

Toward Accurate Camera-based 3D Object Detection via Cascade Depth Estimation and Calibration

Chaoqun Wang^{1*}, Yiran Qin^{1*}, Zijian Kang², Ningning Ma², and Ruimao Zhang^{1†}

Abstract—Recent camera-based 3D object detection is limited by the precision of transforming from image to 3D feature spaces, as well as the accuracy of object localization within the 3D space. This paper aims to address such a fundamental problem of camera-based 3D object detection: How to effectively learn depth information for accurate feature lifting and object localization. Different from previous methods which directly predict depth distributions by using a supervised estimation model, we propose a cascade framework consisting of two depth-aware learning paradigms. First, a depth estimation (DE) scheme leverages relative depth information to realize the effective feature lifting from 2D to 3D spaces. Furthermore, a depth calibration (DC) scheme introduces depth reconstruction to further adjust the 3D object localization perturbation along the depth axis. In practice, the DE is explicitly realized by using both the absolute and relative depth optimization loss to promote the precision of depth prediction, while the capability of DC is implicitly embedded into the detection Transformer through a depth denoising mechanism in the training phase. The entire model training is accomplished through an end-to-end manner. We propose a baseline detector and evaluate the effectiveness of our proposal with +2.2%/+2.7% NDS/mAP improvements on NuScenes benchmark, and gain a comparable performance with 55.9%/45.7% NDS/mAP. Furthermore, we conduct extensive experiments to demonstrate its generality based on various detectors with about +2% NDS improvements.

I. INTRODUCTION

Object detection in 3D space is pivotal for numerous real-world applications, including robotics and autonomous driving systems. Recently, camera-based 3D detectors have attracted increasing attention both in academia and the industry, due to various benefits they offer, such as cost-effectiveness, the long perception range, and the capability to recognize vision-only attributes (e.g., traffic lights and stop signs). However, when compared to LiDAR-based methods, camera-based systems inherently suffer from a limitation tied to 2D images: the absence of depth information, which is crucial for accurate 3D detection.

The literature presents numerous methods aimed at addressing this challenge by extracting depth information. One branch involves directly predicting 3D object boxes from 2D input [1], [2], [3]. In these methods, depth information is implicitly learned by optimizing the 3D box location and sizes, supervised by the provided ground truth. For more precise depth information, another branch of methods [4], [5], [6] explicitly introduces the depth estimation module,

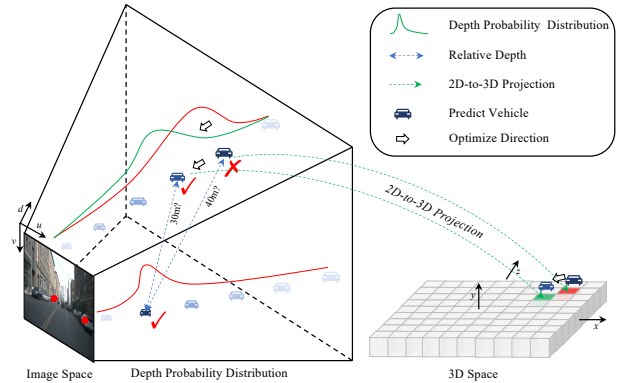


Fig. 1. The red curves indicate the predicted depth probability distribution of objects in red dots. By supervising relative depth in the DE scheme, we can optimize the distribution (from red to green curve), which is more accurate. In the DC scheme, we generate noised anchors (red cube) from the ground truth box (green cube) in 3D space. By reconstructing them, the detector obtains depth calibration capability.

which is trained by the supervision of depth maps sourced from LiDAR inputs. Although significant progress has been made, limited by the natural drawback of camera inputs, depth information prediction remains a challenge due to the following two difficulties.

- Learning the depth usually relies on simple pixel-level supervision and lacks constraints based on the structural context of the scene, leading to high prediction errors.
- Inaccurate depth information can further exacerbate disturbances in the object localization process, leading to a further decline in the accuracy of 3D detection.

To address these coherent issues, we propose a cascade framework leveraging two depth-aware learning paradigms to effectively predict depth information for accurate camera-based 3D object detection. We first introduce a novel depth estimation (DE) scheme, which defines depth relationships through the context of the scene, and uses it as a supervisory signal to train a more robust depth estimator. As illustrated on the left of Fig. 1, given two objects (denoted by red dot) in the image, when directly using depth information for supervision, the error rate of the trained depth estimator often gradually increases with the growing distance of the objects from the camera. However, by introducing the relative depth between objects at different distances as constraints during training, the aforementioned issue can be effectively mitigated (e.g., the depth probability distribution is corrected from the red curve to the green one).

On the other hand, we further introduce a depth calibration (DC) scheme to prevent the decline in object localization accuracy caused by depth estimation errors. During the

*Work done during an internship at NIO.

¹ School of Data Science, Shenzhen Research Institute of Big Data, The Chinese University of Hong Kong, Shenzhen (CUHK-Shenzhen), China. {chaoqunwang@link., yiranqin@link., ruimaozhang@}cuhk.edu.cn

² NIO. {zijian.kang, ningning.ma}@nio.com

†Corresponding author.

training process, we introduce a depth denoising mechanism, embedding this capability within the Transformer model. As shown on the right of Fig. 1, inaccurate depth estimation in images can lead to objects shifting along the depth axis in 3D space. Our proposed DC scheme simulates this object perturbation, generating noised anchors from the ground truth box. By reconstructing the actual boxes from these noised ones, the detector inherently acquires a depth calibration capability.

