UniM²AE: Multi-modal Masked Autoencoders with Unified 3D Representation for 3D Perception in Autonomous Driving

Jian Zou¹, Tianyu Huang¹[®], Guanglei Yang¹[∞], Zhenhua Guo²[®], Tao Luo³[∞], Chun-Mei Feng³[®], and Wangmeng Zuo¹[®]

¹ Harbin Institute of Technology, China

yangguanglei@hit.edu.cn

² Tianyijiaotong Technology Ltd., China ³ Institute of High Performance Computing (IHPC), Agency for Science, Technology

and Research (A*STAR), Singapore

tluo001@e.ntu.edu.sg

Abstract. Masked Autoencoders (MAE) play a pivotal role in learning potent representations, delivering outstanding results across various 3D perception tasks essential for autonomous driving. In real-world driving scenarios, it's commonplace to deploy multiple sensors for comprehensive environment perception. Despite integrating multi-modal features from these sensors can produce rich and powerful features, there is a noticeable challenge in MAE methods addressing this integration due to the substantial disparity between the different modalities. This research delves into multi-modal Masked Autoencoders tailored for a unified representation space in autonomous driving, aiming to pioneer a more efficient fusion of two distinct modalities. To intricately marry the semantics inherent in images with the geometric intricacies of LiDAR point clouds, we propose $UniM^2AE$. This model stands as a potent yet straightforward, multi-modal self-supervised pre-training framework, mainly consisting of two designs. First, it projects the features from both modalities into a cohesive 3D volume space to intricately marry the bird's eye view (BEV) with the height dimension. The extension allows for a precise representation of objects and reduces information loss when aligning multi-modal features. Second, the Multi-modal 3D Interactive Module (MMIM) is invoked to facilitate the efficient inter-modal interaction during the interaction process. Extensive experiments conducted on the nuScenes Dataset attest to the efficacy of UniM²AE, indicating enhancements in 3D object detection and BEV map segmentation by 1.2% NDS and 6.5% mIoU, respectively. The code is available at https://github.com/hollow-503/UniM2AE.

Keywords: Unified representation \cdot sensor fusion \cdot masked autoencoders



Fig. 1: (a) Multi-modal frameworks [6] that align masked input before feature extraction but ignore feature characteristics from two branch. (b) $\text{Uni}M^2\text{AE}$ that interacts multi-modal features with unified representation.

1 Introduction

Autonomous driving marks a transformative leap in transportation, heralding potential enhancements in safety, efficiency, and accessibility [23,25,41,45]. Fundamental to this progression is the vehicle's capability to decode its surroundings [42,48,57]. To tackle the intricacies of real-world contexts, integration of various sensors is imperative: cameras yield detailed visual insights, LiDAR grants exact geometric data, *etc.* Through this multi-sensor fusion, a comprehensive grasp of the environment is achieved [1,28]. Nonetheless, the reliance on extensively labeled data for multi-sensor fusion poses a significant challenge due to the high costs involved.

Masked Autoencoders (MAE) [15] have shown promise in reducing reliance on labeled datasets, especially noted in 2D vision tasks. A natural extension of this success in autonomous driving involves projecting LiDAR point clouds onto the image plane, thereby masking image patches in conjunction with their corresponding LiDAR data. However, this approach faces a significant challenge due to geometric distortions when projecting LiDAR data onto the camera plane, as depicted in Fig. 1(a). This distortion arises due to the inherent spatial discrepancy within the same patch: point clouds that are physically distant or close may appear next to each other on the pixel coordinates, making it impractical to indiscriminately divide and mask these merged image and point cloud patches for reconstruction. Moreover, the discrepancy in the operational range between LiDAR sensors and cameras exacerbates this challenge. Due to the different sensor setups, LiDAR may capture point clouds beyond the camera's field of view, as seen in datasets like Waymo Open Dataset [36], which provides 360-degree Li-DAR views versus the camera's limited frontal and lateral coverage. As a result, a vast amount of LiDAR points are unprojectable. Alternatively, some sensor fusion strategies [35, 39, 40, 53] attempt to interact image and point clouds by employing the camera-to-LiDAR projection. However, this approach is hindered by the inherent sparsity of LiDAR data, resulting in a substantial loss of camera features and, consequently, a detrimental impact on effective feature fusion.

To address the above challenges, we introduce the $\text{Uni}M^2\text{AE}$, depicted in Fig. 1(b), as an innovative self-supervised pre-training framework designed to

3

harmonize the integration of two distinct modalities: images and LiDAR data. $UniM^2AE$ endeavors to establish a unified representation space that enhances the fusion of these modalities. By leveraging the semantic richness of images in tandem with the precise geometric details captured by LiDAR. Uni M^2AE facilitates the generation of robust, cross-modal features. Central to UniM²AE is the innovation of a 3D volume space, achieved by extending the z-axis of the Bird's Eve View (BEV) representation. This crucial expansion preserves height, enabling a more faithful representation of objects that exhibit significant variation in height. It also mitigates the information loss that is typically encountered when projecting features into or retrieving them from the representation space. Unlike the traditional BEV representation [21] and the occupancy representation [37, 42], our unified 3D representation retains sufficient details along the height dimension. This enables accurate re-projection to the original modalities for reconstructing the multi-modal inputs. Moreover, within this enriched 3D volume space, the information contained within each voxel can be derived either from features of its native modality or from features of other modalities. Each modality's decoder then utilizes this unified feature set to reconstruct its specific inputs, pushing the encoder to learn more generalized and cross-modal features that encapsulate a comprehensive understanding of the scene or objects. This architecture not only addresses the previously outlined limitations but also sets a new benchmark for multi-modal integration, advancing the field of autonomous driving by enabling more nuanced and effective utilization of sensor data.

In the expansive and dynamic environments characteristic of autonomous driving, which feature a multitude of objects and intricate inter-instance relationships, a sophisticated approach is required for efficient feature amalgamation. To this end, the Multi-modal 3D Interaction Module (MMIM) is employed to further refine the fusion of modalities within our unified 3D volume space, signifying a pivotal advancement in our framework's capability to facilitate efficient interaction between the dual branches of input data. The MMIM's architecture, built upon stacked 3D deformable self-attention blocks [43,59], enables the modeling of global context at elevated resolutions. This feature is instrumental in boosting the performance of downstream tasks, thereby addressing the complexities inherent in autonomous driving scenarios. The deformable nature of the self-attention mechanism allows for adaptive focus on the most salient features within vast scenes, ensuring that the system is attuned to the nuanced dynamics and relationships present [43]. Consequently, this leads to quicker model convergence and improved pre-training efficiency, marking a significant leap forward in our endeavor to enhance the functionality of autonomous driving systems.

