PF3Det: A Prompted Foundation Feature Assisted Visual LiDAR 3D Detector

Li, Kaidong; Zhang, Tianxiao; Peng, Kuan-Chuan; Wang, Guanghui

TR2025-076 June 07, 2025

Abstract

3D object detection is crucial for autonomous driving, leveraging both LiDAR point clouds for precise depth in- formation and camera images for rich semantic informa- tion. Therefore, the multi-modal methods that combine both modalities offer more robust detection results. However, efficiently fusing LiDAR points and images remains challenging due to the domain gaps. In addition, the performance of many models is limited by the amount of high quality labeled data, which is expensive to create. The recent advances in foundation models, which use large-scale pre- training on different modalities, enable better multi-modal fusion. Combining the prompt engineering techniques for efficient training, we propose the Prompted Foundational 3D Detector (PF3Det), which integrates foundation model encoders and soft prompts to enhance LiDAR-camera feature fusion. PF3Det achieves the state-ofthe-art results under limited training data, improving NDS by 1.19% and mAP by 2.42% on the nuScenes dataset, demonstrating its efficiency in 3D detection.

IEEE Conference on Computer Vision and Pattern Recognition (CVPR) Workshop 2025

© 2025 IEEE. Personal use of this material is permitted. Permission from IEEE must be obtained for all other uses, in any current or future media, including reprinting/republishing this material for advertising or promotional purposes, creating new collective works, for resale or redistribution to servers or lists, or reuse of any copyrighted component of this work in other works.

Mitsubishi Electric Research Laboratories, Inc. 201 Broadway, Cambridge, Massachusetts 02139

PF3Det: A Prompted Foundation Feature Assisted Visual LiDAR 3D Detector

Kaidong Li^{1*} Tianxiao Zhang¹ Kuan-Chuan Peng² Guanghui Wang³

¹University of Kansas

²Mitsubishi Electric Research Laboratories (MERL) ³Toronto Metropolitan University

foronto Metropontan On

likaidong910@gmail.com tianxiao@ku.edu

edu kpeng@merl.com

wangcs@torontomu.ca

Abstract

3D object detection is crucial for autonomous driving, leveraging both LiDAR point clouds for precise depth information and camera images for rich semantic information. Therefore, the multi-modal methods that combine both modalities offer more robust detection results. However, efficiently fusing LiDAR points and images remains challenging due to the domain gaps. In addition, the performance of many models is limited by the amount of high quality labeled data, which is expensive to create. The recent advances in foundation models, which use large-scale pretraining on different modalities, enable better multi-modal fusion. Combining the prompt engineering techniques for efficient training, we propose the Prompted Foundational 3D Detector (PF3Det), which integrates foundation model encoders and soft prompts to enhance LiDAR-camera feature fusion. **PF3Det** achieves the state-of-the-art results under limited training data, improving NDS by 1.19% and mAP by 2.42% on the nuScenes dataset, demonstrating its efficiency in 3D detection.

1. Introduction

3D object detection is an essential task in autonomous driving, involving the localization and classification of objects in 3D space, which is different from 2D object detection [36][37][58][59] that classifies and localizes the objects in 2D visual data. The LiDAR point clouds provide precise depth information. Early LiDAR-only methods like VoxelNet [64] achieved promising results. However, the sparsity of point clouds makes it difficult to detect small or distant objects. In contrast, high-resolution images from cameras provide rich semantic details, but camera-only data do not have precise distance information, which is why they usually do not match the performance of LiDAR-only approaches. Wang *et al.* [46] find that depth prediction error



Figure 1. The illustration of the **PF3Det** pipeline. The blue components are the proposed modules.

contributes 20% to the overall error of camera-only models. Therefore, the complementary strength of LiDAR and camera makes multi-modal detection a more robust solution.

Since the LiDAR-based methods demonstrated stronger performance compared to the camera-based ones, the main stream approach is based on LiDAR detectors, and tries to fuse additional information from the camera modality into the LiDAR models. The main challenge of multi-modal detection is how to fuse the two modalities efficiently. Some multi-modal detectors try to establish point-to-pixel correspondence and query rich camera features from 3D data points [3, 30, 31]. However, the mismatched data densities inevitably result in loss of 2D information. Another paradigm of visual LiDAR detectors first estimates the depth of 2D data and then lifts the 2D pixels into the 3D space [21, 32, 54]. These methods have improved the prediction quality thanks to integrating more information from the 2D data. However, lifting a 2D pixel into 3D requires depth estimation, which is still prone to errors [46]. Also the feature spaces from the two different modalities are not aligned, thus diminishing the effect of the final feature maps before detection.

Another challenge for existing methods is that they heavily rely on extensive and high-quality annotations [12, 24, 34]. Labeling a single scene demands considerable efforts,

^{*} This work was done when Kaidong Li was an intern at MERL.

which makes large-scale labeling impractical for applications such as robotics and autonomous vehicles. Even with existing dataset, Gao *et al.* [12] show the performance of the state-of-the-art (SOTA) models are limited by the amount of data. The situation is worse in industry, since many datasets cannot be used for commercial purposes. Therefore, providing a data-efficient 3D detector is crucial.

In recent years, there has been significant progress in developing multi-modal foundation models, which are trained on large-scale datasets and can be adapted for a wide range of downstream tasks [1, 19, 23, 43, 51]. Their success can be mainly attributed to the capability of large language models (LLMs), which enables massive scaling up on data. The training of LLMs on noisy web-scale data inspires the training on other modalities [43]. These models like GPT-3 [4], with billions of parameters, have demonstrated impressive performance in zero- and few-shot learning, achieving strong results without requiring task-specific data or parameter updates. The success facilitates the training of vision models to the feature space of LLMs. CLIP uses contrastive pre-training with LLMs to generate a model with promising zero-shot classification and detection performance [43]. More modalities soon join the foundation model family, e.g., ULIP [53] proposed a method to create language, image, and point triplet data, and released a point encoder that learns to encode a point cloud into a unified feature space with CLIP. 3D-LLM [16] interacts with 3D scenes using natural language. It is obvious the combined feature encoders can help to fuse features from different modalities more efficiently.

