

PromptDet: A Lightweight 3D Object Detection Framework with LiDAR Prompts

Kun Guo, Qiang Ling*

Dept. of Automation, University of Science and Technology of China, Hefei 230027, China
guok@mail.ustc.edu.cn, qling@ustc.edu.cn

Abstract

Multi-camera 3D object detection aims to detect and localize objects in 3D space using multiple cameras, which has attracted more attention due to its cost-effectiveness trade-off. However, these methods often struggle with the lack of accurate depth estimation caused by the natural weakness of the camera in ranging. Recently, multi-modal fusion and knowledge distillation methods for 3D object detection have been proposed to solve this problem, which are time-consuming during the training phase and not friendly to memory cost. In light of this, we propose PromptDet, a lightweight yet effective 3D object detection framework motivated by the success of prompt learning in 2D foundation model. Our proposed framework, PromptDet, comprises two integral components: a general camera-based detection module, exemplified by models like BEVDet and BEVDepth, and a LiDAR-assisted prompter. The LiDAR-assisted prompter leverages the LiDAR points as a complementary signal, enriched with a minimal set of additional trainable parameters. Notably, our framework is flexible due to our prompt-like design, which can not only be used as a lightweight multi-modal fusion method but also as a camera-only method for 3D object detection during the inference phase. Extensive experiments on nuScenes validate the effectiveness of the proposed PromptDet. As a multi-modal detector, PromptDet improves the mAP and NDS by at most 22.8% and 21.1% with fewer than 2% extra parameters compared with the camera-only baseline. Without LiDAR points, PromptDet still achieves an improvement of at most 2.4% mAP and 4.0% NDS with almost no impact on camera detection inference time.

Code — <https://github.com/lihuashengmax/PromptDet>

Introduction

3D object detection is a cornerstone in autonomous driving, with notable advancements in recent years. Most existing camera-based 3D object detectors aim to deduce the spatial arrangement of objects by analyzing visual cues such as color and texture. The low deployment cost and strong scalability on different hardware devices make camera-based detectors quite popular in both academia and industry. Prominent methods like BEVDet (Huang et al. 2021) and BEVFormer (Li et al. 2022b) excel in transforming image-based

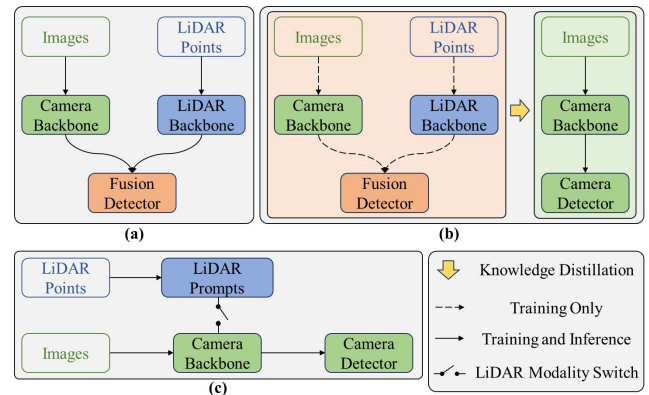


Figure 1: Comparison of our PromptDet with previous detection frameworks. (a) Multi-modal detection needs a more complex network architecture. The model training is time-consuming and occupies a huge memory cost. (b) Though knowledge distillation brings performance gains to camera-only detection, a teacher model needs to be trained first and the whole process is laborious and formidable. (c) Our method uses the LiDAR modality as a flexible prompt with a few additional parameters. PromptDet can conduct multi-modal detection and camera-only detection with better performance than the baseline.

perspective features into a bird’s-eye view, showcasing substantial promise for camera-centric 3D detection tasks. However, unlike LiDAR point clouds, the inherent limitation of images is the absence of precise depth data and 3D perceptive capabilities, which poses a challenge to further enhancing the performance of camera-only 3D object detection methods.

To tackle the limitations of camera-only 3D object detection, several multi-modal approaches (Liang et al. 2022; Liu et al. 2023b; Yan et al. 2023; Deng et al. 2024; Jiao et al. 2023; Li et al. 2022a; Zhang et al. 2023) have been proposed to leverage the multi-modal data, such as LiDAR points and multi-view images, as depicted in Figure 1 (a). The geometric information provided by LiDAR points is instrumental in offsetting the shortcomings of image-based data. These multi-modal methods achieve better results than the single-modality methods, which demonstrate the effectiveness of the fusion strategies. However, these methods often entail

*Corresponding author.

complex network architectures that increase the number of model parameters and computational demands, leading to substantial costs in both the training and inference stages. Moreover, the performance of these multi-modal detectors can significantly degrade in the absence of LiDAR data, potentially falling behind the capabilities of camera-only detectors. Some methods (Liang et al. 2022; Yan et al. 2023; Ge et al. 2023; Wang et al. 2024) focus on solving the problem but usually introduce redundant model architecture or utilize mask-modal data augmentation, which brings more training time or data waste. This highlights the need for more efficient and robust approaches that can maintain high performance even in scenarios where certain data modalities are unavailable.

On the other hand, several methods (Zhou et al. 2023; Chen et al. 2022b; Wang et al. 2023; Zhao et al. 2024) focus on bolstering camera-only 3D object detectors through the application of knowledge distillation (KD) (Hinton, Vinyals, and Dean 2015), as shown in Figure 1 (b). These methods typically employ a sophisticated, pre-trained multi-modal detector as a “teacher” model. This model encapsulates abundant geometric and semantic information about the driving scenario, which is then harnessed to guide and accelerate the learning process of the camera-only “student” model. While these KD-based methods have yielded substantial enhancements in the performance of camera-based detectors, the overall training process remains laborious and time-consuming, which necessitates a pre-trained multi-modal detector. This underscores the need for a more efficient training scheme that can effectively impart the necessary knowledge to camera-only 3D object detectors without the excessive overhead associated with traditional KD strategies.

