

# M3Net: Multimodal Multi-task Learning for 3D Detection, Segmentation, and Occupancy Prediction in Autonomous Driving

Xuesong Chen<sup>1</sup>, Shaoshuai Shi<sup>1\*</sup>, Tao Ma<sup>1</sup>,  
Jingqiu Zhou<sup>1</sup>, Simon See<sup>2</sup>, Ka Chun Cheung<sup>2</sup>, Hongsheng Li<sup>1,3</sup>

<sup>1</sup> MMLab, The Chinese University of HongKong

<sup>2</sup> NVIDIA AI Technology Center

<sup>3</sup> Centre for Perceptual and Interactive Intelligence

{chenxuesong@link, hsl@ee}.cuhk.edu.hk, shaoshuaics@gmail.com

## Abstract

The perception system for autonomous driving generally requires to handle multiple diverse sub-tasks. However, current algorithms typically tackle individual sub-tasks separately, which leads to low efficiency when aiming at obtaining full-perception results. Some multi-task learning methods try to unify multiple tasks with one model, but do not solve the conflicts in multi-task learning. In this paper, we introduce M3Net, a novel multimodal and multi-task network that simultaneously tackles detection, segmentation, and 3D occupancy prediction for autonomous driving and achieves superior performance than single task model. M3Net takes multimodal data as input and multiple tasks via query-token interactions. To enhance the integration of multi-modal features for multi-task learning, we first propose the **Modality-Adaptive Feature Integration (MAFI)** module, which enables single-modality features to predict channel-wise attention weights for their high-performing tasks, respectively. Based on integrated features, we then develop task-specific query initialization strategies to accommodate the needs of detection/segmentation and 3D occupancy prediction. Leveraging the properly initialized queries, a shared decoder transforms queries and BEV features layer-wise, facilitating multi-task learning. Furthermore, we propose a **Task-oriented Channel Scaling (TCS)** module in the decoder to mitigate conflicts between optimizing for different tasks. Additionally, our proposed multi-task querying and TCS module support both Transformer-based decoder and Mamba-based decoder, demonstrating its flexibility to different architectures. M3Net achieves state-of-the-art multi-task learning performance on the nuScenes benchmarks.

**Code** — <https://github.com/Cedarch/M3Net>

## Introduction

The perception module is important for autonomous driving system, which aims to accurately perceive the environment around the self-driving vehicle. To ensure reliable driving, the perception task is divided into different sub-tasks, including detection, map segmentation, occupancy prediction, etc. However, current algorithms generally tackle individual sub-tasks and thus require deploying multiple independent

models to obtain full-perception results. This practice poses notable challenges: 1) It overlooks task similarities, leading to redundant computations, and 2) it impedes the mutual enhancement between the multiple sub-tasks. Therefore, a unified model that can simultaneously tackle multiple perception sub-tasks holds significant promise.

However, integrating multiple tasks on a fully shared backbone often results in low performance. This is because the feature representations required by different tasks vary, leading to gradient conflicts when training the network. Various approaches have been developed to tackle the conflicts in multi-task learning, including shared feature representations (He et al. 2017; Borsa, Graepel, and Shawe-Taylor 2016), task-specific layers (Wallingford et al. 2022), and attention mechanisms (Goncalves et al. 2023; Liu, Johns, and Davison 2019), etc. Shared feature representations involve sharing a backbone to utilize shared knowledge. Some methods (Li et al. 2024; Liu et al. 2023a) utilize the query as a unified representation across tasks to get multi-task results. Task-specific layers equipped with attention mechanisms facilitate task-specific learning and enable the model to handle task-relevant information. Although these strategies perform well in 2D perception tasks, challenges arise in the context of 3D tasks in autonomous driving.

Specifically, tasks of object detection and map segmentation can be effectively tackled using Bird’s Eye View (BEV) features, whereas 3D occupancy estimation requires computational-intensive 3D features for achieving dense 3D prediction. Most previous methods use BEV features to unify the representation of multiple tasks. However, they have difficulties in effectively tackling occupancy prediction and semantic conflicts from the compression of the height dimension in the BEV space. For example, a vehicle may share features with both the trees above it and the road below it. Consequently, the network struggles to learn features to handle multiple classes simultaneously. In response to this challenge, Univision (Hong et al. 2024) maintains both 3D and BEV features and successfully handles detection and occupancy prediction tasks simultaneously, but it comes at the cost of efficiency. Other methods, such as MetaBEV (Ge et al. 2023) and BEVFusion (Liu et al. 2023b), try to handle multi-task learning in one model but overlook the important 3D occupancy prediction task. They only focus on BEV-related tasks, including detection and map segmenta-

\*Corresponding author

Copyright © 2025, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

tion, through task-specific modules such as Mixture of Experts (MoE) and unshared decoders. Despite the efforts, the performance of multi-task learning for autonomous driving perception still falls short of single-task learning.

In this paper, a unified **Multi-Modal** and **Multi-task** learning network, named *M3Net*, is proposed which can handle multiple perception tasks in autonomous driving, including 3D objection detection, BEV map segmentation, and 3D occupancy prediction, and yields superior performance to independently trained single-task models.

M3Net first leverages the well-established backbone of BEVFusion to encode BEV features from camera and LiDAR sensors. Based on the BEV feature, M3Net then unifies multiple tasks into the form of query-token interactions and proposes to utilize 3D and BEV queries in a shared BEV feature decoder for multi-task learning. To accommodate the representation differences between the 3D occupancy prediction and BEV-based detection and segmentation, we propose a task-specific query initialization strategy for different tasks. For detection and segmentation tasks, we choose specific features from the BEV feature map as the initial queries and utilize their positions as the queries' position embedding. For occupancy prediction, each voxel in the 3D space can act as an occupancy query, with its corresponding

Although utilizing the proposed query initialization strategy, a powerful BEV feature that fully exploits the advantages of different modalities is still crucial for multi-task learning. For example, LiDAR point clouds are more suitable for detection tasks due to their precise geometric structure, while images contain more dense semantic information and thus perform better on segmentation tasks. To this end, we propose a **Modality-Adaptive Feature Integration (MAFI)** module. This module enhances the fusion of multi-modal features by allowing single-modality features to predict channel-wise gating weights for their proficient tasks. Specifically, the MAFI module initially fuses two modalities' features using a simple convolution layer, similar to the operation in BEVFusion. Subsequently, LiDAR and image BEV features independently predict lidar-adaptive and camera-adaptive weights to perform transformations on the initial fused features, resulting in modality-adapted multi-modal features respectively. Finally, these adapted features are added to form the final integrated BEV features. Furthermore, the integrated BEV features still face challenges of gradient conflicts between the tasks during optimization. To address this challenge, we propose an effective and efficient **Task-oriented Channel Scaling (TCS)** mechanism, which dynamically predicts channel scaling weights of the shared BEV features for different tasks from the same input. Subsequently, these weights are multiplied with the shared BEV features to derive task-specific BEV features. Through the proposed TCS, the queries of different tasks selectively gather information from the shared BEV feature and effectively mitigate conflicts between multiple tasks. Following several decoder layers, each task's queries are fed into different task heads. With our MAFI and TCS modules, M3Net better integrates multi-modal features and tackles multiple perception tasks in a lightweight manner.