To adapt a pre-trained network to downstream tasks, prompt engineering is a promising research direction. The foundation models have massive amount of parameters, exceeding billions of parameters [35]. It becomes impractical to efficiently fine-tune a foundation model when the data and computational resources are limited [2]. Soft prompts are tensors, which can be fine-tuned during training while keeping other model weights unchanged. It can adapt to different tasks without altering original weights, which helps preserve the foundation models' capabilities. VPT [20] introduces a set of vision patches at different layers of a ViT-based model [10].

Motivated by the capabilities of foundation models, we proposed Prompted Foundational 3D Detector (**PF3Det**), which is a visual LiDAR 3D detector as shown in Fig. 1. It takes the encoder from a foundation model to help extract aligned feature maps from different modalities. The multi-modal soft prompts are inserted at multiple levels to assist the fusion of camera and LiDAR features. **PF3Det** achieves the SOTA performance in efficient 3D detection. According to our experiments, **PF3Det** increases nuScenes detection score (NDS) by 1.19% and mAP by 2.42% on the nuScenes dataset [5] under limited training data.

Our contributions can be summarized below:

- We propose PF3Det, a visual LiDAR 3D detection model architecture that can efficiently learn to predict highquality 3D objects with a small amount of available data.
- To achieve efficient learning, our proposed **PF3Det** takes multi-modal foundational features and bridges the modality domain gaps in the bird-eye-view (BEV) stage by incorporating the soft prompts for convolutional layers.
- Our proposed **PF3Det** achieves the SOTA performance on the 3D object detection task under limited training data. We run extensive experiments to explore different settings and provide guidance on how to select model parameters.

2. Related work

2.1. 3D visual-LiDAR detection

3D visual-LiDAR detection [3, 29, 32, 52, 63] has become increasingly popular in autonomous driving and advanced robotics, taking advantages of the complementary strengths of camera and LiDAR sensors. LiDAR provides precise 3D spatial awareness, while cameras provide rich color and texture details. The sensor fusion techniques can be generally divided into three categories, early fusion, intermediate fusion, and late fusion. These categories are based on when the camera and LiDAR features are fused. Early fusion first extracts information from the 2D camera, then adds the information to LiDAR data before LiDAR detectors. Frustum PointNets [41] uses a 2D detector to generate 2D bounding box (bbox) proposals, and then those 2D bboxes are used to produce frustum proposals to help 3D point cloud segmentation. Du et al. [11] proposed similar frustum-based method, but in later stages, implemented a model-fitting algorithm to refine point segmentation using generalized object models. PointPainting [45] first runs pixel-level segmentation on 2D images, and passes pixel labels to each points via pixel-to-point correlation. However, early fusion methods are usually sequential, which increases latency [38]. Late fusion happens after image and point branches generate their 2D and 3D bboxes separately, which means that the image-based and point-based models can run in parallel. Pang et al. proposed CLOCS [39], which converts 2D and 3D detection candidates into sparse tensors and generates final results from the sparse tensors. However, late fusion does not provide advanced integration of the two modalities, losing the rich semantic features from the other sensor. Intermediate fusion refers to any fusion methods that happen at the stages anywhere between early fusion and late fusion, e.g., at backbone stage [44, 49], at proposal generation stage [7, 25], etc. This type of detectors enables deeper integration of two modalities, generates features with higher quality, and witnesses many SOTA 3D detectors [21, 22, 55]. But one challenge is that the camera and LiDAR sensor data are heterogeneous, and fusion will even cause noises if it is not done properly. TransFusion [3] proposed soft-association to better align the data of these two sensors. MSMDFusion [21] designed two modules, Multi-Depth Unprojection (MDU) and Gated, Modality-Aware Convolution (GMA-Conv), to create alignment and deep fusion between camera and LiDAR's features.

2.2. Multi-modal large language models

Multi-modal Large Language Models (MLLMs) are designed to achieve a comprehensive understanding across various modalities, including audio [18], images [43], point clouds [27, 53], etc. For visual data, Vision Transformers [10][61][60] are usually employed as the visual encoders. The key motivation is that integrating the language modality can significantly enhance a model's ability to interpret complex interactions within diverse input modalities [28, 51]. One approach leverages contrastive learning to align representations from different modalities to joint feature spaces [43, 53]. Notable examples, such as CLIP [43] and ALIGN [19], have demonstrated impressive zero-shot generalization abilities, inspiring a range of applications. For instance, PointCLIP [57] leverages CLIP for zero-shot 3D multi-view classification. Another approach uses language as a tool for interacting with other modalities, enabling versatile interactions between human and deep learning models [53].

To the best of our knowledge, given the capability of MLLMs, there is not yet an MLLM-assisted 3D visual Li-DAR detection model even with the existence of the models like Meta Transformer [62], which aligns a wide range of modalities including 3D point clouds. However, all these MLLMs only align indoor point clouds for simple objects, which are more dense and consistent compared to LiDAR 3D scans. Empirically it is difficult to bridge the domain gap between these 3D point data, which is the main reason such methods do not exist yet.

2.3. Prompt engineering

Prompt engineering is a recent promising method to bridge the domain gap between general foundation models and downstream tasks [26]. This technique gains popularity when large foundation models demonstrate their capability to adapt to a wide range of applications [20]. Comparing to regular fine-tuning, where the knowledge of the foundation models is at the risk of catastrophic forgetting, prompt engineering can keep the original foundation models unchanged [50], and it is proven in many fields to be useful [4, 17, 20, 26, 35, 40, 48]. Generally speaking, prompt engineering can be divided into two categories, hard prompts and soft prompts. **Hard prompt tuning** involves manually engineered and human interpretable inputs, sometimes in the form of texts [16, 42]. For example, *instruction prompting* uses explicit instructions to guide the models to gener-

ate desired response [16, 42]. This approach reduces ambiguity and enhances the model's ability to produce accurate and relevant results. In-context learning is often used where rapid adaptation is needed [4, 9]. This method enables the foundation models to generalize to specific tasks by showing the models just a few examples or prompts, and chainof-thought prompting breaks down complex tasks into a series of simpler steps or instructions within the prompt, allowing the model to process and solve each step sequentially [48]. We refer the readers to these survey papers [13, 26] for more details. This category of methods requires careful crafted prompts and is usually labor-intensive [13]. Soft prompt tuning on the other hand uses learnable parameters [17, 26, 35]. It is a method that has emerged as an effective way to fine-tune large language models (LLMs) for specific tasks by modifying the input prompts rather than adjusting the model's parameters directly. This approach is particularly beneficial when the computational resources or labeled data are limited [13]. These prompt vectors can be added directly to the input or to any model layers [20]. During fine-tuning, the pre-trained model will be frozen, and only the prompt vectors are updated to achieve performance improvements.