To sum up, our contributions can be presented as follows:

- We propose UniM²AE, a multi-modal self-supervised pre-training framework with unified representation in a cohesive 3D volume space. The corresponding representation enables the alignment of the multi-modal features with less information loss, facilitating the reconstruction of multi-modal masked inputs.

- 4 J. Zou et al.
- To better interact semantics and geometries retained in unified 3D volume space, we introduce a Multi-modal 3D Interaction Module (MMIM) to effectively obtain more informative and powerful features.
- We conduct extensive experiments on various 3D downstream tasks, where UniM²AE notably promotes diverse detectors and shows competitive performance.

2 Related Work

2.1 Multi-modal Representation

Multi-modal representation has raised significant interest recently, especially in vision-language pre-training [33,51]. Some works [19,50,55] align point cloud data to 2D vision by depth projection. As for unifying 3D with other modalities, the bird's-eye view (BEV) is a widely-used representation since the transformation to BEV space retains both geometric structure and semantic density. Although many SOTA methods [9,17,25,28] adopt BEV as the unified representation, the lack of height information leads to a poor description of the shape and position of objects, which makes it unsuitable for MAE. In this work, we introduce a unified representation with height dimension in 3D volume space, which captures the detailed height and position of objects.

2.2 Masked Autoencoders

Masked Autoencoders (MAE) [15] are a self-supervised pre-training method, with a pre-text task of predicting masked pixels. With its success, a series of 3D representation learning methods apply masked modeling to 3D data. Some works [26,31,54] reconstruct masked points of indoor objects. Some works [4,29, 38,46,49] predict the masked voxels in outdoor scenes. Recent methods propose multi-modal MAE pre-training: Zhang [56] exploit 2D pre-trained knowledge to the 3D point prediction but fail to exploit the full potential of LiDAR point cloud and image datasets. Chen [6] attempt to tackle it in the indoor scene and conduct a 2D-3D joint prediction by the projection of points but ignore the characteristic of point clouds and images. To address these problems, we propose to predict both masked 2D pixels and masked 3D voxels in a unified representation, focusing on the autonomous driving scenario.

2.3 Multi-modal Fusion in 3D Perception

Recently, multi-modal fusion has been well-studied in 3D perception. Proposallevel fusion methods adopt proposals in 3D and project the proposals to images to extract RoI feature [7,8,30]. Point-level fusion methods usually paint image semantic features onto foreground LiDAR points, which can be classified into input-level decoration [39, 40, 53], and feature-level decoration [22, 24]. However, the camera-to-LiDAR projection is semantically lossy due to the different



Fig. 2: Overview of UniM²AE. The LiDAR branch voxelize the point cloud, while the camera branch divides multiple images into patches, both subsequently randomly masking their inputs. The tokens from the two branches are individually embedded and then passed through the Token-Volume projection, Multi-modal 3D Interaction Module, Volume-Token projection, and eventually the modality-specific decoder. Ultimately, we reconstruct the original inputs using the fused features.

densities of both modalities. Some BEV-based approaches [12] aim to mitigate this problem, but their simple fusion modules fall short in retaining the height information and remain blurry along z-axis, which is disadvantageous for accurate image reconstruction. Accordingly, we design the Multi-modal 3D Interaction Module to effectively fuse the projected 3D volume features.

3 Proposed Method

3.1 Overview Architecture

As shown in Fig. 2, Uni M^2AE learns multi-modal representation by masking inputs (I, V) and jointly combine features projected to the 3D volume space (F_V^{vol}, F_I^{vol}) to accomplish the reconstruction. In our proposed pipeline, the point cloud is first embedded into tokens after voxelization and similarly we embed the images with position encoding after dividing the them into non-overlapping patches. Following this, tokens from both modalities are randomly masked, producing (M_I, M_V) . Separate transformer-based encoders are then utilized to extract features (F_I, F_V) .

To align features from various modalities with the preservation of semantics and geometrics, (F_I, F_V) are separately projected into the unified 3D volume space, which is extended BEV along the height dimension. Specifically, we build a mapping of each voxel to 3D volume space based on its position in the egovehicle coordinates, while for the image tokens, the spatial cross-attention is employed for 2D to 3D conversion. The projected feature (F_V^{vol}, F_I^{vol}) are subsequently passed into the Multi-modal 3D Interaction Module (MMIM), aiming at promoting powerful feature fusion.

Following the cross-modal interaction, we project the fused feature F'_c back to the modality-specific token, denoted F_V^{sp} for LiDAR and $F_I^{proj} \in (C, H, W)$ (which is then reshaped to $F_I^{sp} \in (HW, C)$) for camera. The camera decoder and LiDAR decoder are finally used to reconstruct the original inputs.

3.2 Unified Representation in 3D Volume Space

Different sensors capture data that, while representing the same scene, often provide distinct descriptions. For instance, camera-derived images emphasize the color palette of the environment within their field of view, whereas point clouds primarily capture object locations. Given these variations, selecting an appropriate representation for fusing features from disparate modalities becomes paramount. Such a representation must preserve the unique attributes of multimodal information sourced from various sensors.

In pursuit of capturing the full spectrum of object positioning and appearance, the voxel feature in 3D volume space is adopted as the unified representation, depicted in Fig. 1(b). The 3D volume space uniquely accommodates the height dimension, enabling it to harbor more expansive geometric data and achieve exacting precision in depicting object locations, exemplified by features like elevated traffic signs. This enriched representation naturally amplifies the accuracy of interactions between objects. A salient benefit of the 3D volume space is its capacity for direct remapping to original modalities, cementing its position as an optimal latent space for integrating features. Due to the intrinsic alignment of images and point clouds within the 3D volume space, the Multi-modal 3D Interaction Module can bolster representations across streams, sidestepping the need for additional alignment mechanisms. Such alignment streamlines the transition between pre-training and fine-tuning, producing favorable outcomes for subsequent tasks. Additionally, the adaptability of the 3D volume space leaves the door open for its extension to encompass three or even more modalities.

3.3 Multi-modal Interaction

Projection to 3D Volume Space In the projection of LiDAR to the 3D volume, the voxel embedding is directed to a predefined 3D volume using positions from the ego car coordinate system. This method ensures no geometric distortion is introduced. The resulting feature from this process is denoted as F_V^{vol} . For the image to 3D volume projection, the 2D-3D Spatial Cross-Attention method is employed. Following prior works [23, 42], 3D volume queries for each image is defined as $Q^{vol} \in \mathbb{R}^{C \times H \times W \times Z}$. The corresponding 3D points are then projected to 2D views using the camera's intrinsic and extrinsic parameters. During this projection, the 3D points associate only with specific views, termed as \mathcal{V}_{hit} . The 2D feature is then sampled from the view \mathcal{V}_{hit} at the locations of those projected 3D reference points. The process unfolds as:

$$F_I^{vol} = \frac{1}{|\mathcal{V}_{hit}|} \sum_{i \in \mathcal{V}_{hit}} \sum_{j=1}^{N_{ref}} \text{DeformAttn}\left(Q^{vol}, \mathcal{P}(p, i, j), F_I^i\right)$$
(1)



Fig. 3: Illustration of our Multi-modal 3D Interaction Module. We first concatenate the inputs (F_V^{vol}, F_I^{vol}) and reshape it for the subsequent stacking 3D deformable self-attention blocks. After interaction, we split the output and project them back to feature token. This contributes more generalized and effective feature learning.

where *i* indexes the camera view, *j* indexes the 3D reference points, and N_{ref} is the total number of points for each 3D volume query. F_I^i is the features of the *i*-th camera view. $\mathcal{P}(p, i, j)$ is the project function that gets the *j*-th reference point on the *i*-th view image.