Recently, prompt learning (Lester, Al-Rfou, and Constant 2021; Liu et al. 2023a; Bahng et al. 2022; Radford et al. 2021) has drawn increasing attention due to its effectiveness and transferability in both natural language processing (NLP) and computer vision (CV) fields. Intuitively, the paradigm of prompt learning is inherently well-suited for 3D object detection, which can exhibit several advantages, including less memory cost, adaptation to multi-modal input, retaining the base model’s potential, and so on. In light of this, we propose PromptDet, a novel lightweight 3D object detection framework, as illustrated in Figure 1 (c). Different from previous works, PromptDet defines a lightweight pipeline, that takes multi-modal detection as injecting LiDAR modality into a camera-based detection model. PromptDet exhibits desirable properties that are not achievable by multi-modal and KD-based methods: (i) The use of LiDAR point clouds is flexible. When LiDAR points are available, PromptDet is a lightweight multi-modal detector with fewer than 2% extra parameters compared to the camera-only base model but improves performance by at most 22.8% mAP and 21.1% NDS. Otherwise, PromptDet degenerates into the original camera-based detector, which can also keep the satisfied performance. (ii) PromptDet can further enhance the performance of the camera-only detector in a single-stage training way. We inject the multi-modal information into the camera-based detector online. With the

cross-modal knowledge injection, the camera-based detector learns better feature representations and achieves a performance improvement of at most 2.4% mAP and 4.0% NDS.

Specifically, PromptDet consists of a camera-only detector and our proposed plug-and-play LiDAR-assisted prompter, which conducts Adaptive Hierarchical Aggregation(AHA) and Cross-Modal Knowledge Injection(CMKI) to get the complete model. AHA first fuses features from point clouds and images at different scales to get hierarchical multi-modal representations. Then AHA integrates these different-grained features in an adaptive way with several convolutional layers. Built on this design, CMKI makes sure the camera-only feature learns complementary information from the AHA output. In order to guarantee the quality of camera-only and multi-modal features, both of them are supervised by the ground truth with the same network architecture, which we call the hybrid supervision strategy. There are totally only about 1% additional parameters compared with the camera-only base model during training.

We summarize the contribution of our work as follows:

- We propose a lightweight 3D object detection framework, dubbed PromptDet. The geometry information from point clouds can be injected into the camera-based detector in a flexible way.
- We propose adaptive hierarchical aggregation and cross-modal knowledge injection modules, which can transfer the multi-modal information to the camera-based detector efficiently.
- We extensively validate the proposed PromptDet under multi-modal and camera-only settings, which demonstrate the effectiveness of this design.

Related Work

Camera-based 3D Object Detection

Camera-based 3D object detectors aim to predict the location and class of objects with images only. Numerous recent works (Huang et al. 2021; Li et al. 2022b; Reading et al. 2021) find converting perspective features from images to bird’s-eye view is effective for camera-based detection. BEVDet (Huang et al. 2021) explicitly predicts the depth distribution of image features and projects them to the BEV (Bird’s Eye View) space through the lifting operation from LSS (Phillion and Fidler 2020). The BEVDet-based approach has been further improved by imposing depth supervision (Li et al. 2023) or temporal aggregation (Huang and Huang 2022), resulting in better performance. BEVFormer (Li et al. 2022b) samples image features with grid-shaped queries and adds the semantic information to BEV space. Though great progress has been made, camera-based detectors still fall behind LiDAR-based and fusion-based counterparts by a large margin.

Knowledge Distillation for Camera Detection

Camera-based 3D object detection is quite popular due to its low deployment cost and knowledge distillation can improve the detectors’ performance while keeping the model architecture unchanged. To fully utilize the information in

the 3D world, many researchers employ multi-modal or cross-modal teachers and empower camera-based detectors with more comprehensive representations (Wang et al. 2023; Zhou et al. 2023; Zhao et al. 2024). These cross-modal distillation techniques endow camera-based detectors with better 3D perceptual capability at the cost of a two-stage complex training process. In a KD-based framework, a teacher model needs to be trained first, which is usually a cross-modal detector or multi-modal detector. Then both the teacher and student are sent to memory and the weights for the teacher model are loaded. Finally, the student model starts to be trained and it receives supervision from both the ground truth and the teacher. The multi-stage training procedure is tiresome and in the final stage, there are two models (student and teacher) loaded into the memory, which brings extra memory cost.

Prompt Learning in Computer Vision

Recently, as a new paradigm, prompt learning has shown its effectiveness and efficiency in computer vision. It directly utilizes the large pre-trained foundation models (e.g., ViT (Dosovitskiy et al. 2020)) and applies their powerful representations in various downstream tasks via a few additional parameters. AdaptFormer (Chen et al. 2022a) outperforms fully fine-tuned models on action recognition benchmarks and only adds fewer than 2% extra parameters to the foundation model. Convpass (Jie and Deng 2022) constructs convolutional bypasses as adaptation modules to adapt the large ViT. ViPT (Zhu et al. 2023) learns modal-relevant prompts and adapts the frozen pre-trained foundation model to multi-modal tracking tasks. These prompt-learning works all inject task-relevant or modal-relevant information through a few additional trainable parameters while keeping the original base model unchanged. In this paper, we exploit the feasibility of migrating the idea of prompt learning to 3D perception and propose a lightweight 3D object detection framework named PromptDet.

Methodology

In this section, we describe our proposed PromptDet in detail. In Section , we introduce the mainstream detection paradigms for LiDAR-only and camera-only 3D object detectors, respectively, and briefly present the knowledge distillation used in Unidistill. In Section , we explain the architecture of PromptDet, which comprises a camera-only detector and the LiDAR-assisted prompter. Finally, in Section , we illustrate the superior properties of our method.