3. Prompted foundational 3D detector

We propose Prompted Foundational 3D Detector (PF3Det), a prompted foundation-assisted 3D visual LiDAR detector. This method has two modules, a foundational branch, which utilizes the features from a multi-modal foundation model encoder to assist regular image features, and multimodal soft-prompt adapter, which accepts soft prompts in the visual-LiDAR BEV layers to improve detection performance. The inputs of our detector are LiDAR points and their corresponding single- or multi-view images. Our PF3Det outputs 3D detection bounding boxes. The input images first go through a regular detector backbone and a multi-modal foundational image encoder in parallel. Their features are concatenated to produce combined image foundational features. Then these image features and input point clouds are fed into a MSMDFusion-based [21] fusion architecture to generate multi-modal BEV features. The BEV features are then integrated with the fine-tuned soft prompts to produce the final adapted multi-modal features for the 3D bbox prediction. The overall model architecture is illustrated in Fig. 2.

3.1. Foundational branch

Regular image backbone can be formulated in Eq. 1, where L_i is the *i*-th layer, W_i is the linear projection at residual connection to match dimensions, and x_i is the *i*-th layer output (x_0 is the input images). The multi-scale features x_i will be fused with multi-scale point features.



Figure 2. The architecture of our proposed **PF3Det**. The Foundational branch is added in parallel with the original image backbone. Multi-modal soft-prompt adapter is inserted at the BEV feature level.

$$x_i = L_i(x_{i-1}) + W_i x_{i-1}, \ i = 1, 2, ..., N$$
(1)

Multi-modal foundation models have learnt to extract inputs from different modalities into unified feature spaces. It can extract semantic-rich features to better capture object relations. To help the model extract more high-level information from the input images, we add another branch, the foundational extractor $F(\cdot)$, to obtain the foundational feature vectors, $v = F(x_0)$, $v \in \mathbb{R}^d$, where d is the dimension of the feature vectors. Then the vectors v is concatenated to the feature maps x_N from the last layer of the regular image backbone. Then the concatenated features are fed into a feature pyramid networks (FPNs) [33] to produce the final multi-scale image features.

3.2. Multi-modal soft-prompt adapter

By borrowing the features from the image encoder of a foundation model directly, there will be inevitably some domain gap between its feature vector and the 3D detection features. Therefore, we propose the multi-modal soft-prompt adapter to integrate soft prompts into the detector architecture.

The original detector's architecture after multi-modal BEV features \mathbf{F}_{BEV} can be formulated as Eq. 2. F_{BEV} is obtained by converting the point features, virtual point features, and image features together into the BEV space.

 F_{BEV} is fed into a BEV fusion encoder, E_{Fusion} , which is a combination of four different one-layer CNNs in parallel, followed by another CNN layer to fuse the outputs together. The fused features, $\mathbf{F}_{BEVFused} \in \mathbb{R}^{C \times H \times W}$, then are used to extract multi-scale BEV features, $\mathbf{F}^{i}_{MSBfeat}$. (where *i* is the scale index), through a BEV backbone B_{BEV} and BEV FPN detector P_{BEV} .

$$\mathbf{F}_{BEVFused} = E_{Fusion}(\mathbf{F}_{BEV})$$

$$(\mathbf{F}_{BEVfeat.}^{i}) = B_{BEV}(\mathbf{F}_{BEVFused})$$

$$(\mathbf{F}_{MSBfeat.}^{i}) = P_{BEV}(\mathbf{F}_{BEVfeat.}^{i})$$

$$\mathbf{F}_{BEV}, \quad \mathbf{F}_{BEVFused} \in \mathbb{R}^{C \times H \times W}$$
(2)

Single-level prompts. The prompts are only added to one single layer in the detector's architecture, which requires less computational resources and is faster to train. Only the single-level prompts, $P_S \in \mathbb{R}^{C_p^0 \times H \times W}$, are concatenated to the original BEV feature \mathbf{F}_{BEV} to produce a $(C_p^0 + C) \times H \times W$ BEV feature map. We modify the input dimension of the BEV encoder E_{Fusion} to take the updated BEV feature, denoted as \hat{E}_{Fusion} .

The trainable soft prompts have the same spatial dimension as the feature space. The channel size C_p^0 is a hyperparameter depending on the task complexity. During the training process, the weights of P_S are relaxed, which will learn the task-specific knowledge. The detailed formulation is shown in Eq. 3, where \oplus is the concatenation operation.

Multi-level prompts. Multi-level prompts are added at different levels of the detector. We proposed to add two levels of soft prompts, at F_{BEV} and $F_{BEVFused}$. The dimensions of F_{BEV} and $F_{BEVFused}$ are both $C \times H \times W$. The detailed formulation is summarized in Eq. 3, where \oplus is the concatenation operation.

$$\mathbf{F}_{BEVFused} = \hat{E}_{Fusion} (\mathbf{F}_{BEV} \oplus P_S)$$

$$(\mathbf{F}_{BEVfeat.}^i) = B_{BEV} (Conv (\mathbf{F}_{BEVFused} \oplus P_{S2}))$$
(3)

For the *first level prompts*, we use the same setting as the single level prompts. For the *second level* at $F_{BEVFused}$, the prompts, P_{S2} , are concatenated to the fused BEV features. Then a single layer CNN, $Conv(\cdot)$ is added to align the dimension for original BEV backbone networks. In summary, for the first level soft prompts, we modify the following layers to adapt to the changed channel size, while for second level, a small convolutional layer is added to compress the additional channels.