Our proposed M3Net supports both transformer-based

and mamba-based decoders, featuring an attention mechanism with linear complexity, which demonstrates M3Net's flexibility and the effectiveness of the mamba mechanism in autonomous driving perception tasks.

In a nutshell, our contributions are four-fold: 1) We introduce a unified, multi-modal, and multi-task learning network M3Net, which can simultaneously tackle multiple perception tasks for autonomous driving. 2) We propose a simple and effective modality-adaptive multi-modal BEV feature integration module to effectively leverage the strengths of different modality features. 3) We propose the task-oriented channel scaling mechanism to alleviate gradient conflicts between different tasks during optimization, allowing M3Net to obtain higher performances than single-task models. 4) Our framework is compatible with both transformer and mamba decoders, demonstrating its flexibility and mamba decoder's capability in 3D perception tasks.

## Related Work

### 3D Object Detection and Semantic Map Segmentation.

3D detection involves locating and classifying objects in the 3D world, while map segmentation entails classifying regions within a map, such as drive areas or pedestrian crossing lines. Methods in this field can be categorized into single-modal and multi-modal approaches based on their input data. For single-modal methods, BEVDet (Huang et al. 2021), BEVFormer (Li et al. 2022a) and BEVsegformer (Peng et al. 2023) transform images from surround-view cameras into a single BEV feature map for 3D detection and segmentation. LiDAR-based methods (Shi et al. 2020; Chen et al. 2022, 2023) mainly focus on 3D detection from single-frame or multi-frame point clouds. Multi-modal approaches (Liu et al. 2023b; Ge et al. 2023) utilize inputs from multiple sensors and fuse them in BEV space, adopting strategies similar to LiDAR-based methods for 3D detection and map segmentation.

**3D Occupancy Prediction.** Occupancy prediction involves making judgments about whether a position in a 3D scene is occupied and identifying the occupation category. Early methods in this domain, such as (Song et al. 2017; Cai et al. 2021), relied on vision-based approaches utilizing images and additional geometric information to perform occupancy prediction. Recently, TPVFormer (Huang et al. 2023a) improves the representation by introducing the Tri-Perspective-View (TPV), which enriches the BEV with Z-axis information for a more detailed scene understanding. OccFormer (Zhang, Zhu, and Du 2023) proposes a dual-path transformer network tailored to processing 3D volumes, specifically targeting semantic occupancy prediction tasks. VoxFormer (Li et al. 2023) proposes a transformer-based two-stage network, where the first stage focuses on a sparse set of visible and occupied voxels, and the second stage generates dense 3D voxels from the sparse ones. Additionally, OpenOccupancy (Wang et al. 2023), Occ3D (Tian et al. 2024), and SurroundOcc (Wei et al. 2023) have also contributed by devising pipelines dedicated to generating high-fidelity dense occupancy labels.

**Multi-Task Learning.** Multi-task frameworks aim to efficiently manage various tasks within a singular net-

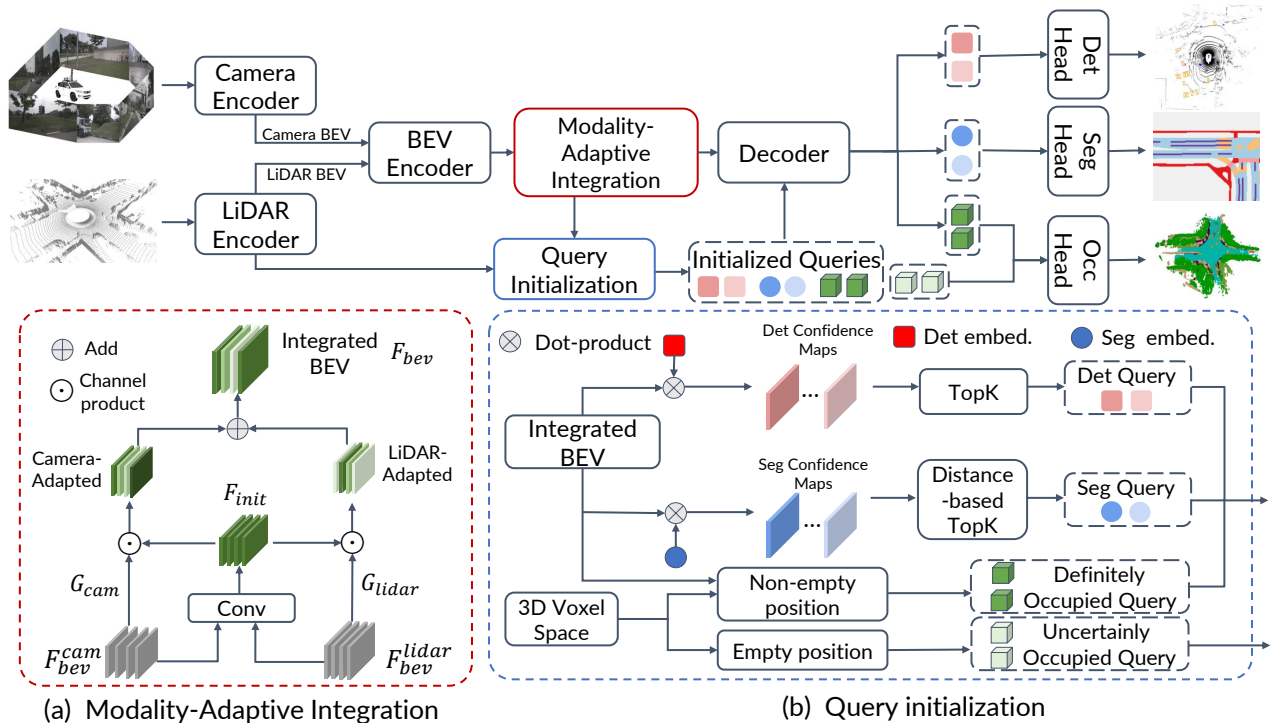


Figure 1: The overall architecture of our proposed M3Net as well as the detailed design of (a) Modality-adaptive Feature Integration module and (b) BEV-based multi-task query initialization.

work. These methods are primarily categorized into two groups (Huang et al. 2023c): optimization techniques (Chen et al. 2018; Mordan et al. 2021; Zhou et al. 2019; Wang et al. 2019) and network architecture enhancement (Ruder et al. 2019; Guo, Lee, and Ulbricht 2020). FULLER (Huang et al. 2023c) proposes a multi-task optimization method to balance the loss weights of different tasks during training, preventing one task from dominating another. On the other hand, LidarMTL (Feng et al. 2021) and LidarMultiNet (Ye et al. 2023) utilize a shared network for tasks encompassing 3D detection, segmentation, and road understanding based on point cloud data. Meanwhile, (Wang et al. 2024) focuses on vision-centric (only using camera feature) multiple perception tasks in autonomous driving. In this paper, M3Net aims to use multimodal data within a unified architecture to address multiple tasks in autonomous driving.