Preliminary

Recent LiDAR-only detectors typically convert unstructured point clouds into regular voxels or pillars through voxelization (Qi et al. 2017; Zhou and Tuzel 2018; Yan, Mao, and Li 2018; Lang et al. 2019), denoted as F_l . These fine-grained features are then extracted and flattened to produce BEV features F_l^{bev} . The BEV encoder takes F_l^{bev} as input and generates deeper features F_l^{en} . The detection head produces response results F_l^{resp} , which include classification and regression heatmaps. Finally, the predictions of the detector are obtained from F_l^{resp} .

Inspired by LiDAR-only 3D object detectors, numerous works attempt to transfer the BEV detection paradigm to camera-based detection. The pixel-wise features of perspective images are projected into 3D space through depth estimation and view transformation. Efficient pooling is employed to obtain pseudo-voxel features F_c , which are analogous to F_l . Following the same procedure as in LiDAR-only detection, F_c^{bev} , F_c^{en} , F_c^{resp} and the final predictions are generated sequentially.

UniDistill is proposed as a cross-modal knowledge distillation framework to align the foreground BEV features of different modalities and three types of distillation losses are calculated. Denoting the detection modality as mod , these losses include feature knowledge distillation on F_{mod}^{bev} , relation knowledge distillation on F_{mod}^{en} , and response knowledge distillation on F_{mod}^{resp} .

PromptDet

Although prompt learning has been proven effective, its direct transfer to 3D object detection is challenging due to the absence of a 3D universal foundation model pre-trained on large-scale datasets. To address this issue, we select the BEVDet-series detectors as the base model and introduce the LiDAR modality as a prompt. The overview of our proposed PromptDet framework is illustrated in Figure 2. The LiDAR-assisted prompter consists of Adaptive Hierarchical Aggregation (AHA) and Cross-Modal Knowledge Injection (CMKI). AHA performs multi-modal fusion to generate more comprehensive representations of the 3D scenario. CMKI injects the fused knowledge into the camera detection pipeline, enhancing the base model’s performance. The entire model is trained in a single stage, with more details provided below.

Adaptive Hierarchical Aggregation Adaptive Hierarchical Aggregation (AHA) aims to generate fine-grained fusion features by fully utilizing both the geometric information from point clouds and the semantic information from images. To achieve this, we perform multi-modal fusion at various granularities and then adaptively aggregate the multi-scale fusion features, as illustrated in Figure 3. For the camera information, the camera-only base model takes multi-view images as input and projects semantic features into 3D space to obtain pseudo voxel features F_c . Point clouds are also voxelized to generate LiDAR voxel features F_l through voxel encoding. Different voxelization scales are considered in both modalities, allowing for comprehensive geometry-semantics interaction at multiple granularities. The multi-scale pseudo voxel features from images are denoted as $F_c^i \in \mathbb{R}^{\frac{D}{2^i S_D} \times \frac{H}{2^i S_H} \times \frac{W}{2^i S_W} \times C}$ and LiDAR voxel features as $F_l^i \in \mathbb{R}^{\frac{D}{2^i S_D} \times \frac{H}{2^i S_H} \times \frac{W}{2^i S_W} \times C}$, where D, H, W is the volumetric size of the entire space, S_D, S_H, S_W is the minimal voxel size in three dimensions and $i = 0, 1, 2$. We dynamically integrate features from F_l^i and F_c^i :

$$W^i = Softmax(C_{lc}^i [C_l^i(F_l^i), C_c^i(F_c^i)]), \quad (1)$$

$$F_{lc}^i = W^i(:, :, :, 0) \odot F_l^i + W^i(:, :, :, 1) \odot F_c^i,$$

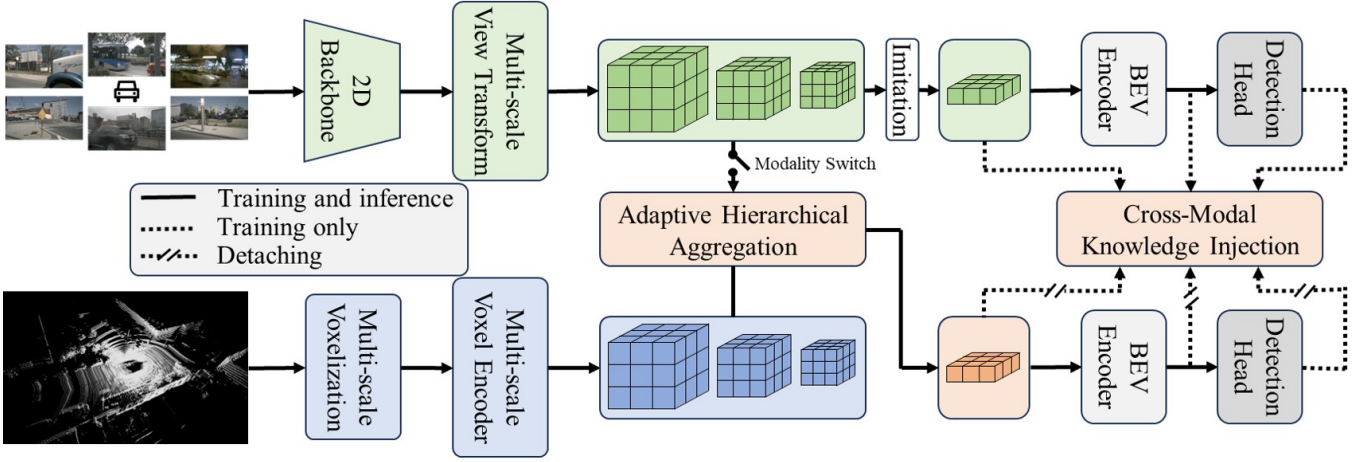


Figure 2: The overview of our proposed PrompDet. The model is composed of a camera-only detector and the LiDAR-assisted prompter including Adaptive Hierarchical Aggregation (AHA) and Cross-Modal Knowledge Injection (CMKI). During model training, LiDAR modality switch is turned off. Multi-modal fusion and online knowledge transfer are performed at the same time. PrompDet supports both LiDAR-camera detection and camera-only detection when inference.

where C_{lc}^i , C_f^i , C_c^i are 3D convolutions, $[\cdot, \cdot]$ is the concatenation along feature channel, $Softmax$ is softmax function along feature channel, $W^i \in \mathbb{R}^{\frac{D}{2^i s_D} \times \frac{H}{2^i s_H} \times \frac{W}{2^i s_W} \times 2}$ is attention weights, \odot represents broadcasting multiplication.