4. Experiment

Using MSMDFusion [21] as our baseline, we carry out extensive experiments to compare the performance of our proposed **PF3Det** and the baseline.

4.1. Dataset

The nuScenes dataset [5] is a comprehensive, large-scale dataset for advancing autonomous driving research. It provides 1000 driving scenes of high-quality sensor data collected from Boston, Pittsburgh, Las Vegas, and Singapore. These scenes are captured in a variety of conditions, including different times of day, varying weather scenarios, and diverse traffic situations. They are divided into training, validation, and testing sets, with 700, 150, and 150 scenes, respectively. Each scene is 20 seconds, with each frame containing one point cloud and six calibrated images.

The entire dataset has around 390k LiDAR point clouds and 1.4M images, including approximately 1.4M bounding boxes of 23 ground truth categories. The 3D detection performance for nuScenes is measured in mean Average Precision (mAP), and nuScenes detection score (NDS), which is a combination of mAP and other metrics including the quality of box location, size, orientation, *etc.* Following MSMDFusion [21], we set the voxel size to (0.075m, 0.075m, 0.2m). To evaluate the model performance under limited amount of training data, we randomly sample 5% from the original nuScenes training set, and keep all data in the validation and testing sets.

4.2. Implementation details

For the full version of **PF3Det**, we use the VoxelNet [64] as the LiDAR backbone. For the image branch, the original image backbone is a ResNet-50 [14] with FPN [33] pre-trained on ImageNet segmentation task [8] in a Mask-RCNN segmentation network [15]. We choose the image encoder with ViT-L [10] as the backbone from CLIP [43] for the foundational image encoder. The original backbone $L_i(\cdot)$ takes 448 × 800 images as input and produces four feature maps at different scales. The foundation encoder $F(\cdot)$ takes 224 × 224 images and produce a feature vector of size C_{ffeat} . The foundational feature wap of the original backbone and concatenated to it.

For prompts, our final **PF3Det** uses 2-level prompts. Two sets of prompt with size [100, 180, 180] and [150, 180, 180] are concatenated to the BEV features and fused BEV features, respectively.

Following the MSMDFusion [21] training strategy, at stage 1, we train the LiDAR branch for 20 epochs using the Transfusion [3]. At stage 2, the foundation model assisted model and the prompted model are trained separately with the LiDAR model for 6 epochs, using a learning rate of 10^{-4} . At stage 3, the weights from the two multi-modal detectors are integrated, and we continue to train another 6 epochs with a learning rate of 10^{-5} . At the inference time, Test-Time Augmentation (TTA) or multi-model ensemble are not used for efficiency.

4.3. Main results

Table 1 shows the results of our main experiments. We first test different settings for each modules individually. The best settings are combined to produce the final **PF3Det**.

From Table 1. We have three conclusions. (1) From Exp. ID 1, 3 and 5 results, the proposed two modules, foundation branch and multi-modal soft-prompt adapter, each can improve visual LiDAR detector performance individually. Furthermore, they can work jointly and further enhance the overall performance from Exp. 6. (2) From the experiments of the foundation features (Exp. ID 2 and 3), we can conclude that the output dimension of the foundation encoders has to suit the original detector. Otherwise, it might reduce the detection performance. CLIP [47] ResNet50 encoder outputs the feature vectors of length 1024, but CLIP ViT-L produces the features of length 768. The ResNet50 feature vectors have larger feature dimensions, which dilute the original information. (3) The prompts require careful design to balance the parameter increase and NDS/mAP score boost. One layer can achieve the majority of performance gain, +0.67% NDS. Adding a second layer results in $2.5\times$ the number of prompt weights, but it only gives a 0.1% NDS gain as shown in Exp. ID=5. More detailed analysis of each module and results are shown in the following subsections.

Exp. ID	Method	Concat-FM	FM Backbone	Prompt Tuning	Prompt Tuning Layers	NDS	mAP
1	MSMDFusion [21]	×	N/A	×	N/A	69.23	64.50
2	PF3Det	\checkmark	ResNet50 [14]	×	N/A	68.36	62.86
3	PF3Det	\checkmark	ViT-L [10]	×	N/A	70.21	66.15
4	PF3Det	×	N/A	\checkmark	1	69.90	65.66
5	PF3Det	×	N/A	\checkmark	2	70.00	66.69
6	PF3Det	\checkmark	ViT-L [10]	\checkmark	2	70.42	66.92

Table 1. Overall experimental results on nuScenes validation dataset. MSMDFusion results are reproduced by training on 5% of nuScenes dataset. When we experiment with the foundation model (FM), we use CLIP [47] with either ResNet50 or ViT-L as its backbone. Acronyms: Exp.: experiment; Concat-FM: concatenation with FM features.



Figure 3. Foundational Point Encoder. The point features are upsampled to match original feature dimensions.

4.4. Foundation feature assisted branch

We run foundation feature encoders from different multimodal foundation models. The results are summarized in Table 2.

Foundation image encoder. As described in the previous section, foundation image features are added to the image branch as shown in Fig. 2. We tested the two backbones provided by CLIP. The larger feature vector from ResNet50 [14] dilutes the original image features and decreases the performance.

Foundation point encoder. We take the point encoder provided by ULIP [53], which is able to train a model that can encode three modalities into a unified feature space. The point encoder is based on the PointBERT [56]. Since the features before BEV features in the LiDAR branch are sparse, the foundation point features are integrated to multimodal BEV features, as shown in Fig. 3. To integrate foundation point feature (fp feat.) vectors $[C_{fpfeat.}, 1, 1]$ to BEV features $[C_{BEV feat.}, W, H]$, we try to 1) directly repeat vector values to extend to size $[C_{fpfeat.}, W, H]$, 2) apply transpose convolution to upsample the vectors.

original BEV features have a channel size of 256, which is the same channel size as the foundation feature vectors. To avoid overwhelming the original information, we add a channel compression layer before concatenation to reduce the foundation point features to size $[C_{comp.fpfeat.}, W, H]$. The ablation study is shown in Table 2. We have three conclusions: 1) The foundation image features work better compared to the point features. 2) The upsampling layers work better compared to repeating feature vectors. 3) $C_{comp.fpfeat.} = 50$ has better performance. The reason for conclusion 1 is the domain gap between the clean indoor point clouds and outdoor LiDAR point clouds. The point training data used for ULIP [53] is from ShapeNet [6], which is clean and evenly sampled from the object surface compared to the LiDAR points which are highly sparse and uneven.

