

MatrixVT: Efficient Multi-Camera to BEV Transformation for 3D Perception

Hongyu Zhou, Zheng Ge, Zeming Li, Xiangyu Zhang*
 MEGVII Technology

{zhouhongyu, gezheng, lizeming, zhangxiangyu}@megvii.com

Abstract

This paper proposes an efficient multi-camera to Bird’s-Eye-View (BEV) view transformation method for 3D perception, dubbed MatrixVT. Existing view transformers either suffer from poor efficiency or rely on device-specific operators, hindering the broad application of BEV models. In contrast, our method generates BEV features efficiently with only convolutions and matrix multiplications (MatMul). Specifically, we propose describing the BEV feature as the MatMul of image feature and a sparse Feature Transporting Matrix (FTM). A Prime Extraction module is then introduced to compress the dimension of image features and reduce FTM’s sparsity. Moreover, we propose the Ring & Ray Decomposition to replace the FTM with two matrices and reformulate our pipeline to reduce calculation further. Compared to existing methods, MatrixVT enjoys a faster speed and less memory footprint while remaining deploy-friendly. Extensive experiments on nuScenes and Waymo benchmarks demonstrate that our method is highly efficient but obtains results on par with the SOTA method in object detection and map segmentation tasks.

1. Introduction

Vision-centric 3D perception in Bird’s-Eye-View (BEV) [17, 20, 24] has recently drawn extensive attention. Apart from their outstanding performance, the compact and unified feature representation in BEV facilitates straight-forward feature fusions [1, 5, 12, 10], and enables various downstream tasks (e.g. object detection [6, 9, 10], map segmentation [17, 30], motion planning, etc.) to be applied thereon easily.

View Transformation (VT) is the key component that converts multi-camera features to BEV, which has been heavily studied in previous works [15, 20, 17, 24, 10, 6, 21, 9]. Existing VT methods can be categorized into *geometry-based* [13, 19, 17, 9, 18, 30] and *learning-based* meth-

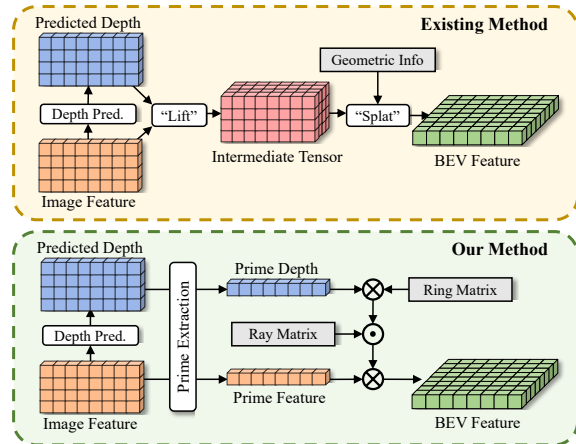


Figure 1. The pipeline of Lift-Splat (upper) and our proposed MatrixVT (lower). We compress image features before VT to reduce memory footprint and calculation. Note that only standard operators are adopted in MatrixVT.

ods [10, 20, 15]. Among these two categories, geometry-based methods show superior performance due to the use of geometric constraints. *Lift-Splat* [17], as a representative geometry-based VT, predicts *categorical depth distribution* for each pixel and “lift” the corresponding features into 3D space according to the predicted depth. These feature vectors are then *gathered* into pre-defined grids on a reference BEV plane (*i.e.*, “splat”) to form the BEV feature (Fig. 1, upper). The Lift-Splat-based VT has shown great potential to produce high-quality BEV features, achieving remarkable performance on object detection [9, 5] and map segmentation tasks.

Despite the effectiveness of Lift-Splat-like VT [9], two issues remain. First, the “splat” operation is not universally feasible. Existing implementations of “splat” relies on either the “cumsum trick” [17, 6] that is highly inefficient, or customized operators [9] that can only be used on specific devices, making application of BEV perception methods costly. Second, the size of “lifted” multi-view image features is *huge*, becoming the memory bottleneck of BEV models. These two issues lead to a heavy burden on BEV methods during both the training and inference phases. As

*Corresponding author. This research was supported by National Key R&D Program of China (No. 2017YFA0700800) and Beijing Academy of Artificial Intelligence (BAAI).

a result, the drawbacks of existing view transformers limit the broad application of autonomous driving technology

In this work, we propose a novel VT method to address the above problems. MatrixVT is proposed based on the fact that the VT can be viewed as a *feature transportation* process. In that case, the BEV feature can be viewed as the MatMul between the “lifted” feature and a transporting matrix, namely *Feature Transporting Matrix (FTM)*. We thus generalize the Lift-Splat VT into a simple form and eliminate specialized operators.