## Methodology

### Overview

M3Net takes LiDAR point cloud and surrounding-view camera images as input and tackles the perception tasks of 3D object detection, BEV map segmentation, and 3D occupancy prediction simultaneously. To tackle multi-task learning, M3Net unifies the multiple perception tasks into the form of query-token interactions and adopts task-specific queries through a shared BEV feature decoder so that different perception tasks can be unified in one model.

As illustrated in Fig. 1, M3Net comprises four sequential sub-networks: a multimodal BEV feature encoder, a modality-adaptive feature integration module, a multi-task query initialization network, and a multi-task decoder.

### Modality-adaptive BEV Feature Integration

We unify both detection, segmentation, and 3D occupancy prediction into the BEV feature space, considering its efficiency, which is different from Univision (Hong et al. 2024) utilizing both BEV and 3D features for multiple tasks. To obtain BEV features, we utilize the feature backbone of BEV Fusion to process inputs from LiDAR and cameras and get corresponding BEV features, i.e.,  $F_{bev}^{lidar}$  and  $F_{bev}^{cam}$ . Subsequently, we introduce the Modality-adaptive Feature Integration (MAFI) module to optimally leverage the strengths of diverse modalities. As illustrated in Fig. 1 (a), the integration module begins with an initial fusion of  $F_{bev}^{lidar}$  and  $F_{bev}^{cam}$  features using a convolution layer, getting  $F_{init}$ . Then, we employ an adaptive gating mechanism to integrate the advantages of LiDAR features in the detection task and the efficacy of image features in the segmentation task. Specifically,  $F_{bev}^{lidar}$  and  $F_{bev}^{cam}$  BEV features predict LiDAR-adaptive gating weights  $G_{lidar}$  and camera-adaptive gating weights  $G_{cam}$  through linear layers followed by a sigmoid activation function, respectively.

These weights enable distinct gating transformations on the initially fused features, resulting in adapted features. The final modality-adapted BEV features are formulated as:

$$F_{bev} = G_{lidar} \odot F_{init} + G_{cam} \odot F_{init}. \quad (1)$$

Because of the better utilization of the advantages of single-modal features, the dense 3D Occupancy prediction task, which combines the information of both foreground detection and background segmentation, can also benefit from our MAFI module. With  $F_{bev}$  in hand, the multi-task queries can be better initialized from the integrated feature.

## BEV-based Multi-task Query Initialization

Query initialization is vital for query-based networks to achieve satisfactory performance. Our experiments show that randomly initialized queries do not carry explicit semantic features and lead to inferior performance. To this end, we propose query initialization strategies based on multimodal BEV features to take advantage of the semantic and location information encoded in the BEV feature map.

**Occupancy query initialization.** We generate a 3D voxel feature space of dimensions  $(H \times W \times Z)$ . Each voxel in this space serves as a query, amounting to a total of  $H \times W \times Z$   $C$ -dimensional occupancy queries. We employ a Multi-Layer Perceptron (MLP) to encode 3D positions into their corresponding  $C$ -dimensional position embeddings. Given that most 3D occupancy queries are unoccupied, we utilize LiDAR point location information to categorize 3D occupancy queries into two groups: definitely occupied (with LiDAR points in the voxel) and uncertainly occupied (without LiDAR points in the voxel). Then different initialization strategies are adopted for them. Specifically, we project the 2D BEV feature to their corresponding 3D positions as the initial feature of the definitely occupied queries. Note that queries at the same 2D position but of varying heights share the same BEV-feature queries but with different positional embeddings. On the other hand, for uncertainly occupied queries, a shared and learnable embedding is employed as their initial feature. At last, we add the queries' initial features and positional embedding to obtain the final initialized occupancy queries.

**Detection and segmentation query initialization.** We choose corresponding features from the integrated multimodal BEV feature map, which are added with positional embeddings, as the initial queries for the detection and segmentation tasks. Such queries are naturally better aligned with the BEV feature map and would be later used in the decoder's cross-attention blocks to perform task-specific transformations. To select better locations to generate queries from the BEV feature map, we employ a unified segmentation approach to predict confidence maps for both object centers and map segmentation. Specifically, we first randomly initialize 16 category embeddings, including 10 detection categories and 6 segmentation categories, and then construct a detection embedding map and a segmentation embedding map. Subsequently, following Mask2Former (Cheng et al. 2021), we apply the dot-product between each category embedding and the corresponding embedding map to obtain the confidence maps. With these confidence maps, for the detection task, we employ  $n_d$  BEV features by selecting top- $n_d$  candidates from the  $H \times W \times 10$  confidence space as detection queries. For the segmentation task, considering the distribution of LiDAR points varies significantly along the forward direction of the vehicle, we introduce a distance-based segmentation query initialization strategy. Specifically, we partition the confidence map evenly into  $S$  blocks (corresponding to different distance ranges) along the forward direction of the vehicle. Within each block, for each segmentation class, the BEV feature with the highest confidence is selected as the query, resulting in a total of  $n_s = S \times 6$  segmentation queries. To this

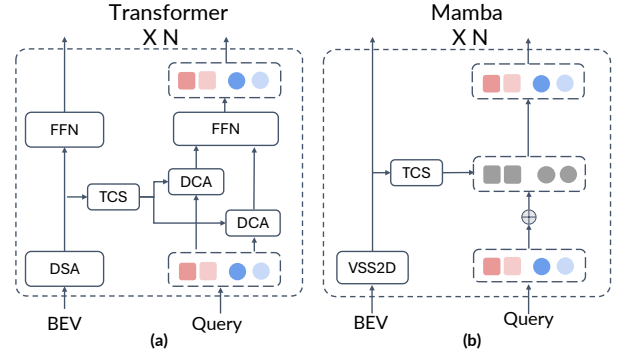


Figure 2: The detailed architectures of our transformer-based and mamba-based decoder layer with the task-oriented channel scaling module. DSA, DCA and VSS2D denote deformable self-attention, cross-attention and the VSS2D block from Vmamba (Liu et al. 2024).

end, queries of different blocks are tasked with segmenting the map within its corresponding distance range.