In practice, our framework integrates the aforementioned auxiliary supervisions, and training in an end-to-end manner to enhance detection performance without added inference costs. In addition, DE and DC schemes can be used independently. Integrating either one into existing camera-based detectors consistently results in a stable performance improvement. The contributions can be summarized as follows.

- We present a cascade framework for camera-based 3D object detection, which gains depth prediction accuracy at both the feature lifting and object localization phases. It provides an effective way to mitigate the accuracy decline caused by cumulative depth prediction errors.
- We introduce depth estimation (DE) and depth calibration (DC) schemes into our framework, which leverage pre-defined depth-aware content within scenes for model training, offering a cost-effective approach that significantly enhances the robustness of the detectors.
- We have implemented a detector and achieved substantial improvements, demonstrating the efficacy of our proposed framework. Moreover, we have carried out extensive experiments using a variety of open-sourced, state-of-the-art detectors to showcase the broad applicability of both DE and DC schemes.

II. RELATED WORKS

Camera-based 3D Object Detection requires predicting 3D boxes only from camera data, in which depth estimation is the most challenging since it is an ill-posed problem. One branch of previous works lies in explicitly using depth information [4], [5], [6], [7], [8], [9], [10]. Some work [4], [5], [6], [7] lift the features from Field-Of-View (FOV) to Bird-Eye-View (BEV) plane, via the prediction depth and camera intrinsics/extrinsic matrix, or mimic the pseudo-LiDAR by projecting the depth map into 3D space, and feed the BEV/pseudo-LiDAR features into detection head to predict 3D box. To gain better performance, BEVDepth [4] utilizes the corresponding LiDAR data to generate depth ground truth and supervise the accurate depth estimation. Besides this, some early works [8], [9], [10] predicted 2D box and depth information from camera inputs and lifted the results into 3D spaces via the geometric properties and constraints to obtain 3D prediction, gaining low generality limited by relying on feature engineering.

Another branch implicitly utilizes depth information. The core of these methods [1], [2], [11] lies in querying objects from image features, through the camera intrinsics/extrinsic matrix. The object queries in 3D space could project to the reference regions from camera inputs and aggregate the object presentation via transformer decoder. Apart from the

previous methods, BEVFormer [3], [12] project 2D features into the BEV plane via the BEV queries, which is aggregated from camera features similar to DETR3D [1].

Depth Estimation is a classical computer vision task. The previous general depth estimation methods mainly consist of single- and multi-view methods. The former [13], [14], [15], [16], [17] regard the depth estimation task as a regression or classification problem of depth distribution. These works adopt an encoder-decoder framework which is supervised by the given dense depth map. The latter [18], [19], [20] usually construct a cost volume to regress disparities based on photometric consistency for effective depth prediction. For camera-based 3D detection, previous works [4], [6], [7] introduce an extra depth estimation module to predict depth for accurate 3D object localization. CaDDN [7] learns a categorical depth distribution and then projects image features into BEV space. BEVDepth [4] supervises depth learning with sparse LiDAR data and proposes a depth correction module to reduce depth error.

Denoising Mechanism aims to integrate the recognition calibration to the transformer head [21], [22], [23] by reconstructing the noised anchors, which is first proposed [21] to accelerate the convergence of DETR-based detectors by introducing noised positional anchors and reconstructing them to the original box. Furthermore, DINO [22] introduces a contrastive denoising training strategy and gains promising improvements on the 2D object detection benchmark. In this paper, we introduce depth calibration, which adopts depth denoising to adjust the 3D object localization perturbation along depth axis for accurate 3D detection.

III. METHOD

A. The Overall Framework

As shown in Fig. 2, we build up a baseline framework consisting of both the DE and DC schemes. It can be briefly divided into three parts. The first is a shared encoder to extract the feature maps of each image. Given multi-camera inputs $\mathbf{X} = \{\mathbf{X}_1, \mathbf{X}_2, \dots, \mathbf{X}_N\}$, where N indicates the number of camera views, we simply adopt ResNet followed FPN as the backbone to extract features $\mathbf{F} = \{\mathbf{F}_n \in \mathbb{R}^{C \times H \times W}\}_{n=1}^N$, where H, W, C denote the height, the width and the channel dimension.

The second part aims to transfer the multi-camera features from 2D to 3D space, *i.e.* from Front-of-View (FOV) to Bird’s-Eye-View (BEV). Following BEVDepth [4], we first employ DepthNet to predict the depth distribution $\mathbf{F}^d = \{\mathbf{F}_n^d \in \mathbb{R}^{C_d \times H \times W}\}_{n=1}^N$ and context features $\mathbf{F}^c = \{\mathbf{F}_n^c \in \mathbb{R}^{C_c \times H \times W}\}_{n=1}^N$ from multi-view image features. Here $C_d = 118$ indicates the probabilities of the pixel-wise depth from 1 to 118 meters and $C_c = 80$ is the contextual channel. Then we lift the multi-view 2D context features into 3D spaces via the prediction depth and camera intrinsics, and obtain one BEV features $\mathbf{F}^b \in \mathbb{R}^{C_b \times H \times W}$ the same as BEVDepth. More details could be referred in BEVDepth [4]. To obtain the accurate depth, we supervise the depth estimation via the proposed DE scheme introduced in Sec. III-B.

