaoxiang Zhang. Super sparse 3d object detection. *IEEE transactions on pattern analysis and machine intelligence*, 45(10):12490–12505, 2023. 1
- [9] Junjie Huang and Guan Huang. Bevdet4d: Exploit temporal cues in multi-camera 3d object detection. *arXiv preprint arXiv:2203.17054*, 2022. 1, 2
- [10] Junjie Huang, Guan Huang, Zheng Zhu, Yun Ye, and Dalong Du. Bevdet: High-performance multi-camera 3d object detection in bird-eye-view. *arXiv preprint arXiv:2112.11790*, 2021. 2
- [11] Lihan Jiang, Yucheng Mao, Linning Xu, Tao Lu, Kerui Ren, Yichen Jin, Xudong Xu, Mulin Yu, Jiangmiao Pang, Feng Zhao, et al. Anysplat: Feed-forward 3d gaussian splatting from unconstrained views. *arXiv preprint arXiv:2505.23716*, 2025. 3
- [12] Xin Jin, Haisheng Su, Kai Liu, Cong Ma, Wei Wu, Fei Hui, and Junchi Yan. Unimamba: Unified spatial-channel representation learning with group-efficient mamba for lidar-based 3d object detection. In *Proceedings of the Computer Vision and Pattern Recognition Conference*, pages 1407–1417, 2025. 1
- [13] Bernhard Kerbl, Georgios Kopanas, Thomas Leimkühler, and George Drettakis. 3d gaussian splatting for real-time radiance field rendering. *ACM Trans. Graph.*, 42(4):139–1, 2023. 3
- [14] Tzofi Klinghoffer, Jonah Philion, Wenzheng Chen, Or Litany, Zan Gojcic, Jungseock Joo, Ramesh Raskar, Sanja Fidler, and Jose M Alvarez. Towards viewpoint robustness in bird’s eye view segmentation. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 8515–8524, 2023. 3
- [15] Abhinav Kumar, Yuliang Guo, Zhihao Zhang, Xinyu Huang, Liu Ren, and Xiaoming Liu. Charm3r: Towards unseen camera height robust monocular 3d detector. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 8777–8788, 2025. 2, 3
- [16] Alex H Lang, Sourabh Vora, Holger Caesar, Lubing Zhou, Jiong Yang, and Oscar Beijbom. Pointpillars: Fast encoders for object detection from point clouds. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 12697–12705, 2019. 1
- [17] Hongyang Li, Chonghao Sima, Jifeng Dai, Wenhai Wang, Lewei Lu, Huijie Wang, Jia Zeng, Zhiqi Li, Jiazhi Yang, Hanming Deng, et al. Delving into the devils of bird’s-eye-view perception: A review, evaluation and recipe. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 46(4):2151–2170, 2023. 1
- [18] Yin hao 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. In *Proceedings of the AAAI conference on artificial intelligence*, pages 1477–1485, 2023. 1, 2, 6
- [19] Yingyan Li, Lue Fan, Yang Liu, Zehao Huang, Yuntao Chen, Naiyan Wang, and Zhaoxiang Zhang. Fully sparse fusion for 3d object detection. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 46(11):7217–7231, 2024. 1
- [20] Ye Li, Wenzhao Zheng, Xiaonan Huang, and Kurt Keutzer. Unidrivr: Towards universal driving perception across camera configurations. *arXiv preprint arXiv:2410.13864*, 2024. 2, 3
- [21] 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 lidar-camera via spatiotemporal transformers. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2024. 1, 2, 6
- [22] Haisong Liu, Yao Teng, Tao Lu, Haiguang Wang, and Limin Wang. Sparsebev: High-performance sparse 3d object detection from multi-camera videos. In *Proceedings of the IEEE/CVF international conference on computer vision*, pages 18580–18590, 2023. 2
- [23] Yingfei Liu, Tiancai Wang, Xiangyu Zhang, and Jian Sun. Petr: Position embedding transformation for multi-view 3d object detection. In *European conference on computer vision*, pages 531–548. Springer, 2022. 1, 2, 6
- [24] Zhijian Liu, Haotian Tang, Alexander Amini, Xinyu Yang, Huizi Mao, Daniela Rus, and Song Han. Bevfusion: Multi-

- task multi-sensor fusion with unified bird’s-eye view representation. *arXiv preprint arXiv:2205.13542*, 2022. 1
- [25] Hao Lu, Tianshuo Xu, Wenzhao Zheng, Yunpeng Zhang, Wei Zhan, Dalong Du, Masayoshi Tomizuka, Kurt Keutzer, and Yingcong Chen. Drivingrecon: Large 4d gaussian reconstruction model for autonomous driving. *arXiv preprint arXiv:2412.09043*, 2024. 3
- [26] Hao Lu, Yunpeng Zhang, Guoqing Wang, Qing Lian, Dalong Du, and Ying-Cong Chen. Towards generalizable multi-camera 3d object detection via perspective rendering. In *Proceedings of the AAAI Conference on Artificial Intelligence*, pages 5811–5819, 2025. 2, 3, 6, 7