Multi-modal 3D Interaction Module To fuse the projected 3D volume features from the camera $(F_I^{vol} \in \mathbb{R}^{C \times H \times W \times Z})$ and the LiDAR $(F_V^{vol} \in \mathbb{R}^{C \times H \times W \times Z})$ branches effectively, the Multi-modal 3D Interaction Module (MMIM) is introduced. As depicted in Fig. 3, MMIM comprises L attention blocks, with L = 3being the default setting.

Given the emphasis on high performance at high resolutions in downstream tasks and the limited scale of token sequences in standard self-attention, deformable self-attention is selected to alleviate computational demands. Each block is composed of deformable self-attention, a feed-forward network, and normalization. Initially, the concatenation of F_V^{vol} and F_I^{vol} is performed along the channel dimension, reshaping the result to form the query token $F_c^{vol} \in \mathbb{R}^{HWZ \times 2C}$. This token is then inputted into the Multi-modal 3D Interaction Module, an extension of [59], to promote effective modal interaction. The interactive process can be described as follows:

$$F'_{c} = \sum_{m=1}^{M} W_{m} \sum_{k=1}^{K} A_{mk} \cdot W'_{m} F^{vol}_{c} \left(p^{vol} + \Delta p^{vol}_{k} \right)$$
(2)

where *m* indexes the attention head, *k* indexes the sampled keys, and *K* is the total sampled key number. Δp_k^{vol} and A_{mk} denote the sampling offset and attention weight of the *k*-th sampling point in the *m*-th attention head, respectively. The attention weight A_{mk} lies in the range [0, 1], normalized by $\sum_{k=1}^{K} A_{mk} = 1$. At the end, F'_c is split along the channel dimension to obtain the modality-specific 3D volume features (F'_V, F'_I) .

Projection to Modality-specific Token By tapping into the advantages of the 3D volume representation, the fused feature can be conveniently projected onto the 2D image plane and 3D voxel token. For the LiDAR branch, the process merely involves sampling the features located at the position of the masked voxel token within the ego-vehicle coordinates. Notably, these features have already been enriched by the fusion module with semantics from the camera branch. Regarding the camera branch, the corresponding 2D coordinate (u, v) can be determined using the projection function T. The 2D-plane feature F_I^{proj} is then obtained by mapping the 3D volume feature in (x, y, z) to the position (u, v). The projection function T_{proj} is defined as :

$$z \begin{bmatrix} u \\ v \\ 1 \end{bmatrix} = T_{proj}(P) = K \cdot R_t \cdot \begin{vmatrix} x \\ y \\ z \\ 1 \end{vmatrix}$$
(3)

where $P \in \mathbb{R}^3$ is the position in 3D volume, $K \in \mathbb{R}^{3 \times 4}$, $R_t \in \mathbb{R}^{4 \times 4}$ are the camera intrinsic and extrinsic matrices.

3.4 Prediction Target

Three distinct reconstruction tasks supervise each modal decoder. A single linear layer is applied to the output of each decoder for each task. The dual-modal reconstruction tasks and their respective loss functions are detailed below. In alignment with Voxel-MAE [16], the prediction focuses on the number of points within each voxel. Supervision for this reconstruction uses the Chamfer distance, which gauges the disparity between two point sets of different scales. Let G_n denote the masked LiDAR point cloud partitioned into voxels. The Chamfer loss \mathcal{L}_c can be presented as:

$$\mathcal{L}_{c} = CD\left(Dec_{V}\left(F_{V}^{sp}\right), G_{n}\right) \tag{4}$$

where $CD(\cdot)$ stands for Chamfer distance function [13], $Dec_V(\cdot)$ denote voxel decoder and F_V^{sp} represents projected voxel features.

In addition to the aforementioned reconstruction task, there is a prediction to ascertain if a voxel is empty. Supervision for this aspect employs the binary cross entropy loss, denoted as \mathcal{L}_{occ} . The cumulative voxel reconstruction loss is thus given as:

$$\mathcal{L}_{voxel} = \mathcal{L}_c + \mathcal{L}_{occ} \tag{5}$$

For the camera branch, the pixel reconstruction is supervised using the Mean Squared Error (MSE) loss, represented as :

$$\mathcal{L}_{img} = \mathcal{L}_{MSE} \left(Dec_I \left(F_I^{sp} \right), G_I \right) \tag{6}$$

where G_I is the original images in pixel space, $Dec_I(\cdot)$ denotes the image decoder and F_I^{sp} represents projected image features.

4 Experiments

4.1 Implementation Details

Dataset and Metrics The nuScenes Dataset [5], a comprehensive autonomous driving dataset, serves as the primary dataset for both pre-training our model and evaluating its performance on multiple downstream tasks. This dataset encompasses 1,000 sequences gathered from Boston and Singapore, with 700 designated for training and 300 split evenly for validation and testing. Each sequence, recorded at 10Hz, spans 20 seconds and is annotated at a frequency of 2Hz. In terms of 3D detection, the principal evaluation metrics employed are mean Average Precision (mAP) and the nuScenes detection score (NDS). For the task of BEV map segmentation, the methodology aligns with the dataset's map expansion pack, using Intersection over Union (IoU) as the assessment metric.

Network Architectures The UniM²AE utilizes SST [14] and Swin-T [27] as the backbones for the LiDAR Encoder and Camera Encoder, respectively. In the Multi-modal 3D Interaction Module, 3 deformable self-attention blocks are stacked, with each attention module comprising 128 input channels and 256 hidden channels. To facilitate the transfer of pre-trained weights for downstream tasks, BEVFusion-SST and TransFusion-L-SST are introduced, with the LiDAR backbone in these architectures being replaced by SST.

Pre-training During this stage, the perception ranges are restricted to [-50m, 50m] for x- and y-axes, [-5m, 3m] for z-axes. Each voxel has dimensions of (0.5m, 0.5m, 4m). In terms of input data masking during the training phase, our experiments have determined that a masking ratio of 70% for the LiDAR branch and 75% for the camera branch yields optimal results. By default, all the MAE methods are trained with a total of 200 epochs on 8 GPUs, a base learning rate of 2.5e-5. Detailed configurations are reported in the supplemental material.