## Decoder with Task-oriented Channel Scaling

Following query initialization, the decoder iteratively updates the multi-task queries and BEV features in a layer-by-layer manner. Notably, only the definitely occupied occupancy queries, detection queries, and segmentation queries are input into the decoder. The uncertainly occupied occupancy queries subsequently interact with the definitely occupied occupancy queries via attention in the occupancy prediction head, jointly predicting the dense 3D occupancy volume. Our decoder supports both transformer-based and mamba-based attention mechanisms for multi-task learning.

**Transformer-based Decoder.** As shown in Fig. 2 (a), in each layer of the transformer-based decoder, we utilize the deformable self-attention layer (Zhu et al. 2020) to propagate information between all features in the BEV space, so that each BEV feature can obtain long-range semantic relations in the whole context. After the updating of the BEV features, task-specific queries can aggregate richer information from the BEV space through a cross-attention layer. However, conflicts are inevitable when queries attempt to aggregate task-specific information from the shared BEV. To this end, we introduce a Task-specific Channel Scaling (TCS) operation to obtain task-specific BEV features.

**Task-oriented Channel Scaling.** We propose the Task-oriented Channel Scaling (TCS) mechanism to perform channel scaling of the shared BEV features so that task-specific queries can selectively gather different information to alleviate conflicts between multiple tasks. Specifically, taking the BEV feature  $F_{bev}$  derived from an FFN module as input, we employ  $n$  separate Linear layers followed by an activation function to obtain  $n$  task-specific embeddings, where  $n$  is the number of tasks and is 3 in our experiments. Then, these embeddings are forwarded into  $n$  different linear layers to obtain  $n$  task-specific channel scaling weights  $W^i \in \mathbb{R}^{H \times W \times C}$ , where  $i = 1, \dots, n$ . Multiplying the shared BEV features by these scaling weights:

$$F_{bev}^i = W^i \cdot F_{bev}, \text{ for } i = 1, \dots, n. \quad (2)$$

Subsequently, task-specific queries can perform cross-attention from corresponding tasks’  $F_{bev}^i$ .

**Mamba-based Decoder.** The Mamba (Gu and Dao 2023) network introduces State Space Models (SSMs) to manage attention computations with linear complexity, capturing long-range dependencies while preserving computational efficiency. We adopt the design of Vmamba (Liu et al. 2024) to scan each BEV feature from different directions. For each BEV location, we scan information from its surrounding positions in four different directions: top, bottom, left, and right. Subsequently, we merge the features scanned from these directions to update the current position’s feature. Utilizing the scan and merge operation, the Vision Selective Scan 2D (VSS2D) block, proposed in Vmamba, can perform spatial information propagation in BEV features,  $F_{mbev} = \text{VSS2D}(F_{bev})$ . Then, similar to the transformer-based decoder, we use TCS operation to obtain the task-specific BEV feature  $F_{mbev}^i$ . Then multi-task queries can interact with  $F_{mbev}^i$  to update their features, as shown in Fig. 2 (b). Specifically, we adopt the index operation to obtain BEV features from the corresponding positions of different queries to update their feature, expressed as  $q^i = q^i + \text{Index}(F_{mbev}^i)$ , where  $q^i$  means task-specific queries and  $i \in \{1, \dots, n\}$ . After updating queries and BEV features within one layer, we stack several layers to gradually complete multi-task learning for both transformer-based and mamba-based decoders.

### Multi-task Heads and Losses

For the detection task, we use the center-based head (Bai et al. 2022a) and the detection loss is composed of the classification loss  $L_{cls}$  and the regression loss  $L_{reg}$ , denoted as  $L_{det} = \lambda_1 \cdot L_{cls} + \lambda_2 \cdot L_{reg}$ , where  $\lambda_1, \lambda_2$  denote hyper-parameters. For the segmentation task, the same operation for the map segmentation confidence prediction is adopted. We utilize the segmentation-specific BEV feature  $F_{bev}^i$  from the last decoder layer as the embedding and we dot-product this embedding with segmentation queries, obtaining map segmentation results. Following BEVFusion, we use the standard focal loss (Lin et al. 2017) as segmentation loss, denoted as  $L_{seg}$ . For the 3D occupancy prediction task, we first merge the updated definitely occupied queries with uncertainly occupied queries to derive the initial voxel features  $F_{3D} \in \mathbb{R}^{H \times W \times Z \times C}$  to obtain dense 3D voxel features. Subsequently, we employ a transformer or mamba block to refine the voxel features. After obtaining refined voxel features, we upsample it and project it to the output space to get the final results (the same shape with ground truth)  $O \in \mathbb{R}^{H' \times W' \times Z' \times (M+1)}$ , where  $M + 1$  denotes  $M$  semantic classes and one empty class. We follow the loss setting in OpenOccupancy (Wang et al. 2023), denoted as  $L_{occ}$ . The total loss is formulated as:  $L = \lambda_{det} \cdot L_{det} + \lambda_{seg} \cdot L_{seg} + \lambda_{occ} \cdot L_{occ}$ , where  $\lambda_{det}, \lambda_{seg}, \lambda_{occ}$  denote hyper-parameters.

## Experiments

### Dataset and Implementation Details

**Dataset.** For detection and segmentation, we evaluate M3Net with the nuScenes (Caesar et al. 2020) dataset, a

large-scale multimodal dataset designed for 3D detection and map segmentation. It encompasses data from various sensors, including 6 cameras, one LiDAR, and five radars. We resize the input views to  $256 \times 704$  resolution and voxelize the point cloud to 0.075m for detection, occupancy prediction, and segmentation tasks. For 3D detection evaluation, we employ the standard nuScenes Detection Score (NDS) and mean Average Precision (mAP). Additionally, for BEV map segmentation evaluation, we utilize the Mean Intersection over Union (mIoU) metric across all six categories, following the setting of BEVFusion (Liu et al. 2023b). For occupancy prediction, we evaluate M3Net on OpenOccupancy (Wang et al. 2023) dataset. The dataset conducts occupancy annotation on the observed scene using multimodal input from the nuScenes dataset.

**Implementation Details.** The implementation of M3Net is built upon the OpenPCDet framework (Team et al. 2020). For optimization, M3Net employs the AdamW (Loshchilov and Hutter 2017) optimizer and cycle learning rate (lr) schedule with max lr setting to 0.0008. To save training resources, we initially train the detection and segmentation tasks for 12 epochs with CBGS sampling (Zhu et al. 2019) and then integrate the occupancy prediction task for the next 15 epochs’ joint training without CBGS sampling. For hyper-parameters, We use  $n_d = 200$  detection queries,  $n_s = 30$  segmentation queries, and  $180 \times 180 \times 5$  occupancy prediction queries. The channel of all queries is 256. We set the number of decoder layers to 6 and the shape of the BEV feature is  $180 \times 180 \times 256$ . The proposed transformer-based and mamba-based M3Net share the same training setting.