4.5. Multi-modal soft-prompt adapter

We run experiments to search for the optimal prompt structures. The soft prompts are added after the stage of multimodal BEV features. This allow maximum effect on both the camera and LiDAR modalities. Most of the soft prompts are added to the attention-based networks [20], allowing easy additional prompt patch. However, the BEV feature part of our 3D detector does not have attention-based structures. We design soft prompts specifically for the convolution-based networks. The visual LiDAR 3D detection is a complicated task, and the difficulty to fuse the two modalities is non-trivial. Thus we make all the layers after F_{BEV} learnable and the layers before it to be frozen, as shown in Fig. 2.

Single-level prompts. First we run experiment to verify the effectiveness of the soft prompts on our foundational encoder assisted networks. Only one layer of soft prompts P_S are added to the feature layer of multi-modal BEV features F_{BEV} . The structure is shown in Fig. 4 without other prompts (P_{S2} , P_{S3} , P_{S4}).

The multi-modal BEV features F_{BEV} have the dimen-

Foundation Model	Modality	Foundation Model Backbone	Channel Size	Upsample	NDS	mAP
CLIP [43]	image	ViT-L [10]	768	×	70.21	66.15
CLIP [43]	image	ResNet50 [14]	1024	×	68.36	62.86
ULIP [53]	point cloud	PointBERT [56]	50	×	68.42	64.33
ULIP [53]	point cloud	PointBERT [56]	100	×	67.89	63.66
ULIP [53]	point cloud	PointBERT [56]	50	\checkmark	69.50	65.53
ULIP [53]	point cloud	PointBERT [56]	100	\checkmark	69.35	65.74

Table 2. Ablation Study with different foundation feature (FM) modality encoders. The column *Channel Size* stands for foundation feature channel size. The column *Upsample* indicates whether upsampling convolutions are used.



Figure 4. Multi-level multi-modal soft-prompt adapter. Three levels with four sets of soft prompts are tested. And weights after the first prompts are set to be learnable to better fuse features and prompts.

sion of [256, 180, 180], where 256 is the BEV feature channel size, and 180 is the BEV width and height. The dimension of the single-level prompt P_S is set to $[C_{P_S}, 180, 180]$ to match the width and height dimension of F_{BEV} . Then P_S is concatenated to F_{BEV} to produce a tensor of size [256 + $C_{P_S}, 180, 180$]. The next network layer is a BEV fusion network, as shown in Fig. 4. Their input dimensions are adjusted accordingly.

The results of different prompt channel sizes are shown in Table 3. The performance of the detector increases as C_{P_S} increases when $C_{P_S} <= 100$. The best performance is achieved when the single prompt channel size $C_{P_S} =$ 100. As C_{P_S} increases further, the added prompts start to overwhelm the original BEV features, which has a size of 256. The detector performance saturates and starts to drop, but we observe that the over-sized prompts do not degrade the performance significantly.

Multi-level prompts. From the prior experiments, we verify that using soft prompts can improve the visual LiDAR 3D detection performance. Further explorations were carried out for multi-level prompts. In total, three levels with four sets of soft prompts, denoted as P_S , P_{S2} , P_{S3}^1 , and P_{S3}^2 , were tested. We show the full architecture in Fig. 4.

For P_S , we follow the single-level prompt setup. The second level soft prompts, P_{S2} , is also set to match the

width and height dimension of $F_{BEVFused}$. Then it is concatenated to the fused BEV features of size [256, 180, 180], to create the combined size of [256 + $C_{P_{S2}}$, 180, 180]. The channel dimension is compressed to the original 256 by a channel align layer $Conv(\cdot)$, instead of modifying the following BEV backbone layers. The third level of prompts P_{S3}^i are added to the backbone BEV features, $F_{BEVfeat.}^i$, after the BEV backbone. As the other prompts, they are concatenated to their corresponding multi-scale BEV features to create features of size $[128 + C_{P_{S3}}^1, 180, 180]$ and $[256 + C_{P_{S3}}^2, 90, 90]$. Like the single level prompts, a modified FPN $\hat{P}_{BEV}(\cdot)$ is used to adapt the changed dimensions. The entire structure shown in Fig. 4 is set to be trainable.

The experimental results are shown in Table 3. We run the multi-level prompts with two different settings. In the *multi-level prompts reloaded* section, the models first load the weights from the trained model corresponding to singlelevel prompts. Only the second level prompts are randomly initialized. The best performance is achieve at $C_{P_{S2}} = 50$. When the channel size of P_{S2} increases to 150, the performance sees a slight decrease, but remains at the same level. The first significant drop occurs when the channel size increases to 200, but compared to the single level prompts, they all have slight performance drop.

Another set of experiments in *multi-level prompts* loads the TransFusion [3] pre-trained weights as MSMDFusion

Method	Channel	NDS	mAP
MSMDFusion [21]	-	69.23	64.50
	10	69.22	64.63
	50	69.68	65.27
PF3Det	100	69.90	65.66
Single-Level	150	69.47	65.16
Frompts	200	69.70	65.37
	500	69.79	65.53
	1000	69.83	65.61
	[100, 200, -, -]	69.68	65.94
PF3Det Multi-Level	[100, 150, -, -]	69.76	65.89
Prompts Reloaded	[100, 100, -, -]	69.76	65.88
	[100, 50, -, -]	69.87	66.14
	[100, 150, -, -]	70.00	66.69
PF3Det	[100, 100, -, -]	68.97	64.50
Multi-Level Prompts	[100, 50, -, -]	69.38	64.94
Tompo	[150, 75, 38, 75]	67.66	64.20
	[50, 25, 12, 25]	67.45	64.07

Table 3. **Detection metrics on different prompt setups.** The column *Channel* indicates the number of channels for each soft prompts. Multi-level prompts reloaded here means the final model is first loaded from a pre-trained single level prompted model and then goes through the training process. Other models in the table are trained from TransFusion [3] pre-trained weights. The results are reported in %.