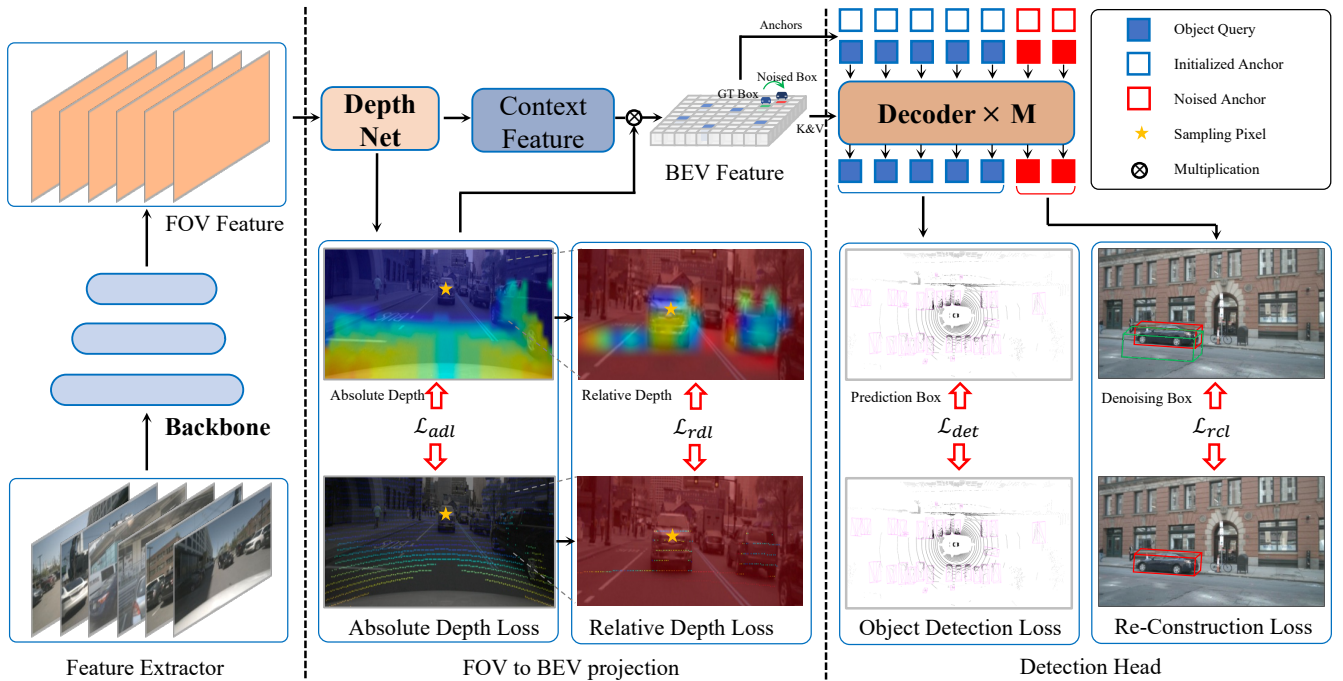


Fig. 2. The overall architecture of our proposed detector consists of three parts: feature extractor, FOV to BEV translation, and detection head. For a given multi-camera input, we extract 2D features via a shared encoder and project them to 3D spaces via predicted depth, and the generated BEV features are fed into the detection head for object localization and recognition. The L_{adl} and L_{det} indicate the absolute depth loss and detection loss while L_{rdl} and L_{rcl} are relative depth loss and reconstruction loss in our proposed DE and DC scheme. Best view in color and more details are referred in Sec. III.

Finally, we adopt a simple transformer decoder-based head to localize the 3D objects from the BEV feature. Specifically, we randomly initialize object query content $\mathbf{Q}_c \in \mathbb{R}^{m \times C}$ and reference learnable anchors $\mathbf{Q}_r \in \mathbb{R}^{m \times 6}$, to aggregate the object presentations from BEV features in the specific reference region via a simple sequential deformable decoder. Here $m = 900$ indicates the query number and $C = 256$ denotes the query dimension. Each reference anchor is represented as 6 dimensions, which denotes the anchor location (x, y, z) and scale (w, l, h) . The updated object queries are further fed into a detection head to predict the box location and a classification head to predict the class label.

Even if the DE scheme effectively improves the depth prediction accuracy, prediction errors are still inevitable, which results in object localization disturbances along the depth axis. To address this, we further introduced the DC scheme to allow the model to adaptively mitigate such errors in Sec. III-C. Specifically, we simulate the depth prediction error and generate the extra object query with noised reference anchors $\mathbf{R}'_r \in \mathbb{R}^{m' \times 6}$, as well as the corresponding query content $\mathbf{Q}'_c \in \mathbb{R}^{m' \times C}$ associated with the object class label. By reconstructing the original object boxes during the training phase, we can implicitly integrate depth calibration capabilities into the detection model the same as [21].

B. Depth Estimation Scheme.

In addition to directly using absolute depth as supervisory information, we also introduce relative depth to describe the depth relationship between two positions in the scene and use it as auxiliary supervision for the depth estimation module. For simplicity, we omit the view index n and process all image views equally. Given a prediction depth probability

distribution $\mathbf{F}^d \in \mathbb{R}^{C_d \times H \times W}$ of a specific image, we first adopt the corresponding LiDAR data and project to image space to obtain the sparse depth map $\mathbf{G}^d \in \mathbb{R}^{C_d \times H \times W}$ as well as a mask $\mathbf{M} \in \mathbb{R}^{H \times W}$. Here, $C_d = 118$ and $\mathbf{F}_i^d[c] \in (0, 1)$ indicates the probability of the i -pixel belong to the c meters depth, where $0 \leq i < H \times W$ and $1 \leq c \leq C_d$. \mathbf{G}_i^d is a one-hot vector that records the real depth value of the i -pixel. \mathbf{M} denotes a mask and $\mathbf{M}_i = 0$ means no LiDAR point projects to corresponding position, otherwise $\mathbf{M}_i = 1$.