### Performance for Detection and Segmentation

Table 1 shows the detection and segmentation results of state-of-the-art methods on nusenes datasets. We present the performances of M3Net based on both transformer and mamba decoders. For transformer-based M3Net, it shows consistency performance enhancements across detection and segmentation tasks, compared to independently trained BEVFusion (Liu et al. 2023b), achieving improvements of 0.5%, 1.0%, and 7.7% in mAP, NDS, and IoU, respectively, which demonstrates that M3Net better utilizes the information of BEV features. Meanwhile, compared with BEVFusion-MTL (Liu et al. 2023b), with two separate BEV encoders, named BEVFusion-MTL (sep), M3Net significantly exceed it by 3.2%, 2.6%, and 11.9% in mAP, NDS, and IoU, proving the effectiveness of our multi-task learning strategy. Moreover, the substantial enhancement in segmentation tasks further demonstrates the effectiveness of our strategy in unifying detection and segmentation tasks into query representation, contrasting with BEVFusion’s strategy that adopts separate query-based detection and convolution-based segmentation heads. When compared to the state-of-the-art multi-task learning method MetaBEV (Ge et al. 2023), which also uses the backbone of BEVFusion, our approach surpasses its shared encoder version by 3.4%, 2.9%, and 3.1% in mAP, NDS, and IoU metrics, respectively. Furthermore, even when compared to the separate encoder version, our method still outperforms it by 3.6%, 2.6%, and 3.5% in mAP, NDS and IoU metrics. These results demonstrate the effectiveness of our MAFI module in better inte-

Methods	Modality	MTL	mAP(val)	NDS(val)	Drivable	Ped.Cross	Walkway	Stop Line	Carpark	Divider	Mean
M <sup>2</sup> BEV(Xie et al. 2022)	C	✗	41.7	47.0	77.2	-	-	-	-	40.5	-
BEVFormer(Li et al. 2022b)	C	✗	41.6	51.7	80.1	-	-	-	-	25.7	-
BEVFusion(Liu et al. 2023b)	C	✗	35.6	41.2	81.7	54.8	58.4	47.4	50.7	46.4	56.6
X-Align(Borse et al. 2023)	C	✗	-	-	82.4	55.6	59.3	49.6	53.8	47.4	58.0
PointPillars(Lang et al. 2019)	L	✗	52.3	61.3	72.0	43.1	53.1	29.7	27.7	37.5	43.8
CenterPoint(Yin, Zhou, and Krahenbuhl 2021)	L	✗	59.6	66.8	75.6	48.4	57.5	36.5	31.7	41.9	48.6
BEVFusion(Liu et al. 2023b)	L	✗	64.7	69.3	75.6	48.4	57.5	36.4	31.7	41.9	48.6
PointPainting(Vora et al. 2020)	L+C	✗	65.8	69.6	75.9	48.5	57.1	36.9	34.5	41.9	49.1
MVP(Yin, Zhou, and Krähenbühl 2021)	L+C	✗	66.1	70.0	76.1	48.7	57.0	36.9	33.0	42.2	49.0
TransFusion(Bai et al. 2022b)	L+C	✗	67.3	71.2	-	-	-	-	-	-	-
FULLER (Huang et al. 2023c)	L+C	✗	67.6	71.3	-	-	-	-	-	-	62.3
BEVFusion(Liu et al. 2023b)	L+C	✗	68.5	71.4	85.5	60.5	67.6	52.0	57.0	53.7	62.7
X-Align(Borse et al. 2023)	L+C	✗	-	-	86.8	65.2	70.0	58.3	57.1	58.2	65.7
FULLER (Huang et al. 2023c)	L+C	✓	60.5	65.3	-	-	-	-	-	-	58.4
BEVFusion-MTL(Liu et al. 2023b) (share)	L+C	✓	-	69.7	-	-	-	-	-	-	54.0
BEVFusion-MTL(Liu et al. 2023b) (sep)	L+C	✓	65.8	69.8	83.9	55.7	63.8	43.4	54.8	49.6	58.5
MetaBEV (Ge et al. 2023)(share)	L+C	✓	65.6	69.5	88.7	64.8	71.5	56.1	58.7	58.1	66.3
MetaBEV (Ge et al. 2023) (sep)	L+C	✓	65.4	69.8	88.5	64.9	71.8	56.7	61.1	58.2	66.9
M3Net (mamba)	L+C	✓	68.5	71.8	<b>90.6</b>	69.1	<b>75.8</b>	62.5	<b>65.3</b>	<b>61.7</b>	<b>70.8</b>
M3Net (transformer)	L+C	✓	<b>69.0</b>	<b>72.4</b>	90.3	<b>69.6</b>	75.8	<b>63.4</b>	62.3	61.1	70.4

Table 1: **Performance comparisons on nuScenes val set.** “share” means multi-task heads share one BEV encoder and “sep” means task heads have separate encoders. M3Net is built upon the backbone of BEVFusion.

Method	Modality			mIoU	Occupancy															
	MTL	Input	Surround		barrier	bicycle	bus	car	const. veh.	motorcycle	pedestrian	traffic cone	trailer	truck	drive. suf.	other flat	sidewalk	terrain	manmade	vegetation
MonoScene (Cao and de Charette 2022)	✗	C	✗	6.9	7.1	3.9	9.3	7.2	5.6	3.0	5.9	4.4	4.9	4.2	14.9	6.3	7.9	7.4	10.0	7.6
TPVFormer (Huang et al. 2023b)	✗	C	✓	7.8	9.3	4.1	11.3	10.1	5.2	4.3	5.9	5.3	6.8	6.5	13.6	9.0	8.3	8.0	9.2	8.2
AICNet (Li et al. 2020)	✗	C&D	✗	10.6	11.5	4.0	11.8	12.3	5.1	3.8	6.2	6.0	8.2	7.5	24.1	13.0	12.8	11.5	11.6	20.2
3DSKetch (Chen et al. 2020)	✗	C&D	✗	10.7	12.0	5.1	10.7	12.4	6.5	4.0	5.0	6.3	8.0	7.2	21.8	14.8	13.0	11.8	12.0	21.2
LMSCNet (Roldao, de Charette, and Verroust-Blondet 2020)	✗	L	✓	11.5	12.4	4.2	12.8	12.1	6.2	4.7	6.2	6.3	8.8	7.2	24.2	12.3	16.6	14.1	13.9	22.2
JS3C-Net (Yan et al. 2021)	✗	L	✓	12.5	14.2	3.4	13.6	12.0	7.2	4.3	7.3	6.8	9.2	9.1	27.9	15.3	14.9	16.2	14.0	24.9
CONet (Wang et al. 2023) (LR)	✗	C&L	✓	15.1	14.3	12.0	15.2	14.9	13.7	15.0	13.1	9.0	10.0	14.5	23.2	17.5	16.1	17.2	15.3	19.5
CONet (Wang et al. 2023) (HR)	✗	C&L	✓	20.1	23.3	13.3	21.2	24.3	15.3	15.9	18.0	13.3	15.3	20.7	<b>33.2</b>	21.0	22.5	21.5	19.6	23.2
M3Net (transformer)	✓	C&L	✓	23.3	27.5	19.3	<b>23.6</b>	27.0	<b>15.5</b>	28.5	34.1	19.1	17.7	23.0	27.8	16.1	22.0	21.0	24.1	26.5
M3Net (mamba)	✓	C&L	✓	<b>24.1</b>	<b>28.7</b>	<b>20.1</b>	23.1	<b>27.3</b>	15.1	<b>28.6</b>	<b>34.6</b>	<b>20.5</b>	<b>18.0</b>	<b>23.0</b>	28.3	<b>21.7</b>	<b>22.5</b>	<b>22.0</b>	<b>24.7</b>	<b>27.2</b>