Fine-tuning Utilizing the encoders from UniM²AE for both camera and Li-DAR, the process then involves fine-tuning and assessing the capabilities of the features learned on tasks that are both single-modal and multi-modal in nature. For tasks that are solely single-modal, one of the pre-trained encoder serves as the feature extraction mechanism. When it comes to multi-modal tasks, which include 3D object detection and BEV map segmentation, both the LiDAR encoder and the camera branch's encoder are capitalized upon as the feature extractors pertinent to these downstream tasks. It's pivotal to note that while the decoder plays a role during the pre-training phase, it's omitted during the fine-tuning process. A comparison of the pre-trained feature extractors with a variety of baselines across different tasks was conducted, ensuring the experimental setup remained consistent. Notably, when integrated into fusion-based methodologies, the pre-trained Multi-modal 3D Interaction Module showcases competitive performance results. Detailed architectures and configurations are in the supplemental material.

Table 1: Data-efficient 3D object detection results of on the nuScenes validation set. Backbones in single-modality and multi-modality are pre-trained using various MAE methods. The model performances are reported using different amounts of fine-tuning data. Random denotes training from scratch. MIM+Voxel-MAE: for the initialization of the model, the image backbone loads the GreenMIM [18] pre-trained weights, whereas the point cloud backbone loads the Voxel-MAE [16] pre-trained weight. L and C represent LiDAR and Camera, respectively. *: our re-implementation.

Data amount Modality		Initialization	mAP	NDS	Car	Truck	C.V.	Bus	Trailer	Barrier	Motor	Bike	Ped.	T.C.
	\mathbf{L}	Random Voxel-MAE* UniM²AE	44.3 48.9 50.0	56.3 59.8 60.0	78.8 80.9 81.0	41.1 47.0 47.8	13.1 12.8 12.0	50.7 59.0 57.3	18.8 23.6 24.0	52.8 61.9 62.3	46.1 47.8 51.2	$17.3 \\ 23.5 \\ 26.8$	75.7 79.9 81.4	49.1 52.4 55.8
20%	C+L	Random MIM+Voxel-MAE PiMAE* [6] UniM²AE	51.5 54.3 52.5 55.9	50.9 51.2 52.0 52.8	84.1 84.3 84.1 85.8	47.6 50.8 48.9 51.1	13.3 18.9 17.5 19.3	49.8 52.3 50.0 54.2	27.9 28.9 28.0 30.6	65.0 68.4 66.2 69.0	53.9 57.4 51.6 61.1	26.8 32.1 26.8 34.3	78.8 80.3 80.8 83.0	67.8 69.2 70.9 70.8
40%	L	Random Voxel-MAE* UniM²AE	50.9 52.6 52.9	60.5 62.2 62.6	81.4 82.4 82.7	47.6 49.1 49.2	13.7 15.4 15.8	58.0 62.2 60.1	24.5 25.8 23.7	61.3 64.5 65.5	57.7 56.8 58.4	30.1 30.4 31.2	80.4 82.3 83.8	54.4 57.5 58.9
	C+L	Random MIM+Voxel-MAE PiMAE* [6] UniM²AE	58.6 60.2 61.4 62.0	61.9 63.5 64.0 64.5	86.2 86.6 86.9 87.0	54.7 56.5 57.8 57.8	21.7 22.5 25.2 22.8	60.0 64.1 64.6 62.7	33.0 33.5 36.2 38.7	70.7 72.2 70.6 69.7	64.2 66.1 69.8 66.8	38.6 41.9 44.3 50.5	83.0 83.8 83.0 86.0	74.3 74.8 75.5 77.9
60%	L	Random Voxel-MAE* UniM²AE	51.9 54.2 54.7	61.7 63.5 63.8	82.2 83.0 83.1	49.0 51.1 51.0	15.6 16.3 17.3	61.2 62.0 62.5	24.9 27.5 26.9	62.9 64.9 65.1	56.3 61.2 62.2	32.1 34.7 35.7	81.6 82.9 83.4	53.1 58.0 59.9
	C+L	Random MIM+Voxel-MAE PiMAE* [6] UniM²AE	61.6 62.1 62.3 62.4	65.2 65.7 65.5 66.1	87.3 87.2 87.3 87.7	58.5 56.7 58.8 59.7	23.9 23.0 24.9 23.8	65.2 65.4 61.6 67.6	35.8 37.0 38.7 37.0	71.9 71.7 72.7 70.5	67.8 70.6 68.5 68.4	46.7 47.3 45.7 48.9	85.7 85.6 86.0 86.6	77.0 76.7 79.1 77.8
80%	L	Random Voxel-MAE* UniM²AE	52.7 55.1 55.6	62.5 64.2 64.6	82.3 83.4 83.4	49.6 51.7 52.9	16.0 18.8 18.2	63.3 64.0 64.2	25.8 28.7 29.4	60.7 63.8 64.7	58.6 62.2 63.1	31.6 35.1 36.1	82.0 84.3 84.5	56.7 58.7 58.8
	C+L	Random MIM+Voxel-MAE PiMAE* [6] UniM²AE	62.5 63.0 63.7 63.9	66.1 66.4 66.6 67.1	87.1 87.6 87.0 87.7	57.6 59.6 61.9 59.6	24.0 24.1 25.3 24.9	66.4 66.1 71.2 69.2	38.1 38.0 38.6 39.8	71.1 71.3 70.3 71.3	68.4 70.2 73.9 71.0	48.8 48.8 49.3 50.2	86.2 86.5 84.6 86.8	77.6 78.1 74.7 78.7
100%	L	Random Voxel-MAE* UniM²AE	53.6 55.3 55.8	63.0 64.1 64.6	82.3 83.2 83.3	49.8 51.2 51.3	16.7 16.8 17.6	64.0 64.6 63.7	26.2 28.3 28.6	60.9 65.0 65.4	61.7 61.8 62.8	33.0 39.6 40.8	82.2 83.6 83.9	58.9 58.9 60.3
	C+L	Random MIM+Voxel-MAE PiMAE* [6] UniM²AE	63.6 63.7 63.9 64.3	67.4 67.7 67.9 68.1	87.7 87.6 87.3 87.9	58.0 58.3 58.7 57.8	26.6 25.3 27.0 24.3	67.8 67.1 67.6 68.6	38.4 38.8 38.7 42.2	72.6 70.8 69.7 71.6	71.8 71.7 71.4 72.5	47.6 51.5 55.4 51.0	87.0 86.7 86.8 87.2	78.9 78.9 76.6 79.5

4.2 Data Efficiency

The primary motivation behind employing MAE is to minimize the dependency on annotated data without compromising the efficiency and performance of the model. In assessing the representation derived from the pre-training with UniM²AE, experiments were conducted on datasets of varying sizes, utilizing different proportions of the labeled data. Notably, for training both the singlemodal and multi-modal 3D Object Detection models, fractions of the annotated dataset, namely $\{20\%, 40\%, 60\%, 80\%, 100\%\}$, are used.