Multi-level fusion features $F_{lc}^i \in \mathbb{R}^{\frac{D}{2^i s_D} \times \frac{H}{2^i s_H} \times \frac{W}{2^i s_W} \times C}$ have different perception ability for objects of different sizes. For example, F_{lc}^0 is more suitable for small object detection because it partitions the entire 3D space with the minimal voxel size and preserves more fine-grained object information. In order to combine fusion features of different voxel scales, we first adjust the grid shape of F_{lc}^0 and F_{lc}^2 to be the same as F_{lc}^1 through interpolation. Then we aggregate the hierarchical features as:

$$W = Softmax(C_f[C_f^0(F_{lc}^0), C_f^1(F_{lc}^1), C_f^2(F_{lc}^2)]),$$

$$F_f = Flatten(\sum_{i=0}^2 W(:, :, :, i) \odot F_{lc}^i), \quad (2)$$

where C_f^0 , C_f^1 , C_f^2 , C_f are 3D convolutions, $W \in \mathbb{R}^{\frac{D}{2^i s_D} \times \frac{H}{2^i s_H} \times \frac{W}{2^i s_W} \times 3}$ is attention weights, $Flatten$ is vertical dimension reduction, $F_f \in \mathbb{R}^{\frac{H}{2^i s_H} \times \frac{W}{2^i s_W} \times C}$ is the fusion BEV feature after adaptive hierarchical aggregation.

Cross-Modal Knowledge Injection Built on fusion features F_f , Cross-Modal Knowledge Injection (CMKI) endows camera detectors with better 3D perception ability. We first get camera-only BEV features from F_c^1 :

$$F_c = Flatten(F_c^1). \quad (3)$$

In the camera-only base model, F_c is sent to the BEV encoder En_{BEV} to extract high-level semantic information of the driving scenario and the detection head H outputs response features:

$$F_c^{en} = En_{BEV}(F_c), F_c^{resp} = H(F_c^{en}). \quad (4)$$

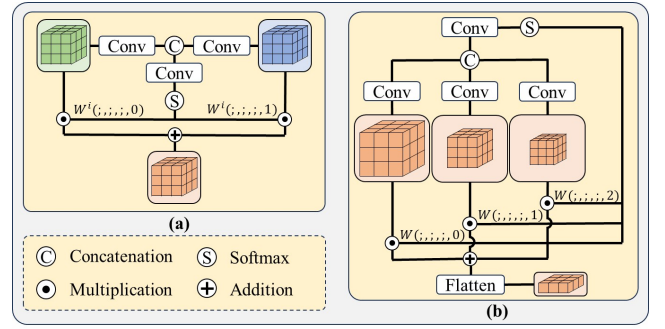


Figure 3: Illustration of Adaptive Hierarchical Aggregation (AHA). LiDAR voxel features and camera pseudo voxel features are fused at different voxel scales, as shown in (a). Hierarchical fusion features are aggregated in a flexible way to obtain the output BEV features, as shown in (b).

We find that though there is a cross-modal domain gap, the BEV encoder and detection head also work for fusion detection as long as multi-modal features F_f join model training:

$$F_f^{en} = En_{BEV}(F_f), F_f^{resp} = H(F_f^{en}). \quad (5)$$

We take this as a hybrid supervision strategy because both kinds of BEV features share the same model parameters and are supervised by the same ground truth.

We conduct CMKI adopting the same approach as Uni-distill with some modifications. For feature knowledge injection, considering F_f is derived from F_c^1 , directly calculating the feature loss between F_f and F_c leads to difficulty in convergence and inferior performance. Therefore, we propose the imitation module to help F_c learn better from F_f and Equation 3 is replaced by:

$$F_c = C_{imi}^{2D}(Flatten(C_{imi}^{3D}(F_c^1))), \quad (6)$$

where C_{imi}^{3D} is a 3D convolutional layer and C_{imi}^{2D} is a 2D convolutional layer. The imitation module reserves during

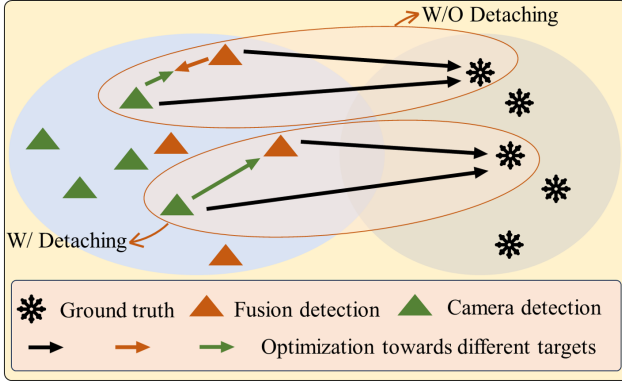


Figure 4: Illustration of data distribution change during PromptDet training. LiDAR-camera detection and camera-only detection are both supervised by the ground truth. If fusion features are not detached in Cross-Modal Knowledge Injection (CMKI), they are also supervised by camera-only ones, which leads to inferior detection performance.