However, transformation with FTM is a kind of degradation — the mapping between the 3D space and BEV grids is *extreme sparse*, leading to the huge size of FTM and poor efficiency. Prior works [9, 18] seek customized operators, successfully avoiding such sparsity. In this paper, we argue that there are other solutions to the problem of sparse mapping. First, we propose Prime Extraction. Motivated by an observation that the height (vertical) dimension of images is *less* informative in autonomous driving (see Sec. 3.2), we compress the image features along this dimension before VT. Second, we adopt matrix decomposition to reduce the sparsity of FTM. The proposed Ring & Ray Decomposition orthogonally decomposes the FTM into two separate matrices, each encoding the distance and direction of the ego-centric polar coordinate. This decomposition also allows us to reformulate our pipeline into a *mathematically equivalent* but more efficient one (Fig. 1, lower). These two techniques reduce memory footprint and calculation during VT by hundreds of times, enabling MatrixVT to be more efficient than existing methods.

The proposed MatrixVT inherits the advantages of the Lift-Splat [17] paradigm while being much more memory efficient and fast. Extensive experimental results show that MatrixVT is 2-to-8 times faster than previous methods [9, 6] and saves up to 97% memory footprint among different settings. Meanwhile, the perception model with MatrixVT achieves 46.6% mAP and 56.2% NDS for object detection and 46.2% mIoU for vehicle segmentation on the nuScenes [2] val set, which is comparable to the state-of-the-art performance [9]. We conclude our main contributions as follows:

- We propose a new description of multi-camera to BEV transformation — using the Feature Transportation Matrix (FTM), which is a more general representation.
- To solve the sparse mapping problem, we propose Prime Extraction and the Ring & Ray Decomposition, boosting VT with FTM by a huge margin.
- Extensive experiments demonstrate that MatrixVT yields comparable performance to the state-of-the-art method on the nuScenes object detection and map segmentation tasks while being more efficient and generally applicable.

2. Related Works

2.1. Visual 3D Perception

Camera-based perception [28, 27, 6, 10, 11] is the most commonly used method for various scenarios due to its low cost and high accessibility. Comparing with 2D perception (object detection [26], semantic segmentation [22], *etc.*), 3D perception requires additional prediction of the depth information which is a naturally ill-posed problem [21].

Existing works either predict the depth information explicitly or implicitly. FCOS3D [28] simply extends the structure of the classic 2D object detector [26], predicting pixel-wise depth explicitly using an extra sub-net. CaDDN [18] proposes treating depth prediction as a classification task rather than a regression task, and projecting image feature into Bird’s-Eye-View (BEV) space to achieve unified modeling of detection and depth prediction. BEVDepth [9] proposes a heavy and well-designed subnet for depth prediction, achieving outstanding performance among other 3D object detectors. Meanwhile, methods like PON [20], BirdGAN [24], and EFT [23] use fully learnable networks to implement feature transformation, learning object depth implicitly using segmentation or detection supervision. Among them, EFT [23] proposes using a learnable matrix for VT, which bears some semblance, in form, to the Feature Transformation Matrix (FTM) we proposed. However, the matrices we designed for MatrixVT contains clear physical meanings with explicit depth modeling, and there are significant differences in their applicable scenarios and methods of use.

2.2. Perception in Bird’s-Eye-View

The concept of BEV is firstly proposed for processing LiDAR point cloud [33, 7, 31, 32], and found effective for fusing multi-view image features. The core component of vision-based BEV paradigm is the view-transformation. OFT [21] firstly propose mapping image feature from Perspective View into BEV using camera parameters. This method projects a reference point from a BEV grid to the image plane, and sample corresponding features back to the BEV grid. Following this work, BEVFormer [10] proposes using Attention module to sample features around the reference point. These methods fail to distinguish BEV grids that are projected to the same position on the image plane, thus showing inferior performance than depth-based methods.