Method	Modality	Voxel Size(m)	$NDS\uparrow$	$\mathrm{mAP}\uparrow$	mATE↓	mASE↓	mAOE↓	$mAVE\downarrow$	$\mathrm{mAAE}{\downarrow}$
CenterPoint [52]	L	[0.075, 0.075, 0.2]	66.8	59.6	29.2	25.5	30.2	25.9	19.4
LargeKernel3D [11]	L	[0.075, 0.075, 0.2]	69.1	63.9	28.6	25.0	35.1	21.1	18.7
TransFusion-L [2]	L	[0.075, 0.075, 0.2]	70.1	65.4	-	-	-	-	-
TransFusion-L-SST	L	[0.15, 0.15, 8]	69.9	65.0	28.0	25.3	30.1	24.1	19.0
TransFusion-L-SST+ $UniM^2AE-L$	L	[0.15, 0.15, 8]	70.4	65.7	28.0	25.2	29.5	23.5	18.6
FUTR3D [8]	C+L	[0.075, 0.075, 0.2]	68.3	64.5	-	-	-	-	-
Focals Conv [10]	C+L	[0.075, 0.075, 0.2]	69.2	64.0	33.2	25.4	27.8	26.8	19.3
MVP [53]	C+L	[0.075, 0.075, 0.2]	70.7	67.0	28.9	25.1	28.1	27.0	18.9
TransFusion [2]	C+L	[0.075, 0.075, 0.2]	71.2	67.3	27.2	25.2	27.4	25.4	19.0
MSMDFusion [20]	C+L	[0.075, 0.075, 0.2]	72.1	69.3	-	-	-	-	-
BEVFusion [28]	C+L	[0.075, 0.075, 0.2]	71.4	68.5	28.7	25.4	30.4	25.6	18.7
BEVFusion-SST	C+L	[0.15, 0.15, 8]	71.5	68.2	27.8	25.3	30.2	23.6	18.9
BEVFusion-SST+UniM ² AE	C+L	[0.15, 0.15, 8]	71.9	68.4	27.2	25.2	28.8	23.2	18.7
BEVFusion-SST+ $UniM^2AE^{\dagger}$	C+L	[0.15, 0.15, 4]	72.7	69.7	26.9	25.2	27.3	23.2	18.9
FocalFormer3D [9]	C+L	[0.075, 0.075, 0.2]	73.1	70.5	-	-	-	-	-
FocalFormer3D*+ $UniM^2AE^{\dagger}$	$\rm C+L$	[0.15, 0.15, 4]	73.8	71.1	26.9	25.0	26.7	19.6	18.9

Table 2: Performances of 3D object detection on the nuScenes validation split. [†]: Finetuned with the pre-trained MMIM. ^{*}: Our re-implementation with SST backbone.

In the realm of single-modal 3D self-supervised techniques, UniM²AE is juxtaposed against Voxel-MAE [16] using an anchor-based detector. Following the parameters set by Voxel-MAE, the detector undergoes training for 288 epochs with a batch size of 4 and an initial learning rate pegged at 1e-5. On the other hand, for multi-modal strategies, evaluations are conducted on a fusion-based detector [28] equipped with a TransFusion head [2]. As per current understanding, this is a pioneering attempt at implementing multi-modal MAE in autonomous driving. For the sake of a comparative analysis, a combination of the pre-trained Swin-T from GreenMIM [18] and SST from Voxel-MAE [16] is utilized.

According to the result in Table 1, the proposed UniM^2AE presents a substantial enhancement to the detector, exhibiting an improvement of 4.4/1.9 mAP/NDS over random initialization and 1.6/1.6 mAP/NDS compared to the basic amalgamation of GreenMIM and Voxel-MAE when trained on just 20% of the labeled data. Impressively, even when utilizing the entirety of the labeled dataset, UniM^2AE continues to outperform, highlighting its superior ability to integrate multi-modal features in the unified 3D volume space during the pretraining phase. Moreover, it's noteworthy to mention that while UniM^2AE isn't specifically tailored for a LiDAR-only detector, it still yields competitive outcomes across varying proportions of labeled data. This underscores the capability of UniM^2AE to derive more insightful representations.

PiMAE [6], as the first framework to explore multi-modal MAE in indoor scenarios, projects point clouds onto the image plane aligned and masked with the image, followed by concatenation of dual modality tokens along the sequence dimension to pre-train the ViT backbone. We compare the UniM²AE to the previous SOTA multi-modal pre-training method *i.e.* PiMAE, on the nuScenes benchmark [5]. As shown in Table 1, when trained with 20% labeled data, Pi-MAE outperforms training from scratch with a marginal increase of 1.1/1.0NDS/mAP due to the loss of depth information in aligning point clouds with images in the pixel coordinates. Conversely, UniM²AE, aligning in the unified

Method	Modality	Drivable	Ped. Cross.	Walkway	Stop Line	Carpark	Divider	mIoU
CVT [58]	С	74.3	36.8	39.9	25.8	35.0	29.4	40.2
OFT [34]	С	74.0	35.3	45.9	27.5	35.9	33.9	42.1
LSS [32]	С	75.4	38.8	46.3	30.3	39.1	36.5	44.4
$M^{2}BEV$ [44]	\mathbf{C}	77.2	-	-	-	-	40.5	-
BEVFusion [*] [28]	С	78.2	48.0	53.5	40.4	45.3	41.7	51.2
$BEVFusion^* + UniM^2AE-C$	\mathbf{C}	79.5	50.5	54.9	42.4	47.3	42.9	52.9
MVP [53]	$\rm C+L$	76.1	48.7	57.0	36.9	33.0	42.2	49.0
PointPainting [39]	C+L	75.9	48.5	57.7	36.9	34.5	41.9	49.1
BEVFusion [28]	C+L	85.5	60.5	67.6	52.0	57.0	53.7	62.7
X-Align [3]	C+L	86.8	65.2	70.0	58.3	57.1	58.2	65.7
BEVFusion-SST	C+L	84.9	59.2	66.3	48.7	56.0	52.7	61.3
$BEVFusion-SST+UniM^2AE$	C+L	85.1	59.7	66.6	48.7	56.0	52.6	61.4
$\mathrm{BEVFusion}\text{-}\mathrm{SST}\text{+}\mathbf{UniM}^{2}\mathbf{AE}^{\dagger}$	$\rm C+L$	88.7	67.4	72.9	59.0	59.0	59.7	67.8

Table 3: Performances of the BEV map segmentation on the nuScenes validation split. [†]: Fine-tuned with the pre-trained MMIM. *: re-implemented by training from scratch.