Following BEVDet [5], we first introduce the absolute depth loss and consider the depth estimation as a pixel-wise classification task, supervised by the cross-entropy loss.

$$\mathcal{L}_{adl} = -\frac{1}{|\mathbf{M}|} \sum_{i=0}^{H \times W} \sum_{c=1}^{C_d} \mathbf{M}_i \cdot \mathbf{G}_i^d[c] \cdot \log(\mathbf{F}_i^d[c]), \quad (1)$$

where $|\mathbf{M}|$ is the summation of \mathbf{M} .

Furthermore, we introduce relative depth as the auxiliary supervision to optimize the depth estimator, maintaining the structural context of the scene. To achieve this, we calculate the depth value from the \mathbf{F}^d and \mathbf{G}^d :

$$\begin{aligned} \tilde{\mathbf{F}}_i^d &= \sum_{c=1}^{C_d} c \cdot \mathbf{F}_i^d[c], \\ \tilde{\mathbf{G}}_i^d &= \text{argmax}(\mathbf{G}_i^d), \end{aligned} \quad (2)$$

Here, we calculate the depth value in an integral manner, and the $\tilde{\mathbf{F}}_i^d$ and $\tilde{\mathbf{G}}_i^d$ denotes the predicted and the ground truth depth values of i -th pixels. Then the relative depth can be formulated as,

$$\mathbf{R}_{jk} = \tilde{\mathbf{F}}_j^d - \tilde{\mathbf{F}}_k^d, \quad (3)$$

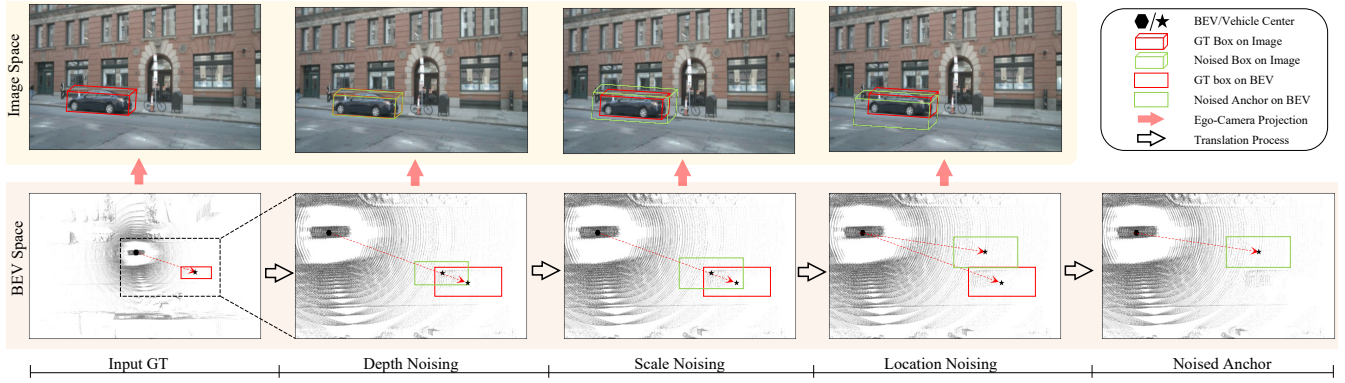


Fig. 3. Noised anchors generation process. Given a ground truth box in the ego coordinate system, we add depth, scale, and location noise to generate the noised reference anchors, which indicate $f_d, f_s,$ and f_l in Eqn. 6 respectively. We visualize the corresponding noised boxes in BEV and image space.

where $\mathbf{R} \in \mathbb{R}^{HW \times HW}$ and \mathbf{R}_{jk} indicates the distance within j -th and k -th pixels along the depth axis. Considering that a specific pixel might have significant relative depth differences with multiple pixels, which may result in numerical instability during the optimization, we introduce a normalization strategy to keep the values of relative depth within a certain range,

$$\mathbf{R}_{jk} = \frac{e^{-|\mathbf{R}_{jk}|/\tau}}{\sum_{k=0}^{HW} e^{-|\mathbf{R}_{jk}|/\tau}}, \quad (4)$$

where τ is the temperature factor. According to the Eqn. 4, the pixel pair within a shorter distance has a higher response. We display the absolute depth as well as relative depth (between the specific pixel in orange star and other all pixels) in Fig. 2. After achieving the normalization values from the depth ground truth, we can guide the learning of such relative depth using a Kullback–Leibler divergence loss.

$$\mathcal{L}_{rdl} = \text{Avg}(\mathbf{R}^g \odot \log(\mathbf{R}^g / \mathbf{R}^p)), \quad (5)$$

where \mathbf{R}^p and \mathbf{R}^g indicate the relative depth from the prediction and ground truth, \odot is element-wise multiplication and $\text{Avg}(\cdot)$ indicates the average value of a matrix.

In practice, considering the memory efficiency and the depth relationships occur in a limited range, we only calculate the relative depth in local patches with the size $p \times p$, which are selected from the whole depth prediction map by using a sliding window, such that the relative depth scale is $p^2 \times p^2$ compared with $HW \times HW$, and the relative depth loss is the average value of Eqn. 5 for all local patches.

C. Depth Calibration scheme.

The proposed DC scheme aims to simulate the depth prediction errors and generate noised anchors, by reconstructing the extra noised queries to the real object location, the depth calibration capability could be implicitly embedded into the detection transformer. In this section, we introduce the noised reference anchor generation process as shown in Fig. 3 for the DC scheme. Specifically, the process mainly disturbs the ground box in three aspects to generate noised reference anchors, including depth, scale, and location. Given a ground truth box in ego coordinate system $\mathcal{B} = (x, y, z, w, l, h)$, the

noised box in Fig. 3 could be formulated as:

$$\mathcal{B}' = f_l \circledast f_s \circledast f_d(\mathcal{B}; \sigma_d, \sigma_s, \sigma_l), \quad (6)$$

where \circledast means assembling translation function f_d, f_s, f_l , which adds depth, scale, and location noise by factor $\sigma_d, \sigma_s, \sigma_l$. $\mathcal{B}' = (x', y', z', w', l', h')$ is the noised reference anchor. We describe the translation function as follows.