[21]. So all prompts are randomly initialized. For two level prompts, the best combination is when $C_{P_S} = 100$ and $C_{P_{S2}} = 150$. Further experiments were carried out to test four levels of prompts. But they all decrease the detection performance significantly by about 1.8%. One of the channel sizes, [50, 25, 12, 25], has less prompt weights in total, compared to the best prompt channel sizes [100, 150, -, -], so the four prompts are not overwhelming the original detector. The decrease indicates that the prompts at later layers, *e.g.*, multi-scale BEV feature layers, affect the detection performance negatively.

5. Conclusion

In this paper, we proposed **PF3Det**, a novel foundation feature assisted, prompted visual LiDAR 3D detector. Our proposed modules integrate foundation features learned from multi-modality pre-training and further bridge domain and modality gaps by inserting soft prompts. Extensive experiments support that our **PF3Det** achieves the state-of-the-art results under limited amount of training data.

References

- Josh Achiam, Steven Adler, Sandhini Agarwal, Lama Ahmad, Ilge Akkaya, Florencia Leoni Aleman, Diogo Almeida, Janko Altenschmidt, Sam Altman, Shyamal Anadkat, et al. Gpt-4 technical report. arXiv preprint arXiv:2303.08774, 2023. 2
- [2] Muhammad Awais, Muzammal Naseer, Salman Khan, Rao Muhammad Anwer, Hisham Cholakkal, Mubarak Shah, Ming-Hsuan Yang, and Fahad Shahbaz Khan. Foundational models defining a new era in vision: A survey and outlook. arXiv preprint arXiv:2307.13721, 2023. 2
- [3] Xuyang Bai, Zeyu Hu, Xinge Zhu, Qingqiu Huang, Yilun Chen, Hongbo Fu, and Chiew-Lan Tai. Transfusion: Robust lidar-camera fusion for 3d object detection with transformers. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 1090–1099, 2022. 1, 2, 3, 5, 7, 8
- [4] Tom B Brown. Language models are few-shot learners. arXiv preprint arXiv:2005.14165, 2020. 2, 3
- [5] Holger Caesar, Varun Bankiti, Alex H Lang, Sourabh Vora, Venice Erin Liong, Qiang Xu, Anush Krishnan, Yu Pan, Giancarlo Baldan, and Oscar Beijbom. nuscenes: A multimodal dataset for autonomous driving. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 11621–11631, 2020. 2, 5
- [6] Angel X Chang, Thomas Funkhouser, Leonidas Guibas, Pat Hanrahan, Qixing Huang, Zimo Li, Silvio Savarese, Manolis Savva, Shuran Song, Hao Su, et al. Shapenet: An information-rich 3d model repository. arXiv preprint arXiv:1512.03012, 2015. 6
- [7] Xiaozhi Chen, Huimin Ma, Ji Wan, Bo Li, and Tian Xia. Multi-view 3d object detection network for autonomous driving. In *Proceedings of the IEEE conference on Computer Vision and Pattern Recognition*, pages 1907–1915, 2017. 2
- [8] Jia Deng, Wei Dong, Richard Socher, Li-Jia Li, Kai Li, and Li Fei-Fei. Imagenet: A large-scale hierarchical image database. In 2009 IEEE conference on computer vision and pattern recognition, pages 248–255. Ieee, 2009. 5
- [9] Qingxiu Dong, Lei Li, Damai Dai, Ce Zheng, Zhiyong Wu, Baobao Chang, Xu Sun, Jingjing Xu, and Zhifang Sui. A survey on in-context learning. arXiv preprint arXiv:2301.00234, 2022. 3
- [10] Alexey Dosovitskiy. An image is worth 16x16 words: Transformers for image recognition at scale. arXiv preprint arXiv:2010.11929, 2020. 2, 3, 5, 6, 7
- [11] Xinxin Du, Marcelo H Ang, Sertac Karaman, and Daniela Rus. A general pipeline for 3d detection of vehicles. In 2018 IEEE International Conference on Robotics and Automation (ICRA), pages 3194–3200. IEEE, 2018. 2
- [12] Biao Gao, Yancheng Pan, Chengkun Li, Sibo Geng, and Huijing Zhao. Are we hungry for 3d lidar data for semantic segmentation? a survey of datasets and methods. *IEEE Trans-*