3D volume space, substantially reduces information loss, boosting downstream model by 1.9/4.4 NDS/mAP, respectively. The utilization of MMIM also accelerates pre-training convergence, which is also the key to improving downstream task performance.

4.3 Comparison on Downstream Tasks

3D Object Detection To demonstrate the capability of the learned representation, we fine-tune various pre-trained detectors on the nuScenes dataset. As shown in Table 2, our UniM²AE substranitally improves both LiDAR-only and fusion-based detection models. Compared to TransFusion-L-SST, the UniM²AE-L registers a 0.5/0.7 NDS/mAP enhancement on the nuScenes validation subset. In the multi-modality realm, the UniM²AE elevates the outcomes of BEVFusion-SST by 1.2/1.5 NDS/mAP and improves FocalFormer3D [9] by 0.7/0.6 NDS/mAP when MMIM is pre-trained. Of note is that superior results are attained even when a larger voxel size is employed. This is particularly significant given that Transformer-centric strategies (*e.g.* SST [14]) generally trail CNN-centric methodologies (*e.g.* VoxelNet [47]).

BEV Map Segmentation Table 3 presents our BEV map segmentation results on the nuScenes dataset based on BEVFusion [28]. For the camera modality, we outperform the results of training from scratch by 1.7 mAP. In the multimodality setting, the UniM²AE boosts the BEVFusion-SST by 6.5 mAP with pre-trained MMIM and achieve 2.1 improvement over state-of-the-art methods X-Align [3], indicating the effectiveness and strong generalization of our UniM²AE.

Mod Camera	ality LiDAR	Interac BEV	tion Space 3D Volume	mAP	NDS
\checkmark	\checkmark			$59.0 \\ 59.7 \\ 60.1$	61.8 62.6 62.6
\checkmark	\checkmark \checkmark \checkmark	\checkmark	\checkmark	60.7 62.0 62.8	63.1 64.3 65.2

Table 4: Comparison of the input modality and the interaction space.

 Table 5: Ablation experiments to validate our settings. Detection results are reported on the nuScenes validation split

(a) iiij	pact of the m	asking rat		(b) impac	of of th	c numbe	of layer	s along 2-axis.	
Masking Ratio		mAD	NDC	Z-layers	NDS	mAP	FLOPs	GPU memory	
Camera	LiDAR	- IIIAI	NDS	1	56.6	54.9	77G	5.84G	
60%	60%	63.6	66.5	2	58.3	55.0	79G	6.29G	
70%	70%	64.2	67.3	4	58.2	55.0	88G	7.16G	
75%	70%	64.5	67.3	8	58.4	55.0	105G	8.80G	
80%	80%	63.3	66.4						

(b) Impact of the number of layers along z avis

4.4 Ablation Study

(a) Impact of the masking ratio

Multi-modal Pre-training To underscore the importance of the Multi-modal 3D Interaction Module (MMIM) in the 3D volume space for dual modalities, ablation studies were performed with single modal input and were compared with other interaction techniques. As presented in Table 4, the 1st row shows the result of training from scratch. The 2nd and 3rd rows show fine-tuning results with parameters pre-trained on a single modality. The 4th row displays outcomes using the two uni-modal pre-training parameters. The 5th and 6th rows present fine-tuning results with weights pre-trained with different feature spaces for integration. Results reveal that by utilizing MMIM to integrate features from two branches within a unified 3D volume space, the UniM²AE model achieves a remarkable enhancement in performance. Specifically, there's a 3.4 NDS improvement for training from scratch, 2.6 NDS for LiDAR-only and camera-only pre-training, and 2.1 NDS when simply merging the two during the initialization of downstream detectors. Additionally, a noticeable decline in performance becomes evident when substituting the 3D volume space with BEV, as showcased in the concluding rows of Table 4. This drop is likely attributed to features mapped onto BEV losing essential geometric and semantic details, especially along the height axis, resulting in a less accurate representation of an object's true height and spatial positioning. These findings conclusively highlight the crucial role of concurrently integrating camera and LiDAR features within the unified 3D volume space and further validate the effectiveness of the MMIM.

Masking Ratio In Table 5a, an examination of the effects of the masking ratio reveals that optimal performance is achieved with a high masking ratio (70% and 75%). This not only offers advantages in GPU memory savings but also ensures commendable performance. On the other hand, if the masking ratio is set too low or excessively high, there is a notable decline in performance, akin to the results observed with the single-modal MAE.

Number of Layers along z-axis As demonstrated in Table 5b, an increase in the number of layers beyond two does not uniformly enhance performance, yet it invariably escalates the computational cost. This observation can be intuitively understood by considering the typical consistency of road conditions and the specific altitudes where objects are located within a scene(*e.g.* street lights are found above, whereas cars are on the road). Such spatial distribution suggests that a high resolution along the z-axis may not always contribute to a meaning-ful improvement, especially when the additional computational burden is taken into account. Therefore, to strike an optimal balance between computational efficiency and model effectiveness, we opted for two layers along z-axis.

5 Conclusion

The disparity in the multi-modal integration of MAE methods for practical driving sensors is identified. With the introduction of UniM^2AE , a multi-modal self-supervised model is brought forward that adeptly marries image semantics to LiDAR geometries. Two principal innovations define this approach: firstly, the fusion of dual-modal attributes into an augmented 3D volume, which incorporates the height dimension absent in BEV; and secondly, the deployment of the Multi-modal 3D Interaction Module that guarantees proficient cross-modal communications. Benchmarks conducted on the nuScenes Dataset reveal substantial enhancements in 3D object detection by 1.2/1.5 NDS/mAP and in BEV map segmentation by 6.5 mAP, reinforcing the potential of UniM²AE in advancing autonomous driving perception.

Limitations UniM²AE uses a random masking strategy for sensor data from different modalities, which fails to consider their interconnections. Employing simultaneous or complementary masking across multi-modal inputs may enhance feature learning. Moreover, dataset in autonomous driving exhibits significant temporal continuity, with adjacent frames frequently resembling each other. This aspect is omitted in our UniM²AE, leading to pre-training on redundant data, which diminishes pre-training efficiency.

Acknowledgments This work was partially supported by National Key RD Program of China under Grant No. 2021ZD0112100. Guanglei Yang is supported by the Postdoctoral Fellowship Program of CPSF under Grant Number GZC20233458. Tao Luo is supported by the A*STAR Funding under Grant Number C230917003.