First, since the depth prediction occurs in 2D images, the newly added depth noise will simultaneously affect the anchor's location and scale in 3D space, which can be formulated as:

$$f_d(\mathcal{B}, \sigma_d) = \sigma_d \cdot [x, y, z, w, l, h]^T, \quad (7)$$

To prove this, we first consider projecting any points (x, y, z) to the 2D space with the depth information (u, v, d) via the translation parameter \mathbf{P} :

$$d [u, v, 1]^T = \mathbf{P} [x, y, z]^T, \quad (8)$$

where (u, v) is the pixel position the 2D space and d denotes the corresponding depth value of pixel (u, v) . Adding depth noise with factor σ_d will change the depth $d' = \sigma_d \cdot d$, and the corresponding 3D points with depth noise is:

$$[x', y', z']^T = \mathbf{P}^{-1} d' [u, v, 1]^T = \sigma_d \cdot [x, y, z]^T, \quad (9)$$

Based on the above translation, we conclude that adding depth noising will lead to constant changes in the point location. So, the box scale, which is the distances of two corner points will constantly change with the same factor, which can be presented as follows:

$$[w', l', h']^T = \sigma_d \cdot [w, l, h]^T, \quad (10)$$

Next, we add scale noise and location noise, which simply modified the object scale and location as follows,

$$\begin{aligned} f_s(\mathcal{B}, \sigma_s) &= [x, y, z, \sigma_s w, \sigma_s l, \sigma_s h]^T, \\ f_l(\mathcal{B}, \sigma_l) &= [\sigma_l x, \sigma_l y, \sigma_l z, w, l, h]^T, \end{aligned} \quad (11)$$

To increase the diversity of noised anchors, the noised factors are random values in a specific range controlled by hyper-parameters, e.g., $\sigma_d = \text{Rand}(1 - \delta_d, 1 + \delta_d)$. Here we set $\delta_d = 0.5, \delta_s, \delta_l = 0.1$

Finally, as described in Sec. III-A, we produce noised reference anchors \mathbf{Q}'_r via the above formulation and feed extra noised queries to the shared detection head. We constrain the reconstruction loss with the following formulation:

$$\mathcal{L}_{rcl} = \mathcal{L}_{det}(\mathcal{H}(\mathbf{Q}'_c, \mathbf{Q}'_r)), \quad (12)$$

where \mathcal{H} is the decoder head and \mathcal{L}_{det} is the detection loss that includes the classification cross-entropy loss and regression L1 loss.

D. Supervision.

Following the aforementioned architecture in Sec. III-A, as well as the proposed depth estimation scheme in Sec. III-B and depth calibration in Sec. III-C, the supervision in our proposed including the two basic supervision such as detection loss and absolute depth loss, as well as two auxiliary supervision including relative depth loss and reconstruction loss:

$$\mathcal{L} = \mathcal{L}_{adl} + \mathcal{L}_{det}(\mathcal{H}(\mathbf{Q}_c, \mathbf{Q}_r)) + \alpha \mathcal{L}_{rdl} + \beta \mathcal{L}_{rcl} \quad (13)$$

where $\alpha = 0.1$ and $\beta = 1.0$ are manually set balance factors.

IV. EXPERIMENTS

A. Experiment settings.

Dataset and Metric. We conduct experiments based on NuScenes [24] benchmarks following previous works [2], [3], [5], [29]. NuScenes is a large-scale autonomous driving dataset, which contains 700/150/150 sequences for training, testing, and validation. We follow previous works to take the official evaluation metrics, including mean average precision (mAP), ATE, ASE, AOE, AVE, and AAE for our evaluations.

Implementation Details. The detector is trained with 20 epochs using AdamW optimizer with a base learning rate of 2×10^{-4} on 8 NVIDIA Tesla A100 GPUs. To demonstrate the effectiveness, we apply the DE to multiple open-sourced methods that leverage depth prediction in their frameworks, such as BEVDepth [4] and BEVStereo [6]. We also realize DC in various transformer-based detectors, such as DETR3D [1], BEVFormer [3], and PETR [2]. All the models in ablation studies adopt ResNet50 [25] as the backbone for higher efficiency. To compare with previous state-of-the-art detectors, we adopt Swin-Base [26] as our backbone and improve the performance via CBGS [27], Exponential Moving Average (EMA) [28], and sequential long-term frames following BEVStereo [6].

B. Ablation Study for Depth Estimation Scheme.

To evaluate the depth estimation scheme, we conduct experiments with several settings as shown in Tab. I. Here, the RD indicates using the relative depth in Eqn. 3 to define the relationships and supervised by L1 loss in contrastive to KL loss, the SW indicates select relative regions in a slide window manner with window side p , in contrast to global manner, RD^\dagger and KL means using depth relations in Eqn. 4 with temperature factor τ and supervised by the Kullback–Leibler divergence loss. With our proposed DE scheme, the best setting yields results with 51.0%/40.9%

TABLE I

INVESTIGATION OF THE DE SCHEME WITH DIFFERENT SETTINGS. RD, SW, KL INDICATE RELATIVE DEPTH, SLIDING WINDOW, AND KL LOSS. p AND τ INDICATE PATCH SIZE AND TEMPERATURE FACTOR.