- [27] Jinyung Park, Chenfeng Xu, Shijia Yang, Kurt Keutzer, Kris Kitani, Masayoshi Tomizuka, and Wei Zhan. Time will tell: New outlooks and a baseline for temporal multi-view 3d object detection. *arXiv preprint arXiv:2210.02443*, 2022. 2
- [28] 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. 2
- [29] Julius Plücker. *Analytisch-geometrische Entwicklungen*. GD Baedeker, 1828. 5
- [30] Nikhila Ravi, Jeremy Reizenstein, David Novotny, Taylor Gordon, Wan-Yen Lo, Justin Johnson, and Georgia Gkioxari. Accelerating 3d deep learning with pytorch3d. *arXiv preprint arXiv:2007.08501*, 2020. 6
- [31] Thomas Roddick, Alex Kendall, and Roberto Cipolla. Orthographic feature transform for monocular 3d object detection. *arXiv preprint arXiv:1811.08188*, 2018. 2
- [32] Pei Sun, Henrik Kretzschmar, Xerxes Dotiwalla, Aurelien Chouard, Vijaysai Patnaik, Paul Tsui, James Guo, Yin Zhou, Yuning Chai, Benjamin Caine, et al. Scalability in perception for autonomous driving: Waymo open dataset. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 2446–2454, 2020. 2, 6
- [33] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. Attention is all you need. *Advances in neural information processing systems*, 30, 2017. 5
- [34] Ignacio Vizzo, Tiziano Guadagnino, Jens Behley, and Cyrill Stachniss. Vdbfusion: Flexible and efficient tsdf integration of range sensor data. *Sensors*, 22(3):1296, 2022. 6
- [35] Haotian Wang, Meng Yang, Xihu Zheng, and Gang Hua. Scale propagation network for generalizable depth completion. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2024. 6
- [36] Shihao Wang, Yingfei Liu, Tiancai Wang, Ying Li, and Xiangyu Zhang. Exploring object-centric temporal modeling for efficient multi-view 3d object detection. In *Proceedings of the IEEE/CVF international conference on computer vision*, pages 3621–3631, 2023. 2
- [37] Shuo Wang, Xinhai Zhao, Hai-Ming Xu, Zehui Chen, Dameng Yu, Jiahao Chang, Zhen Yang, and Feng Zhao. Towards domain generalization for multi-view 3d object detection in bird-eye-view. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 13333–13342, 2023. 2, 3, 6, 7
- [38] 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. 1, 2
- [39] Dongxu Wei, Zhiqi Li, and Peidong Liu. Omni-scene: Omni-gaussian representation for ego-centric sparse-view scene reconstruction. In *Proceedings of the Computer Vision and Pattern Recognition Conference*, pages 22317–22327, 2025. 3
- [40] Aming Wu and Cheng Deng. Single-domain generalized object detection in urban scene via cyclic-disentangled self-distillation. In *Proceedings of the IEEE/CVF Conference on computer vision and pattern recognition*, pages 847–856, 2022. 7
- [41] Yunzhi Yan, Haotong Lin, Chenxu Zhou, Weijie Wang, Haiyang Sun, Kun Zhan, Xianpeng Lang, Xiaowei Zhou, and Sida Peng. Street gaussians: Modeling dynamic urban scenes with gaussian splatting. In *European Conference on Computer Vision*, pages 156–173. Springer, 2024. 3
- [42] Chenyu Yang, Yuntao Chen, Hao Tian, Chenxin Tao, Xizhou Zhu, Zhaoxiang Zhang, Gao Huang, Hongyang Li, Yu Qiao, Lewei Lu, et al. Bevformer v2: Adapting modern image backbones to bird’s-eye-view recognition via perspective supervision. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 17830–17839, 2023. 1, 2
- [43] Fan Yang, Hui Chen, Yuwei He, Sicheng Zhao, Chenghao Zhang, Kai Ni, and Guiguang Ding. Geometry-guided domain generalization for monocular 3d object detection. In *Proceedings of the AAAI Conference on Artificial Intelligence*, pages 6467–6476, 2024. 2, 3
- [44] Lei Yang, Tao Tang, Jun Li, Kun Yuan, Kai Wu, Peng Chen, Li Wang, Yi Huang, Lei Li, Xinyu Zhang, et al. Bevheight++: Toward robust visual centric 3d object detection. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2025. 2
- [45] Tianwei Yin, Xingyi Zhou, and Philipp Krahenbuhl. Center-based 3d object detection and tracking. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 11784–11793, 2021. 1
- [46] Jason Y Zhang, Amy Lin, Moneish Kumar, Tzu-Hsuan Yang, Deva Ramanan, and Shubham Tulsiani. Cameras as rays: Pose estimation via ray diffusion. *arXiv preprint arXiv:2402.14817*, 2024. 5
- [47] Liangtao Zheng, Yicheng Liu, Yue Wang, and Hang Zhao. Cross-dataset sensor alignment: Making visual 3d object detector generalizable. In *Conference on Robot Learning*, pages 1903–1929. PMLR, 2023. 2, 3