References

- Athar, A., Li, E., Casas, S., Urtasun, R.: 4d-former: Multimodal 4d panoptic segmentation. In: Conference on Robot Learning. pp. 2151–2164. PMLR (2023)
- Bai, X., Hu, Z., Zhu, X., Huang, Q., Chen, Y., Fu, H., Tai, C.L.: Transfusion: Robust lidar-camera fusion for 3d object detection with transformers. In: Proceedings of the IEEE/CVF conference on computer vision and pattern recognition. pp. 1090–1099 (2022)
- Borse, S., Klingner, M., Kumar, V.R., Cai, H., Almuzairee, A., Yogamani, S., Porikli, F.: X-align: Cross-modal cross-view alignment for bird's-eye-view segmentation. In: Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision. pp. 3287–3297 (2023)
- Boulch, A., Sautier, C., Michele, B., Puy, G., Marlet, R.: Also: Automotive lidar self-supervision by occupancy estimation. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 13455–13465 (2023)
- Caesar, H., Bankiti, V., Lang, A.H., Vora, S., Liong, V.E., Xu, Q., Krishnan, A., Pan, Y., Baldan, G., Beijbom, O.: nuscenes: A multimodal dataset for autonomous driving. In: Proceedings of the IEEE/CVF conference on computer vision and pattern recognition. pp. 11621–11631 (2020)
- Chen, A., Zhang, K., Zhang, R., Wang, Z., Lu, Y., Guo, Y., Zhang, S.: Pimae: Point cloud and image interactive masked autoencoders for 3d object detection. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 5291–5301 (2023)
- Chen, X., Ma, H., Wan, J., Li, B., Xia, T.: Multi-view 3d object detection network for autonomous driving. In: Proceedings of the IEEE conference on Computer Vision and Pattern Recognition. pp. 1907–1915 (2017)
- Chen, X., Zhang, T., Wang, Y., Wang, Y., Zhao, H.: Futr3d: A unified sensor fusion framework for 3d detection. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 172–181 (2023)
- Chen, Y., Yu, Z., Chen, Y., Lan, S., Anandkumar, A., Jia, J., Alvarez, J.M.: Focalformer3d: focusing on hard instance for 3d object detection. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. pp. 8394–8405 (2023)
- Chen, Y., Li, Y., Zhang, X., Sun, J., Jia, J.: Focal sparse convolutional networks for 3d object detection. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 5428–5437 (2022)
- Chen, Y., Liu, J., Zhang, X., Qi, X., Jia, J.: Largekernel3d: Scaling up kernels in 3d sparse cnns. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 13488–13498 (2023)
- Chi, X., Liu, J., Lu, M., Zhang, R., Wang, Z., Guo, Y., Zhang, S.: Bev-san: Accurate bev 3d object detection via slice attention networks. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 17461– 17470 (2023)
- Fan, H., Su, H., Guibas, L.J.: A point set generation network for 3d object reconstruction from a single image. In: Proceedings of the IEEE conference on computer vision and pattern recognition. pp. 605–613 (2017)
- Fan, L., Pang, Z., Zhang, T., Wang, Y.X., Zhao, H., Wang, F., Wang, N., Zhang, Z.: Embracing single stride 3d object detector with sparse transformer. In: Proceedings of the IEEE/CVF conference on computer vision and pattern recognition. pp. 8458–8468 (2022)

- 16 J. Zou et al.
- He, K., Chen, X., Xie, S., Li, Y., Dollár, P., Girshick, R.: Masked autoencoders are scalable vision learners. In: Proceedings of the IEEE/CVF conference on computer vision and pattern recognition. pp. 16000–16009 (2022)
- Hess, G., Jaxing, J., Svensson, E., Hagerman, D., Petersson, C., Svensson, L.: Masked autoencoder for self-supervised pre-training on lidar point clouds. In: Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision. pp. 350–359 (2023)
- Hu, H., Wang, F., Su, J., Wang, Y., Hu, L., Fang, W., Xu, J., Zhang, Z.: Ea-lss: Edge-aware lift-splat-shot framework for 3d bev object detection. arXiv preprint arXiv:2303.17895 (2023)
- Huang, L., You, S., Zheng, M., Wang, F., Qian, C., Yamasaki, T.: Green hierarchical vision transformer for masked image modeling. Advances in Neural Information Processing Systems 35, 19997–20010 (2022)
- Huang, T., Dong, B., Yang, Y., Huang, X., Lau, R.W., Ouyang, W., Zuo, W.: Clip2point: Transfer clip to point cloud classification with image-depth pretraining. arXiv preprint arXiv:2210.01055 (2022)
- Jiao, Y., Jie, Z., Chen, S., Chen, J., Ma, L., Jiang, Y.G.: 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. pp. 21643–21652 (2023)
- Li, H., Sima, C., Dai, J., Wang, W., Lu, L., Wang, H., Zeng, J., Li, Z., Yang, J., Deng, H., et al.: Delving into the devils of bird's-eye-view perception: A review, evaluation and recipe. IEEE Transactions on Pattern Analysis and Machine Intelligence (2023)
- 22. Li, Y., Yu, A.W., Meng, T., Caine, B., Ngiam, J., Peng, D., Shen, J., Lu, Y., Zhou, D., Le, Q.V., 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. pp. 17182–17191 (2022)
- Li, Z., Wang, W., Li, H., Xie, E., Sima, C., Lu, T., Qiao, Y., Dai, J.: Bevformer: Learning bird's-eye-view representation from multi-camera images via spatiotemporal transformers. In: European conference on computer vision. pp. 1–18. Springer (2022)
- Liang, M., Yang, B., Wang, S., Urtasun, R.: Deep continuous fusion for multi-sensor 3d object detection. In: Proceedings of the European conference on computer vision (ECCV). pp. 641–656 (2018)
- Liang, T., Xie, H., Yu, K., Xia, Z., Lin, Z., Wang, Y., Tang, T., Wang, B., Tang, Z.: Bevfusion: A simple and robust lidar-camera fusion framework. Advances in Neural Information Processing Systems 35, 10421–10434 (2022)
- Liu, H., Cai, M., Lee, Y.J.: Masked discrimination for self-supervised learning on point clouds. In: European Conference on Computer Vision. pp. 657–675. Springer (2022)
- Liu, Z., Lin, Y., Cao, Y., Hu, H., Wei, Y., Zhang, Z., Lin, S., Guo, B.: Swin transformer: Hierarchical vision transformer using shifted windows. In: Proceedings of the IEEE/CVF international conference on computer vision. pp. 10012–10022 (2021)
- Liu, Z., Tang, H., Amini, A., Yang, X., Mao, H., Rus, D.L., Han, S.: Bevfusion: Multi-task multi-sensor fusion with unified bird's-eye view representation. In: 2023 IEEE International Conference on Robotics and Automation (ICRA). pp. 2774– 2781. IEEE (2023)
- Min, C., Zhao, D., Xiao, L., Nie, Y., Dai, B.: Voxel-mae: Masked autoencoders for pre-training large-scale point clouds. arXiv preprint arXiv:2206.09900 (2022)