Method	NDS \uparrow	mAP \uparrow
Baseline	0.495	0.392
Baseline+RD	0.501	0.400
Baseline+RD+SW($p = 5$)	0.503	0.401
Baseline+RD+SW($p = 7$)	0.500	0.398
Baseline+RD † +SW($p = 5$)+KL($\tau = 4$)	0.505	0.403
Baseline+RD † +SW($p = 5$)+KL($\tau = 8$)	0.510	0.409
Baseline+RD † +SW($p = 5$)+KL($\tau = 16$)	0.508	0.406

TABLE II

NDS AND MAP ON NUSCENE VAL SET. α INDICATES BALANCE FACTOR.

Method	NDS \uparrow	mAP \uparrow
Baseline	0.495	0.392
Baseline+DE($\alpha = 0.01$)	0.502	0.403
Baseline+DE($\alpha = 0.1$)	0.510	0.409
Baseline+DE($\alpha = 0.5$)	0.506	0.404

TABLE III

GENERALIZATION ABILITY OF DE SCHEME ON MULTIPLE DETECTORS.

Method	NDS \uparrow	mAP \uparrow	mATE \downarrow	mASE \downarrow	mAOE \downarrow	mAVE \downarrow	mAAE \downarrow
BEVDepth	0.484	0.363	0.612	0.278	0.482	0.402	0.202
BEVDepth+DE	0.493	0.370	0.600	0.280	0.477	0.371	0.189
BEVStereo	0.497	0.381	0.584	0.281	0.473	0.385	0.207
BEVStereo+DE	0.507	0.393	0.576	0.275	0.476	0.374	0.190
Baseline	0.495	0.392	0.627	0.327	0.460	0.396	0.204
Baseline+DE	0.510	0.409	0.615	0.298	0.471	0.374	0.193

in terms of NDS and mAP, exceeding the baseline model 49.5%/39.2%, and gains +1.5%/ + 1.7% improvements. Besides, we conduct experiments of the hyper-parameters for balance factor α in Tab. II. As an auxiliary supervision, the $\alpha = 0.1$ could gain the best improvements.

To demonstrate the broad applicability of the proposed DE, we apply the DE scheme to multiple open-sourced methods that adopt the depth estimator. As reported in Tab. III, DE brings consistent improvements and demonstrates its capability across different models. Without bells and whistles, the DE scheme improves BEVDepth and BEVStereo for +0.9%/ + 0.7% and +1.0%/ + 1.2% NDS/mAP improvements without introducing any extra inference cost.

C. Ablation Study for Depth Calibration Scheme.

As shown in Tab. V, we conduct experiments with several settings and hyper-parameters to analyze the depth calibration scheme. Here, the DN, LN, and SN indicate the depth noising, location noising, and scale noising respectively. From the results, we can observe that adding the depth noising could gain significant improvements by +0.9%/ + 1.1% in terms of NDS and mAP. Furthermore, denoising the location and scale errors could continuously boost the performances and yield +1.6%/ + 2.0% improvements.

To demonstrate the broad applicability of the proposed DC, we also apply DC to multiple open-sourced methods that employ transformer-based detection heads. As illustrated in Tab. VI, the DC scheme consistently enhances various methods. Specifically, the DE scheme boosts DETR3D, BEVFormer, and PETR by +1.8%/+2.1%, +2.6%/+2.4%, and +1.7%/ + 1.9% NDS/mAP.

Method	Publication	Backbone	LiDAR	NDS \uparrow	mAP \uparrow	mATE \downarrow	mASE \downarrow	mAOE \downarrow	mAVE \downarrow	mAAE \downarrow
BEVDet [5]	Arxiv	R50	\times	0.379	0.298	0.725	0.279	0.589	0.860	0.245
DETR3D [1]	CoRL2021	R50	\times	0.373	0.302	0.811	0.282	0.493	0.979	0.212
PETR [2]	ECCV2022	R50	\times	0.403	0.339	0.748	0.273	0.539	0.907	0.203
BEVFormer [3]	ECCV2022	R50	\times	0.354	0.252	0.900	0.294	0.655	0.657	0.216
PETrv2 [30]	ICCV2023	R50	\times	0.494	0.398	0.690	0.273	0.467	0.424	0.195
BEVDet4D [31]	Arxiv	Swin-T	\checkmark	0.476	0.338	0.672	0.274	0.460	0.337	0.185
BEVDepth [4]	AAAI2023	R50	\checkmark	0.478	0.359	0.612	0.269	0.507	0.409	0.201
BEVStereo [6]	AAAI2023	R50	\checkmark	0.500	0.372	0.598	0.270	0.438	0.367	0.190
Ours	ICRA2024	R50	\checkmark	0.517	0.419	0.611	0.276	0.465	0.368	0.211
DETR3D ‡ [1]	CoRL2021	R101	\times	0.434	0.349	0.716	0.268	0.379	0.842	0.200
PETR ‡ [2]	ECCV2022	R101	\times	0.442	0.370	0.711	0.267	0.383	0.865	0.201
BEVFormer [3]	ECCV2022	R101-DCN	\times	0.517	0.416	0.672	0.274	0.369	0.397	0.191
PETrv2 ‡ [30]	ICCV2023	R101	\times	0.524	0.421	0.681	0.267	0.357	0.377	0.186
PolarFormer [32]	AAAI2023	R101-DCN	\times	0.528	0.432	0.648	0.270	0.348	0.409	0.201
BEVDet4D [31]	Arxiv	Swin-B	\checkmark	0.545	0.421	0.579	0.258	0.329	0.301	0.191
FCOS3D [33]	ICCV2021	R101	\checkmark	0.415	0.343	0.725	0.263	0.422	1.292	0.153
BEVDepth [4]	AAAI2023	R101	\checkmark	0.535	0.412	0.565	0.266	0.358	0.331	0.190
Ours	ICRA2024	Swin-B	\checkmark	0.559	0.457	0.543	0.259	0.354	0.349	0.193