model inference and the added two convolutional layers almost have no impact on inference time. For relation knowledge injection and response knowledge injection, we respectively calculate the relation loss between F_c^{en} and F_f^{en} and the response loss between F_c^{resp} and F_f^{resp} . To make sure the multi-modal features are not influenced by the camera-only branch, we detach the multi-modal features to stop the backward gradient during CMKI, as illustrated in Figure 4. The CMKI loss is shown as:

$$\begin{aligned} \mathcal{L}_{fea} &= FeaDistill(F_f.detach(), F_c, GT), \\ \mathcal{L}_{rel} &= RelDistill(F_f^{en}.detach(), F_c^{en}, GT), \\ \mathcal{L}_{resp} &= RespDistill(F_f^{resp}.detach(), F_c^{resp}, GT), \end{aligned} \quad (7)$$

where $.detach()$ means detached from the computational graph, $FeaDistill()$, $RelDistill()$, $RespDistill()$ are the knowledge distillation methods mentioned in Unidistill, GT is the ground truth, which contains the category and location of objects.

Training and Inference During PromptDet training, the LiDAR modality switch in Figure 2 is turned off. Multi-modal BEV features generated by AHA are supervised by the ground truth. Camera-only BEV features are supervised by both the ground truth and fusion branch. The whole model training is completed in a single stage and the overall training objective \mathcal{L} is defined as:

$$\mathcal{L} = \mathcal{L}_{det}^f + \mathcal{L}_{det}^c + \lambda_1 \mathcal{L}_{fea} + \lambda_2 \mathcal{L}_{rel} + \lambda_3 \mathcal{L}_{resp}, \quad (8)$$

where \mathcal{L}_{det}^f and \mathcal{L}_{det}^c are multi-modal detection loss and camera-only detection loss respectively, which apply the same setting as the baseline. We empirically set λ_1 , λ_2 , λ_3 , and experiments on the sensitivity analysis are shown in supplementary.

There are two situations during PromptDet inference. When images and LiDAR points are both available, the LiDAR modality switch is turned off and PromptDet conducts multi-modal detection. When there are only images

for inference, the LiDAR modality switch is turned on and PromptDet degenerates into the camera-only base model but achieves better performance.

Properties of PromptDet

As a novel 3D object detection framework, PromptDet offers three main advantages: (1) Integration of LiDAR Modality: PromptDet introduces LiDAR data to the camera-based detector, resulting in significant performance gains with only a minimal increase in parameters—about 1% of the base model’s total. This increase primarily comes from convolutional layers in the Adaptive Hierarchical Aggregation and the imitation module. (2) Enhanced Perception Through Cross-Modal Knowledge Injection: PromptDet enhances the camera detector’s perception capabilities via online cross-modal knowledge transfer. This approach is easy and simple during model training. (3) Flexibility of LiDAR Modality During Inference: The LiDAR modality in PromptDet is flexible during inference, thanks to an adjustable LiDAR modality switch. This feature ensures that PromptDet remains a robust 3D object detector even in cases of LiDAR malfunction.

Experiment

Experimental Setup

Dataset and Metrics We evaluate our framework on the nuScenes dataset (Caesar et al. 2020), one of the most challenging benchmarks in autonomous driving. The dataset consists of 1,000 driving scenarios, divided into 700 for training, 150 for validation, and 150 for testing. Each scenario includes synchronized LiDAR point clouds and images from six surrounding cameras, enabling comprehensive multi-modal perception. In alignment with the official evaluation protocol, we report the nuScenes Detection Score (NDS), mean Average Precision (mAP), and five additional metrics across ten popular classes.

Implementation Details We use three camera-only BEVDet-series detectors—BEVDet, BEVDepth, and BEVDet4D—as the base models to evaluate the generalization of our method. All experiments are conducted in PyTorch using 4 NVIDIA A40 GPUs (45GB memory), based on the MMDetection3D (Contributors 2020) codebase. Unless otherwise specified, we use ResNet-50 (He et al. 2016) as the image backbone, with images resized to 256×704. The detection range is set to [-51.2m, 51.2m] × [-51.2m, 51.2m], and we use dynamic point cloud voxelization (Zhou et al. 2020) to reduce unnecessary computational costs. The moderate voxel size is set to [0.8m, 0.8m, 0.8m]. AdamW (Loshchilov and Hutter 2017) is used as the optimizer with a step-scheduled learning rate. When compared with other methods, we train the model for 30 epochs with CBGS (Zhu et al. 2019); other experiments are trained for 36 epochs without CBGS. We apply synchronous BEV data augmentations to both LiDAR and camera modalities for better generalization. For BEVDet4D, we perform temporal camera fusion at the voxel level.

Method	Mode	Backbone	mAP \uparrow	NDS \uparrow	mATE \downarrow	mASE \downarrow	mAOE \downarrow	mAVE \downarrow	mAAE \downarrow
BEVFusion (Liang et al. 2022)	L&C	VoxelNet SwinT	67.9	71.0	-	-	-	-	-
BEVFusion* (Liang et al. 2022)	L&C	PointPillars SwinT	53.5	60.4	-	-	-	-	-
PointPillars (Lang et al. 2019)	L	PointPillars	45.3	30.5	-	-	-	-	-
FCOS3D (Wang et al. 2021)	C	R101	34.3	41.5	72.5	26.3	42.2	129.2	15.3
PETR (Liu et al. 2022)	C	R101	35.7	42.1	71.0	27.0	49.0	88.5	22.4
PETR	C	SwinT	36.1	43.1	73.2	27.3	49.7	80.8	18.5
DETR3D (Wang et al. 2022)	C	R101	34.9	43.4	71.6	26.8	37.9	84.2	20.0
BEVFusion-C (Liu et al. 2023b)	C	SwinT	35.6	41.2	66.8	27.3	56.1	89.6	25.9
BEVFormer (Li et al. 2022b)	C	R101	41.6	51.7	67.3	27.4	37.2	39.4	19.8
BEVDepth (Li et al. 2023)	C	R101	41.2	53.5	-	-	-	-	-
PromptDet	C	R101	43.3	56.9	44.5	23.6	35.1	32.1	12.6
	L&C	R101	56.2	64.1	39.8	24.1	34.6	29.5	12.1

Table 1: Performance comparison on the nuScenes validation set. L and C represent the input modality, *i.e.*, LiDAR and camera. * means the PointPillars version of BEVFusion.