- Nabati, R., Qi, H.: Centerfusion: Center-based radar and camera fusion for 3d object detection. In: Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision. pp. 1527–1536 (2021)
- Pang, Y., Wang, W., Tay, F.E., Liu, W., Tian, Y., Yuan, L.: Masked autoencoders for point cloud self-supervised learning. In: European conference on computer vision. pp. 604–621. Springer (2022)
- Philion, J., Fidler, S.: Lift, splat, shoot: Encoding images from arbitrary camera rigs by implicitly unprojecting to 3d. In: Computer Vision–ECCV 2020: 16th European Conference, Glasgow, UK, August 23–28, 2020, Proceedings, Part XIV 16. pp. 194–210. Springer (2020)
- Radford, A., Kim, J.W., Hallacy, C., Ramesh, A., Goh, G., Agarwal, S., Sastry, G., Askell, A., Mishkin, P., Clark, J., et al.: Learning transferable visual models from natural language supervision. In: International conference on machine learning. pp. 8748–8763. PMLR (2021)
- 34. Roddick, T., Kendall, A., Cipolla, R.: Orthographic feature transform for monocular 3d object detection. arXiv preprint arXiv:1811.08188 (2018)
- Sindagi, V.A., Zhou, Y., Tuzel, O.: Mvx-net: Multimodal voxelnet for 3d object detection. In: 2019 International Conference on Robotics and Automation (ICRA). pp. 7276–7282. IEEE (2019)
- 36. Sun, P., Kretzschmar, H., Dotiwalla, X., Chouard, A., Patnaik, V., Tsui, P., Guo, J., Zhou, Y., Chai, Y., Caine, B., et al.: Scalability in perception for autonomous driving: Waymo open dataset. In: Proceedings of the IEEE/CVF conference on computer vision and pattern recognition. pp. 2446–2454 (2020)
- Tian, X., Jiang, T., Yun, L., Wang, Y., Wang, Y., Zhao, H.: Occ3d: A largescale 3d occupancy prediction benchmark for autonomous driving. arXiv preprint arXiv:2304.14365 (2023)
- Tian, X., Ran, H., Wang, Y., Zhao, H.: Geomae: Masked geometric target prediction for self-supervised point cloud pre-training. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 13570–13580 (2023)
- Vora, S., Lang, A.H., Helou, B., Beijbom, O.: Pointpainting: Sequential fusion for 3d object detection. In: Proceedings of the IEEE/CVF conference on computer vision and pattern recognition. pp. 4604–4612 (2020)
- Wang, C., Ma, C., Zhu, M., Yang, X.: Pointaugmenting: Cross-modal augmentation for 3d object detection. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 11794–11803 (2021)
- 41. Wang, H., Tang, H., Shi, S., Li, A., Li, Z., Schiele, B., Wang, L.: Unitr: A unified and efficient multi-modal transformer for bird's-eye-view representation. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. pp. 6792–6802 (2023)
- Wei, Y., Zhao, L., Zheng, W., Zhu, Z., Zhou, J., Lu, J.: Surroundocc: Multi-camera 3d occupancy prediction for autonomous driving. arXiv preprint arXiv:2303.09551 (2023)
- Xia, Z., Pan, X., Song, S., Li, L.E., Huang, G.: Vision transformer with deformable attention. In: Proceedings of the IEEE/CVF conference on computer vision and pattern recognition. pp. 4794–4803 (2022)
- 44. Xie, E., Yu, Z., Zhou, D., Philion, J., Anandkumar, A., Fidler, S., Luo, P., Alvarez, J.M.: M²bev: Multi-camera joint 3d detection and segmentation with unified birdseye view representation. arXiv preprint arXiv:2204.05088 (2022)
- Xie, S., Kong, L., Zhang, W., Ren, J., Pan, L., Chen, K., Liu, Z.: Robobev: Towards robust bird's eye view perception under corruptions. arXiv preprint arXiv:2304.06719 (2023)

- 18 J. Zou et al.
- 46. Xu, R., Wang, T., Zhang, W., Chen, R., Cao, J., Pang, J., Lin, D.: Mv-jar: Masked voxel jigsaw and reconstruction for lidar-based self-supervised pre-training. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 13445–13454 (2023)
- Yan, Y., Mao, Y., Li, B.: Second: Sparsely embedded convolutional detection. Sensors 18(10), 3337 (2018)
- 48. Yang, C., Chen, Y., Tian, H., Tao, C., Zhu, X., Zhang, Z., Huang, G., Li, H., Qiao, Y., Lu, L., et al.: Bevformer v2: Adapting modern image backbones to bird's-eyeview recognition via perspective supervision. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 17830–17839 (2023)
- 49. Yang, H., He, T., Liu, J., Chen, H., Wu, B., Lin, B., He, X., Ouyang, W.: Gd-mae: generative decoder for mae pre-training on lidar point clouds. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 9403–9414 (2023)
- Yang, Z., Chen, J., Miao, Z., Li, W., Zhu, X., Zhang, L.: Deepinteraction: 3d object detection via modality interaction. Advances in Neural Information Processing Systems 35, 1992–2005 (2022)
- 51. Yao, L., Huang, R., Hou, L., Lu, G., Niu, M., Xu, H., Liang, X., Li, Z., Jiang, X., Xu, C.: Filip: Fine-grained interactive language-image pre-training. arXiv preprint arXiv:2111.07783 (2021)
- Yin, T., Zhou, X., Krahenbuhl, P.: Center-based 3d object detection and tracking. In: Proceedings of the IEEE/CVF conference on computer vision and pattern recognition. pp. 11784–11793 (2021)
- Yin, T., Zhou, X., Krähenbühl, P.: Multimodal virtual point 3d detection. Advances in Neural Information Processing Systems 34, 16494–16507 (2021)
- Yu, X., Tang, L., Rao, Y., Huang, T., Zhou, J., Lu, J.: 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. pp. 19313– 19322 (2022)
- 55. Zhang, R., Guo, Z., Zhang, W., Li, K., Miao, X., Cui, B., Qiao, Y., Gao, P., Li, H.: Pointclip: Point cloud understanding by clip. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 8552–8562 (2022)
- 56. Zhang, R., Wang, L., Qiao, Y., Gao, P., Li, H.: Learning 3d representations from 2d pre-trained models via image-to-point masked autoencoders. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 21769–21780 (2023)
- 57. Zhang, Y., Zhu, Z., Du, D.: Occformer: Dual-path transformer for vision-based 3d semantic occupancy prediction. arXiv preprint arXiv:2304.05316 (2023)
- Zhou, B., Krähenbühl, P.: Cross-view transformers for real-time map-view semantic segmentation. In: Proceedings of the IEEE/CVF conference on computer vision and pattern recognition. pp. 13760–13769 (2022)
- Zhu, X., Su, W., Lu, L., Li, B., Wang, X., Dai, J.: Deformable detr: Deformable transformers for end-to-end object detection. arXiv preprint arXiv:2010.04159 (2020)