actions on Intelligent Transportation Systems, 23(7):6063–6081, 2021. 1, 2

- [13] Jindong Gu, Zhen Han, Shuo Chen, Ahmad Beirami, Bailan He, Gengyuan Zhang, Ruotong Liao, Yao Qin, Volker Tresp, and Philip Torr. A systematic survey of prompt engineering on vision-language foundation models. *arXiv preprint arXiv:2307.12980*, 2023. 3
- [14] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 770–778, 2016. 5, 6, 7
- [15] Kaiming He, Georgia Gkioxari, Piotr Dollár, and Ross Girshick. Mask r-cnn. In Proceedings of the IEEE international conference on computer vision, pages 2961–2969, 2017. 5
- [16] Yining Hong, Haoyu Zhen, Peihao Chen, Shuhong Zheng, Yilun Du, Zhenfang Chen, and Chuang Gan. 3d-llm: Injecting the 3d world into large language models. *Advances* in Neural Information Processing Systems, 36:20482–20494, 2023. 2, 3
- [17] Edward J Hu, Yelong Shen, Phillip Wallis, Zeyuan Allen-Zhu, Yuanzhi Li, Shean Wang, Lu Wang, and Weizhu Chen. Lora: Low-rank adaptation of large language models. arXiv preprint arXiv:2106.09685, 2021. 3
- [18] Rongjie Huang, Mingze Li, Dongchao Yang, Jiatong Shi, Xuankai Chang, Zhenhui Ye, Yuning Wu, Zhiqing Hong, Jiawei Huang, Jinglin Liu, et al. Audiogpt: Understanding and generating speech, music, sound, and talking head. arXiv preprint arXiv:2304.12995, 2023. 3
- [19] Chao Jia, Yinfei Yang, Ye Xia, Yi-Ting Chen, Zarana Parekh, Hieu Pham, Quoc Le, Yun-Hsuan Sung, Zhen Li, and Tom Duerig. Scaling up visual and vision-language representation learning with noisy text supervision. In *International conference on machine learning*, pages 4904–4916. PMLR, 2021. 2, 3
- [20] Menglin Jia, Luming Tang, Bor-Chun Chen, Claire Cardie, Serge Belongie, Bharath Hariharan, and Ser-Nam Lim. Visual prompt tuning. In *European Conference on Computer Vision*, pages 709–727. Springer, 2022. 2, 3, 6
- [21] Yang Jiao, Zequn Jie, Shaoxiang Chen, Jingjing Chen, Lin Ma, and Yu-Gang Jiang. 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, pages 21643– 21652, 2023. 1, 2, 3, 5, 6, 8
- [22] Yecheol Kim, Konyul Park, Minwook Kim, Dongsuk Kum, and Jun Won Choi. 3d dual-fusion: Dual-domain dual-query camera-lidar fusion for 3d object detection. *arXiv preprint arXiv:2211.13529*, 2022. 2
- [23] Alexander Kirillov, Eric Mintun, Nikhila Ravi, Hanzi Mao, Chloe Rolland, Laura Gustafson, Tete Xiao, Spencer Whitehead, Alexander C Berg, Wan-Yen Lo, et al. Segment anything. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 4015–4026, 2023. 2
- [24] Lingdong Kong, Xiang Xu, Jiawei Ren, Wenwei Zhang, Liang Pan, Kai Chen, Wei Tsang Ooi, and Ziwei Liu. Multimodal data-efficient 3d scene understanding for autonomous driving. arXiv preprint arXiv:2405.05258, 2024. 1

- [25] Jason Ku, Melissa Mozifian, Jungwook Lee, Ali Harakeh, and Steven L Waslander. Joint 3d proposal generation and object detection from view aggregation. In 2018 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS), pages 1–8. IEEE, 2018. 2
- [26] Brian Lester, Rami Al-Rfou, and Noah Constant. The power of scale for parameter-efficient prompt tuning. *arXiv preprint arXiv:2104.08691*, 2021. 3
- [27] Kaidong Li, Tianxiao Zhang, Cuncong Zhong, Ziming Zhang, and Guanghui Wang. Robust 3d point clouds classification based on declarative defenders. *Neural Computing and Applications*, pages 1–13, 2024. 3
- [28] Liunian Harold Li, Mark Yatskar, Da Yin, Cho-Jui Hsieh, and Kai-Wei Chang. Visualbert: A simple and performant baseline for vision and language. arXiv preprint arXiv:1908.03557, 2019. 3
- [29] Yingwei Li, Adams Wei Yu, Tianjian Meng, Ben Caine, Jiquan Ngiam, Daiyi Peng, Junyang Shen, Yifeng Lu, Denny Zhou, Quoc V Le, et al. Deepfusion: Lidar-camera deep fusion for multi-modal 3d object detection. In *Proceedings* of the IEEE/CVF conference on computer vision and pattern recognition, pages 17182–17191, 2022. 2
- [30] Ming Liang, Bin Yang, Shenlong Wang, and Raquel Urtasun. Deep continuous fusion for multi-sensor 3d object detection. In Proceedings of the European conference on computer vision (ECCV), pages 641–656, 2018. 1
- [31] Ming Liang, Bin Yang, Yun Chen, Rui Hu, and Raquel Urtasun. Multi-task multi-sensor fusion for 3d object detection. In Proceedings of the IEEE/CVF conference on computer vision and pattern recognition, pages 7345–7353, 2019.
- [32] Tingting Liang, Hongwei Xie, Kaicheng Yu, Zhongyu Xia, Zhiwei Lin, Yongtao Wang, Tao Tang, Bing Wang, and Zhi Tang. Bevfusion: A simple and robust lidar-camera fusion framework. Advances in Neural Information Processing Systems, 35:10421–10434, 2022. 1, 2
- [33] Tsung-Yi Lin, Piotr Dollár, Ross Girshick, Kaiming He, Bharath Hariharan, and Serge Belongie. Feature pyramid networks for object detection. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 2117–2125, 2017. 4, 5
- [34] Minghua Liu, Yin Zhou, Charles R Qi, Boqing Gong, Hao Su, and Dragomir Anguelov. Less: Label-efficient semantic segmentation for lidar point clouds. In *European conference* on computer vision, pages 70–89. Springer, 2022. 1
- [35] Xiao Liu, Yanan Zheng, Zhengxiao Du, Ming Ding, Yujie Qian, Zhilin Yang, and Jie Tang. Gpt understands, too. AI Open, 2023. 2, 3
- [36] Wenchi Ma, Kaidong Li, and Guanghui Wang. Locationaware box reasoning for anchor-based single-shot object detection. *Ieee Access*, 8:129300–129309, 2020. 1
- [37] Wenchi Ma, Tianxiao Zhang, and Guanghui Wang. Miti-detr: Object detection based on transformers with mitigatory self-attention convergence. arXiv preprint arXiv:2112.13310, 2021. 1
- [38] Jiageng Mao, Shaoshuai Shi, Xiaogang Wang, and Hongsheng Li. 3d object detection for autonomous driving: A comprehensive survey. *International Journal of Computer Vision*, 131(8):1909–1963, 2023. 2