Method	LaP	Mode	mAP	NDS
BEVDet	-	C	27.7	35.3
	✓	L&C	50.5	56.4
	✓	C	30.1	39.3
BEVDet4D	-	C	28.9	41.4
	✓	L&C	52.3	58.5
	✓	C	30.7	44.5
BEVDepth	-	C	32.8	44.1
	✓	L&C	55.9	62.7
	✓	C	35.2	48.0

Table 2: Performance comparison of different camera detectors with our proposed LiDAR-assisted Prompter, which is shortened as LaP.

Main Results

We begin by using BEVDepth as the base model to compare our method, PromptDet, with state-of-the-art methods on the nuScenes validation set. The results are shown in Table 1. Using ResNet-101 (He et al. 2016) as the image backbone, our method, PromptDet-C (with images only), improves the baseline by 2.1% mAP and 3.4% NDS. When point clouds are available, the lightweight multi-modal detector, PromptDet-LC, outperforms the PointPillars version of BEVFusion (Liang et al. 2022) by 2.7% mAP and 3.7% NDS, without using the time-consuming LiDAR backbone. Next, we compare detection performance before and after introducing the LiDAR-assisted prompter into three representative camera-based models to evaluate the generalization of our method. The experiments, conducted on the nuScenes validation set, are summarized in Table 2. We observe significant performance gains when point clouds are available, with both mAP and NDS improving substantially for all three baselines. The largest performance margin (*i.e.*, 22.8% mAP and 21.1% NDS) is achieved by adding the LiDAR-assisted prompter to BEVDet with few additional parameters compared with the base model. Even when only images are used for inference, PromptDet still enhances these camera-based models. For the base model BEVDet, the most significant improvement is 2.4% mAP and 4.0% NDS. For the temporal fusion model BEVDet4D, we still see gains of 1.8% mAP and 3.1% NDS.

	AHA	CMKI	Det.	mAP	NDS	mAP*	NDS*
a				27.1	34.8	49.7	55.1
b	✓			27.3	34.2	50.9	55.9
c		✓		29.1	38.1	48.8	54.3
d		✓	✓	29.6	39.0	49.8	55.1
e	✓	✓		29.3	38.9	49.4	56.0
f	✓	✓	✓	30.1	39.3	50.5	56.4

Table 3: Ablation study of the model architecture and training setting. Det. means detaching fusion features in Cross-Modal Knowledge Injection (CMKI). * means multi-modal detection results of PromptDet-LC.

Ablation Studies

We conduct the ablations on the detailed components of PromptDet with BEVDet as the base model.

Effect of Individual Component We ablate the contribution of each component in PromptDet, as shown in Table 3. In baseline (a), we fuse multi-modal features at a single scale and use hybrid supervision without CMKI. In models (c) and (d), we introduce CMKI and compare the performance with and without detaching fusion features in CMKI. CMKI improves PromptDet-C by 2.0% mAP and 3.3% NDS but reduces PromptDet-LC by 0.9% mAP and 0.8% NDS. The detachment operation reduces this performance drop, supporting our hypothesis. Finally, we introduce AHA and perform multi-modal fusion at different scales, as shown in (b), (e), and (f). Compared to “without AHA”, PromptDet-LC shows consistent improvement. The most significant gains are in model (f), with 3.0% mAP and 4.5% NDS for PromptDet-C, and 0.8% mAP and 1.3% NDS for PromptDet-LC. To better understand the impact of CMKI during training, we show the trends of CMKI loss for models (a) and (d) as well as the detailed CMKI loss for model (d) in Figure 5. Without CMKI’s constraints, the loss tends to diverge during training and the data distribution difference between PromptDet-LC and PromptDet-C increases. Even with CMKI active, the loss decreases slowly, likely because PromptDet-LC and PromptDet-C are trained together to align with the ground truth. CMKI keeps the two data distributions closer and aids in cross-modal knowledge transfer.




Figure 5: Illustration to show the effect of Cross-Modal Knowledge Injection (CMKI). With CMKI working, the data distribution difference between fusion features and camera features narrows down and all three kinds of CMKI loss decrease slowly.

	Conv3D	Conv2D	mAP	NDS	mAP*	NDS*
a			28.4	36.8	50.4	56.5
b	✓		29.6	39.0	50.1	56.2
c		✓	28.9	37.5	50.1	56.7
d	✓	✓	30.1	39.3	50.5	56.4

Table 4: Ablation study of different imitation module designs. * means multi-modal detection results of PromptDet-LC.

Design of Imitation Module The imitation module is designed to force camera-only features to learn from fusion features. It consists of a 3D convolutional layer, a flatten operation, and a 2D convolutional layer. We tested different designs of the imitation module (*i.e.*, without any convolutional layers, with only a 3D convolutional layer, with only a 2D convolutional layer, and the complete architecture). As shown in Table 4, different imitation modules have little impact on PromptDet-LC, but omitting all convolutional layers leads to poorer performance for PromptDet-C. Both the 2D and 3D convolutional layers help camera-only features better represent the 3D scene, with the complete design providing the most significant improvement of 1.7% mAP and 2.5% NDS.