Table 2: **Performance comparison on nuScenes-Occupancy val set.** The *C, D, L* denotes camera, depth, LiDAR and *MTL* means Multi-Task Learning. *LR* means low resolution feature and *HR* denotes high resolution feature.

grating multimodal features and our TCS operation in mitigating conflicts of multi-task learning, outperforming the MoE strategy utilized by MetaBEV.

For mamba-based M3Net, it shows similar performance to our transformer model and surpasses the MetaBEV (sep), by 3.1%, 2.0%, and 3.9% in terms of mAP, NDS, and IoU, demonstrating the flexibility of M3Net to different attention mechanisms. Meanwhile, the mamba model has fewer parameters than the transformer-based one by removing FFN.

### Performance for 3D Occupancy Prediction

Table 2 shows the results of 3D occupancy prediction of different methods on the nuScenes-based OpenOccupancy (Wang et al. 2023) dataset. Our transformer-based and mamba-based M3Net significantly surpass the state-of-the-art method CONet (HR) by 3.2% and 4.0% respectively in terms of mIoU. Thanks to the joint training with the detection task, M3Net has achieved great improvements in occupancy prediction of foreground objects, especially small objects. For example, mamba-based M3Net significantly surpasses CONet (HR) by about 12% and 16% on motorcycle and pedestrian, respectively, proving the superiority of multi-task learning. On the other hand, the resolution of our adopted 3D feature space is much lower than that of CONet (HR), which utilizes the resolution of  $512 \times 512 \times 40$  (the

same as ground truth’s resolution). However, M3Net works on a 3D feature space of resolution  $180 \times 180 \times 5$ , and obtains the  $512 \times 512 \times 40$  prediction results via up-sampling, which greatly reduces the computational load.

### Ablation Studies

We conduct ablation studies to verify the effectiveness of each component of M3Net. To save the training cost, we initially train M3Net for 6 epochs for detection and segmentation, and then add the occupancy prediction task to jointly fine-tune the three tasks for another 10 epochs. All experiment results are trained on the transformer-based M3Net.

**Multi-task model vs. single-task models.** As shown in the 3rd to 5th rows of Table 3, we independently train the M3Net-based single-task models using the same training epochs of multi-task learning. Then we train M3Net without our proposed TCS module and MAFI module but with our proposed multi-task query initialization, named M3Net multi-query, to perform multi-task learning, as shown in the 6th row in Table 3. Without TCS and MAFI modules, although different tasks are assigned to task-specific queries, conflicts between tasks still exist on the shared BEV features, resulting in the detection and segmentation tasks performance loss of 0.8%, 0.6% (total 1.4%) in terms of NDS

Method	Det NDS	Seg IoU	Occ mIoU	$\Delta$ MTL
BEVFusion detection-only	71.4	*	*	*
BEVFusion segmentation-only		62.7	*	*
M3Net detection-only	72.0	*	*	*
M3Net segmentation-only	*	66.3	*	*
M3Net occupancy-only	*	*	15.7	*
M3Net multi-query	71.2	65.7	21.7	+4.6
M3Net multi-query + MAFI	71.5	66.3	22.1	+5.9
M3Net multi-query + MAFI + TCS	72.2	66.7	22.4	+7.3

Table 3: Performances of M3Net based single-task models and different multi-task learning methods.  $\Delta$  MTL represents the sum of the improvements of the proposed method over the independently trained single-task models.

Query No.	Layout	IoU	Method	FFN	Para.(M)
6	$1 \times 6$	65.6	Tr. Dec	$\checkmark$	18.9
30	$5 \times 6$	66.3	Ma. Dec	$\times$	9.0
60	$10 \times 6$	66.4			

Table 4: Performance comparison of different segmentation query init. strategies.

Table 5: Comparison of parameters for different attention mechanisms in decoder.

and IoU. However, this performance loss is much lower than that of BEVFusion, where NDS and IoU decrease by 0.7% and 8.7% (total 9.4%) respectively. On the other hand, our M3Net multi-query baseline still achieves comparable detection performance and higher segmentation performance than independently trained BEVFusion, demonstrating the effectiveness of our proposed approach of tackling multiple tasks simultaneously with task-specific query-token interaction in a unified model. For the occupancy, there is a notable improvement of 6% on mIoU to single-task model because the occupancy queries can leverage shared knowledge.

**Effects of the modality-adaptive feature integration.** The 7th row (M3Net multi-query + MAFI) of Table 3, shows the effects of the MAFI module. The MAFI module brings improvements of 0.3%, 0.6%, and 0.4% in terms of NDS, IoU, and mIoU, respectively, to the multi-query baseline and a total performance improvement  $\Delta$  MTL of 5.9% for single-task models. Allowing single-modality features to predict channel-wise gating weights for their proficient tasks, the MAFI module better preserves the features that are beneficial to different single tasks during feature integration. As a result, the 3D occupancy prediction task also benefits from better-integrated BEV features.

**Effects of the task-oriented channel scaling.** The results in Table 3 demonstrate the TCS module’s efficacy in mitigating multi-task learning conflicts. Building upon MAFI’s multimodal feature integration, TCS further elevates performance by 7.3% ( $\Delta$  MTL) compared to single-task models. Specifically, TCS improves NDS by 0.7% in detection, IoU by 0.4% in segmentation, and mIoU by 0.3% in occupancy prediction relative to the MAFI baseline. Notably, these metrics surpass their single-task counterparts, with occupancy prediction showing a substantial 6.7% mIoU gain. These improvements validate TCS’s effectiveness.

**Effects of the segmentation query initialization.** Table 4 evaluates our distance-based query initialization strategy for segmentation. The query layout is represented as (M,

N), where M denotes the number of map blocks and N indicates the class count. The 6-query configuration assigns one query per category for global (100m $\times$ 100m) segmentation. The 30-query setting divides the range into 5 blocks, with 6 queries per (20m $\times$ 100m) block. Results show that the 6-query configuration underperforms due to limited global semantic representation capacity. The 30-query setting improves performance by 0.7%, validating our distance-based initialization approach. The marginal 0.1% gain from 60 queries leads us to adopt 30 queries as the optimal parameter-performance trade-off.