- [39] Su Pang, Daniel Morris, and Hayder Radha. Clocs: Cameralidar object candidates fusion for 3d object detection. In 2020 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS), pages 10386–10393. IEEE, 2020. 2
- [40] Lant Pritchett and Justin Sandefur. Learning from experiments when context matters. *American Economic Review*, 105(5):471–475, 2015. 3
- [41] Charles R Qi, Wei Liu, Chenxia Wu, Hao Su, and Leonidas J Guibas. Frustum pointnets for 3d object detection from rgbd data. In Proceedings of the IEEE conference on computer vision and pattern recognition, pages 918–927, 2018. 2
- [42] Alec Radford, Jeffrey Wu, Rewon Child, David Luan, Dario Amodei, Ilya Sutskever, et al. Language models are unsupervised multitask learners. *OpenAI blog*, 1(8):9, 2019. 3
- [43] Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal, Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, et al. Learning transferable visual models from natural language supervision. In *International conference on machine learning*, pages 8748–8763. PMLR, 2021. 2, 3, 5, 7
- [44] Vishwanath A Sindagi, Yin Zhou, and Oncel Tuzel. Mvxnet: Multimodal voxelnet for 3d object detection. In 2019 International Conference on Robotics and Automation (ICRA), pages 7276–7282. IEEE, 2019. 2
- [45] Sourabh Vora, Alex H Lang, Bassam Helou, and Oscar Beijbom. Pointpainting: Sequential fusion for 3d object detection. In Proceedings of the IEEE/CVF conference on computer vision and pattern recognition, pages 4604–4612, 2020. 2
- [46] Tai Wang, ZHU Xinge, Jiangmiao Pang, and Dahua Lin. Probabilistic and geometric depth: Detecting objects in perspective. In *Conference on Robot Learning*, pages 1475– 1485. PMLR, 2022. 1
- [47] Zihao Wang, Wei Liu, Qian He, Xinglong Wu, and Zili Yi. Clip-gen: Language-free training of a text-to-image generator with clip. arXiv preprint arXiv:2203.00386, 2022. 5, 6
- [48] Jason Wei, Xuezhi Wang, Dale Schuurmans, Maarten Bosma, Fei Xia, Ed Chi, Quoc V Le, Denny Zhou, et al. Chain-of-thought prompting elicits reasoning in large language models. Advances in neural information processing systems, 35:24824–24837, 2022. 3
- [49] Danfei Xu, Dragomir Anguelov, and Ashesh Jain. Pointfusion: Deep sensor fusion for 3d bounding box estimation. In Proceedings of the IEEE conference on computer vision and pattern recognition, pages 244–253, 2018. 2
- [50] Lingling Xu, Haoran Xie, Si-Zhao Joe Qin, Xiaohui Tao, and Fu Lee Wang. Parameter-efficient fine-tuning methods for pretrained language models: A critical review and assessment. arXiv preprint arXiv:2312.12148, 2023. 3
- [51] Runsen Xu, Xiaolong Wang, Tai Wang, Yilun Chen, Jiangmiao Pang, and Dahua Lin. Pointllm: Empowering large language models to understand point clouds. arXiv preprint arXiv:2308.16911, 2023. 2, 3
- [52] Xinli Xu, Shaocong Dong, Tingfa Xu, Lihe Ding, Jie Wang, Peng Jiang, Liqiang Song, and Jianan Li. Fusionrcnn: Lidarcamera fusion for two-stage 3d object detection. *Remote Sensing*, 15(7):1839, 2023. 2

- [53] Le Xue, Mingfei Gao, Chen Xing, Roberto Martín-Martín, Jiajun Wu, Caiming Xiong, Ran Xu, Juan Carlos Niebles, and Silvio Savarese. Ulip: Learning a unified representation of language, images, and point clouds for 3d understanding. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 1179–1189, 2023. 2, 3, 6, 7
- [54] Tianwei Yin, Xingyi Zhou, and Philipp Krähenbühl. Multimodal virtual point 3d detection. Advances in Neural Information Processing Systems, 34:16494–16507, 2021. 1
- [55] Jin Hyeok Yoo, Yecheol Kim, Jisong Kim, and Jun Won Choi. 3d-cvf: Generating joint camera and lidar features using cross-view spatial feature fusion for 3d object detection. In Computer vision–ECCV 2020: 16th European conference, Glasgow, UK, August 23–28, 2020, proceedings, part XXVII 16, pages 720–736. Springer, 2020. 2
- [56] Xumin Yu, Lulu Tang, Yongming Rao, Tiejun Huang, Jie Zhou, and Jiwen Lu. Point-bert: Pre-training 3d point cloud transformers with masked point modeling. In *Proceedings* of the IEEE/CVF conference on computer vision and pattern recognition, pages 19313–19322, 2022. 6, 7
- [57] Renrui Zhang, Ziyu Guo, Wei Zhang, Kunchang Li, Xupeng Miao, Bin Cui, Yu Qiao, Peng Gao, and Hongsheng Li. Pointclip: Point cloud understanding by clip. In Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, pages 8552–8562, 2022. 3
- [58] Tianxiao Zhang, Wenchi Ma, and Guanghui Wang. Sixchannel image representation for cross-domain object detection. In *Image and Graphics: 11th International Conference, ICIG 2021, Haikou, China, August 6–8, 2021, Proceedings, Part I 11*, pages 171–184. Springer, 2021. 1
- [59] Tianxiao Zhang, Bo Luo, Ajay Sharda, and Guanghui Wang. Dynamic label assignment for object detection by combining predicted ious and anchor ious. *Journal of Imaging*, 8(7): 193, 2022. 1
- [60] Tianxiao Zhang, Bo Luo, and Guanghui Wang. Improving vision transformers by overlapping heads in multi-head selfattention. arXiv preprint arXiv:2410.14874, 2024. 3
- [61] Tianxiao Zhang, Wenju Xu, Bo Luo, and Guanghui Wang. Depth-wise convolutions in vision transformers for efficient training on small datasets. *Neurocomputing*, 617:128998, 2025. 3
- [62] Yiyuan Zhang, Kaixiong Gong, Kaipeng Zhang, Hongsheng Li, Yu Qiao, Wanli Ouyang, and Xiangyu Yue. Metatransformer: A unified framework for multimodal learning. arXiv preprint arXiv:2307.10802, 2023. 3
- [63] Xin Zhou, Jinghua Hou, Tingting Yao, Dingkang Liang, Zhe Liu, Zhikang Zou, Xiaoqing Ye, Jianwei Cheng, and Xiang Bai. Diffusion-based 3d object detection with random boxes. In *Chinese Conference on Pattern Recognition and Computer Vision (PRCV)*, pages 28–40. Springer, 2023. 2
- [64] Yin Zhou and Oncel Tuzel. Voxelnet: End-to-end learning for point cloud based 3d object detection. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 4490–4499, 2018. 1, 5