Further Discussion

Model Efficiency In Table 5, we report the model efficiency, including parameters and inference speed. PromptDet-C adds only the imitation module to the baseline, consisting of two convolutional layers, which has minimal impact on inference time. Therefore, we focus on PromptDet-LC when point clouds are available, with all inferences conducted on a single NVIDIA A40 GPU. The LiDAR-assisted prompter adds fewer than 2% extra parameters during training and inference. PromptDet-LC, based on BEVDepth, takes 115.4 ms for inference, running much faster than the PointPillars version of BEVFusion (Liang et al. 2022) while achieving better detection performance.

Qualitative Results In Figure 6, we visualize the prediction results from both the multi-camera view and LiDAR top view. Compared to BEVDet as the baseline, PromptDet-LC significantly improves perception and detects objects missed by the baseline. Even without LiDAR input, PromptDet-C reduces false positives and better localizes objects, thanks to

Methods	Mode	Params(M)	Latency(ms)
BEVFusion	L&C	90.19	1453.8
BEVFusion*	L&C	89.46	1428.6
BEVDet	C	44.25	41.9
+LaP	L&C	44.86(+1.38%)	100.5
BEVDet4D	C	52.48	73.0
+LaP	L&C	53.49(+1.92%)	158.4
BEVDepth	C	66.67	55.7
+LaP	L&C	67.28(+0.91%)	115.4

Table 5: Comparison of model efficiency. * means the PointPillars version of BEVFusion. LaP denotes our proposed LiDAR-assisted Prompter.




Figure 6: The comparison of BEVDet (Huang et al. 2021) and our proposed PromptDet. The red boxes and green boxes denote the ground truth and detection results, respectively. In the LiDAR top view, We use the yellow arrows to mark the objects neglected by baseline and orange arrows to mark better localization in PromptDet-LC and PromptDet-C.

our Cross-Modal Knowledge Injection during training.

Conclusion

In this paper, we propose a lightweight 3D object detection framework dubbed PromptDet, which consists of a camera detector and the LiDAR-assisted prompter. PromptDet conducts LiDAR-camera fusion through AHA and is a lightweight multi-modal detector if both images and LiDAR points are available. Thanks to CMKI, PromptDet still outperforms the baseline with only images as input. AHA and CMKI constitute the plug-and-play LiDAR-assisted prompter and the whole framework training is simple with few extra model parameters. Extensive experiments on the challenging nuScenes benchmark validate the effectiveness of our method. We believe the paradigm of PromptDet is compatible with diverse camera-based detectors and can be extended to more multi-camera perception tasks, such as occupancy prediction, BEV segmentation, and so on.

Acknowledgments

This work was supported by USTC-JAC Smart Electric Vehicle Joint Lab, and by the Key Science and Technology Program of Anhui under Grant 202203a05020012.

References

- Bahng, H.; Jahanian, A.; Sankaranarayanan, S.; and Isola, P. 2022. Exploring visual prompts for adapting large-scale models. *arXiv preprint arXiv:2203.17274*.
- 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*, 11621–11631.
- Chen, S.; Ge, C.; Tong, Z.; Wang, J.; Song, Y.; Wang, J.; and Luo, P. 2022a. Adaptformer: Adapting vision transformers for scalable visual recognition. *Advances in Neural Information Processing Systems*, 35: 16664–16678.
- Chen, Z.; Li, Z.; Zhang, S.; Fang, L.; Jiang, Q.; and Zhao, F. 2022b. Bevdistill: Cross-modal bev distillation for multi-view 3d object detection. *arXiv preprint arXiv:2211.09386*.
- Contributors, M. 2020. MMDetection3D: OpenMMLab next-generation platform for general 3D object detection.
- Deng, J.; Zhang, S.; Dayoub, F.; Ouyang, W.; Zhang, Y.; and Reid, I. 2024. PolFusion: Multi-Modal 3D Object Detection via Fusion at Points of Interest. *arXiv preprint arXiv:2403.09212*.
- Dosovitskiy, A.; Beyer, L.; Kolesnikov, A.; Weissenborn, D.; Zhai, X.; Unterthiner, T.; Dehghani, M.; Minderer, M.; Heigold, G.; Gelly, S.; et al. 2020. An image is worth 16x16 words: Transformers for image recognition at scale. *arXiv preprint arXiv:2010.11929*.
- Ge, C.; Chen, J.; Xie, E.; Wang, Z.; Hong, L.; Lu, H.; Li, Z.; and Luo, P. 2023. Metabev: Solving sensor failures for 3d detection and map segmentation. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 8721–8731.
- He, K.; Zhang, X.; Ren, S.; and Sun, J. 2016. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, 770–778.
- Hinton, G.; Vinyals, O.; and Dean, J. 2015. Distilling the knowledge in a neural network. *arXiv preprint arXiv:1503.02531*.
- Huang, J.; and Huang, G. 2022. Bevdet4d: Exploit temporal cues in multi-camera 3d object detection. *arXiv preprint arXiv:2203.17054*.
- Huang, J.; Huang, G.; Zhu, Z.; Ye, Y.; and Du, D. 2021. Bevdet: High-performance multi-camera 3d object detection in bird-eye-view. *arXiv preprint arXiv:2112.11790*.
- Jiao, Y.; Jie, Z.; Chen, S.; Chen, J.; Ma, L.; and Jiang, Y.-G. 2023. Msmdfusion: Fusing lidar and camera at multiple scales with multi-depth seeds for 3d object detection. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 21643–21652.
- Jie, S.; and Deng, Z.-H. 2022. Convolutional bypasses are better vision transformer adapters. *arXiv preprint arXiv:2207.07039*.
- 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*, 12697–12705.
- Lester, B.; Al-Rfou, R.; and Constant, N. 2021. The power of scale for parameter-efficient prompt tuning. *arXiv preprint arXiv:2104.08691*.
- Li, Y.; Chen, Y.; Qi, X.; Li, Z.; Sun, J.; and Jia, J. 2022a. Unifying voxel-based representation with transformer for 3d object detection. *Advances in Neural Information Processing Systems*, 35: 18442–18455.
- Li, Y.; Ge, Z.; Yu, G.; Yang, J.; Wang, Z.; Shi, Y.; Sun, J.; and Li, Z. 2023. Bevdepth: Acquisition of reliable depth for multi-view 3d object detection. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 37, 1477–1485.
- Li, Z.; Wang, W.; Li, H.; Xie, E.; Sima, C.; Lu, T.; Qiao, Y.; and Dai, J. 2022b. Bevformer: Learning bird’s-eye-view representation from multi-camera images via spatiotemporal transformers. In *European conference on computer vision*, 1–18. Springer.
- Liang, T.; Xie, H.; Yu, K.; Xia, Z.; Lin, Z.; Wang, Y.; Tang, T.; Wang, B.; and Tang, Z. 2022. Bevfusion: A simple and robust lidar-camera fusion framework. *Advances in Neural Information Processing Systems*, 35: 10421–10434.
- Liu, P.; Yuan, W.; Fu, J.; Jiang, Z.; Hayashi, H.; and Neubig, G. 2023a. Pre-train, prompt, and predict: A systematic survey of prompting methods in natural language processing. *ACM Computing Surveys*, 55(9): 1–35.
- Liu, Y.; Wang, T.; Zhang, X.; and Sun, J. 2022. Petr: Position embedding transformation for multi-view 3d object detection. In *European Conference on Computer Vision*, 531–548. Springer.
- Liu, Z.; Tang, H.; Amini, A.; Yang, X.; Mao, H.; Rus, D. L.; and Han, S. 2023b. Bevfusion: Multi-task multi-sensor fusion with unified bird’s-eye view representation. In *2023 IEEE international conference on robotics and automation (ICRA)*, 2774–2781. IEEE.
- Loshchilov, I.; and Hutter, F. 2017. Decoupled weight decay regularization. *arXiv preprint arXiv:1711.05101*.
- Phillion, J.; and Fidler, S. 2020. Lift, splat, shoot: Encoding images from arbitrary camera rigs by implicitly unprojecting to 3d. In *Computer Vision—ECCV 2020: 16th European Conference, Glasgow, UK, August 23–28, 2020, Proceedings, Part XIV 16*, 194–210. Springer.
- Qi, C. R.; Su, H.; Mo, K.; and Guibas, L. J. 2017. Pointnet: Deep learning on point sets for 3d classification and segmentation. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, 652–660.
- Radford, A.; Kim, J. W.; Hallacy, C.; Ramesh, A.; Goh, G.; Agarwal, S.; Sastry, G.; Askell, A.; Mishkin, P.; Clark, J.;