**Effects of the segmentation query initialization. Analysis of transformer-based and mamba-based decoders.** We present the parameter amount of transformer-based (Tr. Dec) and mamba-based (Ma. Dec) decoder of M3Net in Table 5. The mamba decoder exhibits significantly fewer parameters compared to the deformable transformer attention, primarily due to its independence from FFN. Meanwhile, our mamba-based M3Net achieves comparable performance to the transformer for detection task and obtains better results for map segmentation and occupancy prediction tasks, demonstrating the potential of using a mamba decoder in addressing perception tasks for autonomous driving.

## Conclusion

In this work, we propose a multi-modal and multi-task network M3Net, which unifies different tasks into a query-token interaction manner and can tackle 3D objection, map segmentation, and 3D occupancy prediction simultaneously. We first propose a modality-adaptive feature integration module to better combine the advantages of single-modal features and obtain the modality-adapted multi-modal BEV feature. Based on the integrated BEV feature, we propose a BEV-based task-specific query initialization strategy to accommodate the representation disparities for different tasks. With the initialized queries, we adopt a shared decoder to update queries, facilitating multi-task learning. Furthermore, a task-oriented channel scaling module is proposed to tackle the conflicts among tasks. The experiments on nuScenes datasets demonstrate that our approach effectively integrates the features of the two modalities and alleviates the conflicts of different tasks, outperforming state-of-the-art multi-task learning methods with remarkable margins.

## Acknowledgements

This project is funded in part by National Key R&D Program of China Project 2022ZD0161100, by Shanghai Artificial Intelligence Laboratory (Grant No. 2022ZD0160104), by the Centre for Perceptual and Interactive Intelligence (CPII) Ltd under the Innovation and Technology Commission (ITC)’s InnoHK, by General Research Fund of Hong Kong RGC Project 14204021, by Smart Traffic Fund PSRI/76/2311/PR for algorithm framework design and dataset curation. Hongsheng Li is a PI of CPII under the InnoHK.