TABLE IV

PERFORMANCES COMPARISON ON NUSCENE VAL SET. ‡ INDICATES USING PRE-TRAINED MODEL FROM FCOS3D. LiDAR INDICATES SUPERVISING THE DEPTH ESTIMATION BY GROUND TRUTH GENERATED FROM LiDAR DATA. THE BEST AND SECOND ARE IN **BOLD** AND UNDERLINED.

TABLE V

INVESTIGATION OF THE DC SCHEME WITH DIFFERENT SETTINGS. DN, LN, AND SN INDICATE DEPTH, LOCATION, AND SCALE NOISE.

Method	NDS \uparrow	mAP \uparrow
Baseline	0.495	0.392
Baseline+DN($\delta_d = 0.1$)	0.499	0.400
Baseline+DN($\delta_d = 0.5$)	0.504	0.403
Baseline+DN($\delta_d = 0.5$)+LN($\delta_l = 0.1$)	0.506	0.406
Baseline+DN($\delta_d = 0.5$)+LN($\delta_l = 0.5$)	0.506	0.402
Baseline+DN($\delta_d = 0.5$)+LN($\delta_l = 0.1$)+SN($\delta_s = 0.1$)	0.511	0.412
Baseline+DN($\delta_d = 0.5$)+LN($\delta_l = 0.1$)+SN($\delta_s = 0.5$)	0.508	0.407

TABLE VI

GENERALIZATION ABILITY OF DC SCHEME ON MULTIPLE DETECTORS.

Method	NDS \uparrow	mAP \uparrow	mATE \downarrow	mASE \downarrow	mAOE \downarrow	mAVE \downarrow	mAAE \downarrow
DETR3D	0.373	0.302	0.811	0.282	0.493	0.979	0.212
DETR3D+DC	0.391	0.323	0.799	0.259	0.495	0.952	0.203
BEVFormer	0.354	0.252	0.900	0.294	0.655	0.657	0.216
BEVFormer+DC	0.380	0.276	0.880	0.295	0.588	0.596	0.221
PETR	0.403	0.339	0.748	0.273	0.539	0.907	0.203
PETR+DC	0.420	0.358	0.725	0.269	0.514	0.887	0.192
Baseline	0.495	0.392	0.627	0.327	0.460	0.396	0.204
Baseline+DC	0.511	0.412	0.601	0.318	0.443	0.371	0.215

D. Ablation Study for Module Pruning.

In Tab. I and Tab. V, we analyze the impacts of hyper-parameters and settings in the proposed DE and DC schemes, as well as the performance improvements based on ResNet50 backbone over the baseline model. Furthermore, we conduct experiments to implement them in a cascade manner as shown in Tab. IV. We can observe that the proposed detector gains 51.7%/41.9% NDS/mAP, and achieves +2.2%/ +2.7% improvements with ResNet50, which is higher than the module pruning version, demonstrating that the DE and DC could boost the performances from different aspects and jointly promote the detection results to a certain extent.

E. Comparison with the State-of-The-Art Methods.

We aim to prove the efficacy of our proposed approach by making a comparison with previous state-of-the-art methods as shown as Tab. IV. For a fair comparison, all the detectors adopt R50/Swin-T and R101/Swin-B backbones. From the results, we can observe that our proposal could gain the

best NDS/mAP performance. Specifically, our proposal could gain 51.7%/41.9% NDS/mAP with ResNet50 backbone, and achieve +1.7%/ +4.7% improvements against BEVStereo. The significant mAP improvements demonstrate the superiority of object localization ability. When enlarging the backbone to Swin-Base, our proposal gains 55.9%/45.7% NDS/mAP, and gains a significant margin over the previous SOTA methods, such as +3.1%/ +2.5% over PolarFormer and +1.4%/ +3.6% over BEVDet4D. The exciting performance and significant promotion demonstrate the effectiveness of our proposal.

V. CONCLUSIONS

In this paper, we introduce a cascade framework to improve the accuracy of camera-based 3D object detection by promoting depth prediction in both feature lifting and object localization. Our method comprises two depth-aware learning paradigms named depth estimation and depth calibration. The depth estimation (DE) employs depth relationships as the constraints in the training phase to ensure precise depth estimation. Conversely, the depth calibration (DC) scheme integrates a depth denoising mechanism, equipping the detection transformer with depth calibration capabilities. Both the DE and DC can be seamlessly integrated into the proposed detectors and previous cutting-edge detectors to significantly elevate detection performance without extra inference costs.

ACKNOWLEDGMENT

The work is partially supported by the Young Scientists Fund of the National Natural Science Foundation of China under grant No.62106154, by the Natural Science Foundation of Guangdong Province, China (General Program) under grant No.2022A1515011524, and by Shenzhen Science and Technology Program JCYJ20220818103001002 and ZDSYS20211021111415025, and by the Guangdong Provincial Key Laboratory of Big Data Computing, The Chinese University of Hong Kong (Shenzhen).