- et al. 2021. Learning transferable visual models from natural language supervision. In *International conference on machine learning*, 8748–8763. PMLR.
- 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*, 8555–8564.
- Wang, S.; Caesar, H.; Nan, L.; and Kooij, J. F. 2024. Unibev: Multi-modal 3d object detection with uniform bev encoders for robustness against missing sensor modalities. In *2024 IEEE Intelligent Vehicles Symposium (IV)*, 2776–2783. IEEE.
- Wang, T.; Zhu, X.; Pang, J.; and Lin, D. 2021. Fcos3d: Fully convolutional one-stage monocular 3d object detection. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 913–922.
- Wang, Y.; Guizilini, V. C.; Zhang, T.; Wang, Y.; Zhao, H.; and Solomon, J. 2022. Detr3d: 3d object detection from multi-view images via 3d-to-2d queries. In *Conference on Robot Learning*, 180–191. PMLR.
- Wang, Z.; Li, D.; Luo, C.; Xie, C.; and Yang, X. 2023. Distillbev: Boosting multi-camera 3d object detection with cross-modal knowledge distillation. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 8637–8646.
- Yan, J.; Liu, Y.; Sun, J.; Jia, F.; Li, S.; Wang, T.; and Zhang, X. 2023. Cross modal transformer: Towards fast and robust 3d object detection. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 18268–18278.
- Yan, Y.; Mao, Y.; and Li, B. 2018. Second: Sparsely embedded convolutional detection. *Sensors*, 18(10): 3337.
- Zhang, Y.; Zhang, Q.; Hou, J.; Yuan, Y.; and Xing, G. 2023. Unleash the potential of image branch for cross-modal 3d object detection. In *Advances in Neural Information Processing Systems*, volume 36, 51562–51583.
- Zhao, H.; Zhang, Q.; Zhao, S.; Chen, Z.; Zhang, J.; and Tao, D. 2024. Simdistill: Simulated multi-modal distillation for bev 3d object detection. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 38, 7460–7468.
- Zhou, S.; Liu, W.; Hu, C.; Zhou, S.; and Ma, C. 2023. Uni-Distill: A Universal Cross-Modality Knowledge Distillation Framework for 3D Object Detection in Bird’s-Eye View. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 5116–5125.
- Zhou, Y.; Sun, P.; Zhang, Y.; Anguelov, D.; Gao, J.; Ouyang, T.; Guo, J.; Ngiam, J.; and Vasudevan, V. 2020. End-to-end multi-view fusion for 3d object detection in lidar point clouds. In *Conference on Robot Learning*, 923–932. PMLR.
- Zhou, Y.; and Tuzel, O. 2018. Voxelnet: End-to-end learning for point cloud based 3d object detection. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, 4490–4499.
- Zhu, B.; Jiang, Z.; Zhou, X.; Li, Z.; and Yu, G. 2019. Class-balanced grouping and sampling for point cloud 3d object detection. *arXiv preprint arXiv:1908.09492*.
- Zhu, J.; Lai, S.; Chen, X.; Wang, D.; and Lu, H. 2023. Visual prompt multi-modal tracking. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 9516–9526.