## References

- Bai, X.; Hu, Z.; Zhu, X.; Huang, Q.; Chen, Y.; Fu, H.; and Tai, C.-L. 2022a. Transfusion: Robust lidar-camera fusion for 3d object detection with transformers. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 1090–1099.
- Bai, X.; Hu, Z.; Zhu, X.; Huang, Q.; Chen, Y.; Fu, H.; and Tai, C.-L. 2022b. Transfusion: Robust lidar-camera fusion for 3d object detection with transformers.
- Borsa, D.; Graepel, T.; and Shawe-Taylor, J. 2016. Learning shared representations in multi-task reinforcement learning. *arXiv preprint arXiv:1603.02041*.
- Borse, S.; Klingner, M.; Kumar, V. R.; Cai, H.; Almuzairee, A.; Yogamani, S.; and Porikli, F. 2023. X-Align: Cross-Modal Cross-View Alignment for Bird’s-Eye-View Segmentation.
- 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.
- Cai, Y.; Chen, X.; Zhang, C.; Lin, K.-Y.; Wang, X.; and Li, H. 2021. Semantic scene completion via integrating instances and scene in-the-loop. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 324–333.
- Cao, A.-Q.; and de Charette, R. 2022. Monoscene: Monocular 3d semantic scene completion. In *CVPR*.
- Chen, X.; Lin, K.-Y.; Qian, C.; Zeng, G.; and Li, H. 2020. 3d sketch-aware semantic scene completion via semi-supervised structure prior. In *CVPR*.
- Chen, X.; Shi, S.; Zhang, C.; Zhu, B.; Wang, Q.; Cheung, K. C.; See, S.; and Li, H. 2023. Trajectoryformer: 3d object tracking transformer with predictive trajectory hypotheses. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 18527–18536.
- Chen, X.; Shi, S.; Zhu, B.; Cheung, K. C.; Xu, H.; and Li, H. 2022. Mppnet: Multi-frame feature intertwining with proxy points for 3d temporal object detection. In *European Conference on Computer Vision*, 680–697. Springer.
- Chen, Z.; Badrinarayanan, V.; Lee, C.-Y.; and Rabinovich, A. 2018. Gradnorm: Gradient normalization for adaptive loss balancing in deep multitask networks. In *International Conference on Machine Learning*, 794–803.
- Cheng, B.; Choudhuri, A.; Misra, I.; Kirillov, A.; Girdhar, R.; and Schwing, A. G. 2021. Mask2former for video instance segmentation. *arXiv preprint arXiv:2112.10764*.
- Feng, D.; Zhou, Y.; Xu, C.; Tomizuka, M.; and Zhan, W. 2021. A simple and efficient multi-task network for 3d object detection and road understanding. In *2021 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 7067–7074. IEEE.
- 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.
- Goncalves, D. N.; Junior, J. M.; Zamboni, P.; Pistori, H.; Li, J.; Nogueira, K.; and Goncalves, W. N. 2023. MTLSegFormer: Multi-task Learning with Transformers for Semantic Segmentation in Precision Agriculture. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 6289–6297.
- Gu, A.; and Dao, T. 2023. Mamba: Linear-time sequence modeling with selective state spaces. *arXiv preprint arXiv:2312.00752*.
- Guo, P.; Lee, C.-Y.; and Ulbricht, D. 2020. Learning to branch for multi-task learning. In *International Conference on Machine Learning*, 3854–3863.
- He, K.; Gkioxari, G.; Dollár, P.; and Girshick, R. 2017. Mask r-cnn. In *Proceedings of the IEEE international conference on computer vision*, 2961–2969.
- Hong, Y.; Liu, Q.; Cheng, H.; Ma, D.; Dai, H.; Wang, Y.; Cao, G.; and Ding, Y. 2024. UniVision: A Unified Framework for Vision-Centric 3D Perception. *arXiv preprint arXiv:2401.06994*.
- 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*.
- Huang, Y.; Zheng, W.; Zhang, Y.; Zhou, J.; and Lu, J. 2023a. Tri-perspective view for vision-based 3d semantic occupancy prediction. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 9223–9232.
- Huang, Y.; Zheng, W.; Zhang, Y.; Zhou, J.; and Lu, J. 2023b. Tri-Perspective View for Vision-Based 3D Semantic Occupancy Prediction. *arXiv preprint arXiv:2302.07817*.
- Huang, Z.; Lin, S.; Liu, G.; Luo, M.; Ye, C.; Xu, H.; Chang, X.; and Liang, X. 2023c. FULLER: Unified Multi-modality Multi-task 3D Perception via Multi-level Gradient Calibration. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 3502–3511.
- Lang, A. H.; Vora, S.; Caesar, H.; Zhou, L.; Yang, J.; and Beijbom, O. 2019. Pointpillars: Fast encoders for object detection from point clouds.
- Li, J.; Han, K.; Wang, P.; Liu, Y.; and Yuan, X. 2020. Anisotropic convolutional networks for 3d semantic scene completion. In *CVPR*.
- Li, X.; Yuan, H.; Li, W.; Ding, H.; Wu, S.; Zhang, W.; Li, Y.; Chen, K.; and Loy, C. C. 2024. OMG-Seg: Is One Model Good Enough For All Segmentation? *CVPR*.
- Li, Y.; Yu, Z.; Choy, C.; Xiao, C.; Alvarez, J. M.; Fidler, S.; Feng, C.; and Anandkumar, A. 2023. Voxformer: Sparse voxel transformer for camera-based 3d semantic scene completion. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 9087–9098.
- Li, Z.; Wang, W.; Li, H.; Xie, E.; Sima, C.; Lu, T.; Qiao, Y.; and Dai, J. 2022a. Bevformer: Learning bird’s-eye-view representation from multi-camera images via spatiotemporal transformers. In *European conference on computer vision*, 1–18. Springer.
- 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.
- Lin, T.-Y.; Goyal, P.; Girshick, R.; He, K.; and Dollár, P. 2017. Focal loss for dense object detection. In *Proceedings of the IEEE international conference on computer vision*, 2980–2988.
- Liu, S.; Johns, E.; and Davison, A. J. 2019. End-to-end multi-task learning with attention. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 1871–1880.
- Liu, Y.; Tian, Y.; Zhao, Y.; Yu, H.; Xie, L.; Wang, Y.; Ye, Q.; and Liu, Y. 2024. VMamba: Visual State Space Model. *arXiv preprint arXiv:2401.10166*.
- Liu, Y.; Yan, J.; Jia, F.; Li, S.; Gao, A.; Wang, T.; and Zhang, X. 2023a. Petrv2: A unified framework for 3d perception from multi-camera images. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 3262–3272.
- 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*.
- Mordan, T.; Cord, M.; Pérez, P.; and Alahi, A. 2021. Detecting 32 pedestrian attributes for autonomous vehicles. *IEEE Transactions on Intelligent Transportation Systems*, 23(8): 11823–11835.
- Peng, L.; Chen, Z.; Fu, Z.; Liang, P.; and Cheng, E. 2023. BEVSegFormer: Bird’s Eye View Semantic Segmentation From Arbitrary Camera Rigs. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision*, 5935–5943.
- Roldao, L.; de Charette, R.; and Verroust-Blondet, A. 2020. Lmscnet: Lightweight multiscale 3d semantic completion. In *3DV*.
- Ruder, S.; Bingel, J.; Augenstein, I.; and Søgaard, A. 2019. Latent multi-task architecture learning. In *Proceedings of the AAAI Conference on Artificial Intelligence*, 4822–4829.
- Shi, S.; Guo, C.; Jiang, L.; Wang, Z.; Shi, J.; Wang, X.; and Li, H. 2020. Pv-rcnn: Point-voxel feature set abstraction for 3d object detection. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 10529–10538.
- Song, S.; Yu, F.; Zeng, A.; Chang, A. X.; Savva, M.; and Funkhouser, T. 2017. Semantic Scene Completion from a Single Depth Image. In *Conference on Computer Vision and Pattern Recognition (CVPR)*.
- Team, O.; et al. 2020. Openpcdet: An open-source toolbox for 3d object detection from point clouds. *OD Team*.
- Tian, X.; Jiang, T.; Yun, L.; Mao, Y.; Yang, H.; Wang, Y.; Wang, Y.; and Zhao, H. 2024. Occ3d: A large-scale 3d occupancy prediction benchmark for autonomous driving. *Advances in Neural Information Processing Systems*, 36.
- Vora, S.; Lang, A. H.; Helou, B.; and Beijbom, O. 2020. Pointpainting: Sequential fusion for 3d object detection.
- Wallingford, M.; Li, H.; Achille, A.; Ravichandran, A.; Fowlkes, C.; Bhotika, R.; and Soatto, S. 2022. Task adaptive parameter sharing for multi-task learning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 7561–7570.
- Wang, X.; Zhou, J.; Mou, S.; and Corless, M. J. 2019. A distributed algorithm for least squares solutions. *IEEE Transactions on Automatic Control*, 64(10): 4217–4222.
- Wang, X.; Zhu, Z.; Xu, W.; Zhang, Y.; Wei, Y.; Chi, X.; Ye, Y.; Du, D.; Lu, J.; and Wang, X. 2023. Openoccupancy: A large scale benchmark for surrounding semantic occupancy perception. *arXiv preprint arXiv:2303.03991*.
- Wang, Y.; Chen, Y.; Liao, X.; Fan, L.; and Zhang, Z. 2024. Panoocc: Unified occupancy representation for camera-based 3d panoptic segmentation. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 17158–17168.
- Wei, Y.; Zhao, L.; Zheng, W.; Zhu, Z.; Zhou, J.; and Lu, J. 2023. Surroundocc: Multi-camera 3d occupancy prediction for autonomous driving. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 21729–21740.
- Xie, E.; Yu, Z.; Zhou, D.; Phillion, J.; Anandkumar, A.; Fidler, S.; Luo, P.; and Alvarez, J. M. 2022. M<sup>2</sup>bev: Multi-camera joint 3d detection and segmentation with unified birds-eye view representation. *arXiv preprint arXiv:2204.05088*.
- Yan, X.; Gao, J.; Li, J.; Zhang, R.; Li, Z.; Huang, R.; and Cui, S. 2021. Sparse single sweep lidar point cloud segmentation via learning contextual shape priors from scene completion. In *AAAI*.
- Ye, D.; Zhou, Z.; Chen, W.; Xie, Y.; Wang, Y.; Wang, P.; and Foroosh, H. 2023. Lidarmultinet: Towards a unified multi-task network for lidar perception. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 37, 3231–3240.
- Yin, T.; Zhou, X.; and Krahenbuhl, P. 2021. Center-based 3d object detection and tracking.
- Yin, T.; Zhou, X.; and Krähenbühl, P. 2021. Multimodal virtual point 3d detection.
- Zhang, Y.; Zhu, Z.; and Du, D. 2023. OccFormer: Dual-path Transformer for Vision-based 3D Semantic Occupancy Prediction. *arXiv preprint arXiv:2304.05316*.
- Zhou, J.; Wang, X.; Mou, S.; and Anderson, B. D. 2019. Finite-time distributed linear equation solver for solutions with minimum  $l_1$ -norm. *IEEE Transactions on Automatic Control*, 65(4): 1691–1696.
- 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, X.; Su, W.; Lu, L.; Li, B.; Wang, X.; and Dai, J. 2020. Deformable detr: Deformable transformers for end-to-end object detection. *arXiv preprint arXiv:2010.04159*.