REFERENCES

- [1] Wang, Yue, et al. "Detr3d: 3d object detection from multi-view images via 3d-to-2d queries." *Conference on Robot Learning*. PMLR, 2022.
- [2] Liu, Yingfei, et al. "Petr: Position embedding transformation for multi-view 3d object detection." *European Conference on Computer Vision*. Cham: Springer Nature Switzerland, 2022.
- [3] Li, Zhiqi, et al. "Bevformer: Learning bird's-eye-view representation from multi-camera images via spatiotemporal transformers." *European conference on computer vision*. Cham: Springer Nature Switzerland, 2022.
- [4] Li, Yinhao, et al. "Bevdepth: Acquisition of reliable depth for multi-view 3d object detection." *Proceedings of the AAAI Conference on Artificial Intelligence*. Vol. 37. No. 2. 2023.
- [5] Huang, Junjie, et al. "Bevdet: High-performance multi-camera 3d object detection in bird-eye-view." *arXiv preprint arXiv:2112.11790* (2021).
- [6] Li, Yinhao, et al. "Bevstereo: Enhancing depth estimation in multi-view 3d object detection with dynamic temporal stereo." *arXiv preprint arXiv:2209.10248* (2022).
- [7] Reading, Cody, et al. "Categorical depth distribution network for monocular 3d object detection." *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. 2021
- [8] Shi, Xuepeng, et al. "Geometry-based distance decomposition for monocular 3d object detection." *Proceedings of the IEEE/CVF International Conference on Computer Vision*. 2021.
- [9] Li, Buyu, et al. "Gs3d: An efficient 3d object detection framework for autonomous driving." *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. 2019.
- [10] Li, Peiliang, Xiaozhi Chen, and Shaojie Shen. "Stereo r-cnn based 3d object detection for autonomous driving." *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. 2019.
- [11] Wang, Shihao, et al. "Exploring Object-Centric Temporal Modeling for Efficient Multi-View 3D Object Detection." *arXiv preprint arXiv:2303.11926* (2023).
- [12] Yang, Chenyu, et al. "BEVFormer v2: Adapting Modern Image Backbones to Bird's-Eye-View Recognition via Perspective Supervision." *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. 2023.
- [13] Bhat, Shariq Farooq, Ibraheem Alhashim, and Peter Wonka. "Adabins: Depth estimation using adaptive bins." *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. 2021.
- [14] Zhou, Ling, et al. "Pattern-structure diffusion for multi-task learning." *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. 2020.
- [15] Eigen, David, Christian Puhrsch, and Rob Fergus. "Depth map prediction from a single image using a multi-scale deep network." *Advances in neural information processing systems* 27 (2014).
- [16] Fu, Huan, et al. "Deep ordinal regression network for monocular depth estimation." *Proceedings of the IEEE conference on computer vision and pattern recognition*. 2018
- [17] Ranftl, René, Alexey Bochkovskiy, and Vladlen Koltun. "Vision transformers for dense prediction." *Proceedings of the IEEE/CVF international conference on computer vision*. 2021.
- [18] Zhu, Qingtian, et al. "Deep learning for multi-view stereo via plane sweep: A survey." *arXiv preprint arXiv:2106.15328* (2021).
- [19] Wei, Zizhuang, et al. "Aa-rmvnet: Adaptive aggregation recurrent multi-view stereo network." *Proceedings of the IEEE/CVF International Conference on Computer Vision*. 2021.
- [20] Wei, Zizhuang, et al. "Bidirectional hybrid LSTM based recurrent neural network for multi-view stereo." *IEEE Transactions on Visualization and Computer Graphics* (2022).
- [21] Li, Feng, et al. "Dn-detr: Accelerate detr training by introducing query denoising." *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. 2022.
- [22] Zhang, Hao, et al. "Dino: Detr with improved denoising anchor boxes for end-to-end object detection." *arXiv preprint arXiv:2203.03605* (2022).
- [23] Yang, Jie, et al. "Neural interactive keypoint detection." *Proceedings of the IEEE/CVF International Conference on Computer Vision*. 2023.
- [24] Caesar, Holger, et al. "nuscenes: A multimodal dataset for autonomous driving." *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*. 2020.
- [25] He, Kaiming, et al. "Deep residual learning for image recognition." *Proceedings of the IEEE conference on computer vision and pattern recognition*. 2016.
- [26] Liu, Ze, et al. "Swin transformer: Hierarchical vision transformer using shifted windows." *Proceedings of the IEEE/CVF international conference on computer vision*. 2021.
- [27] Zhu, Benjin, et al. "Class-balanced grouping and sampling for point cloud 3d object detection." *arXiv preprint arXiv:1908.09492* (2019).
- [28] Tarvainen, Antti, and Harri Valpola. "Mean teachers are better role models: Weight-averaged consistency targets improve semi-supervised deep learning results." *Advances in neural information processing systems* 30 (2017).
- [29] Li, Hongyang, et al. "DFA3D: 3D Deformable Attention For 2D-to-3D Feature Lifting." *arXiv preprint arXiv:2307.12972* (2023).
- [30] Liu, Yingfei, et al. "PetrV2: A unified framework for 3d perception from multi-camera images." *arXiv preprint arXiv:2206.01256* (2022).
- [31] Huang, Junjie, and Guan Huang. "Bevdet4d: Exploit temporal cues in multi-camera 3d object detection." *arXiv preprint arXiv:2203.17054* (2022).
- [32] Jiang, Yanqin, et al. "Polarformer: Multi-camera 3d object detection with polar transformer." *Proceedings of the AAAI Conference on Artificial Intelligence*. Vol. 37. No. 1. 2023.
- [33] Wang, Tai, et al. "Fcos3d: Fully convolutional one-stage monocular 3d object detection." *Proceedings of the IEEE/CVF International Conference on Computer Vision*. 2021.