Depth-based methods, represented by LSS [17] and CaDDN [18], predict categorical depth for each pixel, the extracted image feature on a specific pixel is then projected into 3D space by doing per-pixel outer product with corresponding depth. The projected high-dimensional tensor is then “collapsed” to a BEV reference plane using convolution [18], Pillar Pooling [17, 6], or Voxel Pooling [9]. Lift-

Splat based methods [17, 9, 6] show outstanding performance for downstream tasks but introduce two extra problems. Firstly, the intermediate representation of image features is large and in-efficient, making training and application of these methods difficult. Secondly, the Pillar Pooling [17] introduces random memory access, which is slow and device-demanding (extremely slow on general-purpose devices). In this paper, we propose a new depth-based view transformation to overcome these problems while retaining the ability to produce high-quality BEV features.

3. MatrixVT

The MatrixVT is a simple view transformer based on the depth-based VT paradigm. In Sec. 3.1, we first revisit existing Lift-Splat transformation [17, 9] and introduce the concept of Feature Transporting Matrix (FTM) together with the sparse mapping problem. Then, techniques proposed to solve the problem of sparse mapping *i.e.*, Prime Extraction (Sec. 3.2) and Ring & Ray Decomposition (Sec. 3.3), are introduced. In Sec. 3.4, we designate the novel VT method utilizing the aforementioned techniques as MatrixVT, and elaborate its overall pipeline.

For clarity, we use letters of the normal script (*e.g.* F) for tensors, the Roman script for matrices (*e.g.* M), and Bold script (*e.g.* v) for vectors. Besides, $A \cdot B$ denotes Matrix Multiplication (MatMul), $A \odot B$ denotes Hadamard Product, and $a \otimes b$ denotes outer product.

3.1. Background

Lift-Splat-like methods exploit the idea of pseudo-LiDAR [29] and LiDAR-based BEV methods [33, 7] for View Transformation. Image features are “lifted” to 3D space and processed like point cloud. We define N_c as the number of cameras for a specific scene; N_d as the number of depth bins; H_I and W_I to be the height and width of image features; H_B and W_B be the height and width of BEV features (*i.e.*, shape of BEV grids); C to be the number of feature channels. Consequently, let multi-view image features to be $F \in \mathbb{R}^{N_c \times H_I \times W_I \times C}$; categorical depth to be $D \in \mathbb{R}^{N_c \times H_I \times W_I \times N_d}$; BEV features to be $F_{BEV} \in \mathbb{R}^{H_B \times W_B \times C}$ (initialized to zeros). The Lift-Splat VT first “lift” the F into 3D space by doing per-pixel outer product with D , obtaining the high-dimensional intermediate tensor $F_{inter} \in \mathbb{R}^{N_c \times H_I \times W_I \times N_d \times C}$.

$$F_{inter} = \{F_{inter}^{h,w}\}_{H_I, W_I} = \{F^{h,w} \cdot (D^{h,w})^T\}_{H_I, W_I} \quad (1)$$

The intermediate tensor can be treated as $N_c \times H_I \times W_I \times N_d$ feature vectors, each vector corresponding to a geometric coordinate. The “splat” operation is then adopted using operators like Pillar Pooling [17], during which each feature vector is summed to a BEV grid according to geometric coordinates.

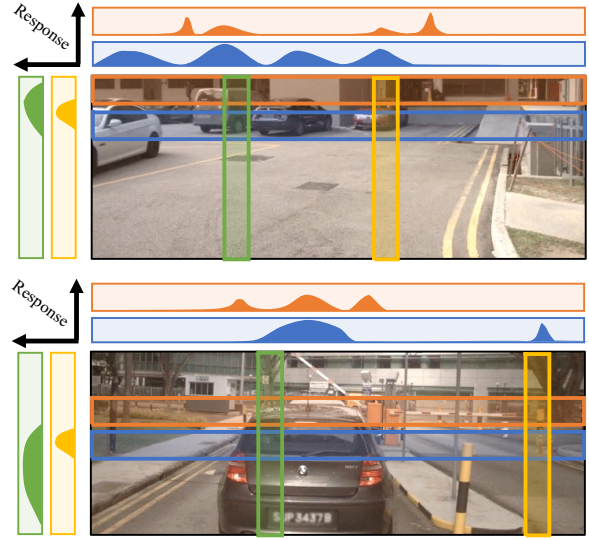


Figure 2. Response strengths along vertical and horizontal dimensions. For autonomous driving, the variance of response on the width dimension is higher than height dimension.

We find that the “splat” operation is a fixed mapping between the intermediate tensor and BEV grids. Therefore, we use a Feature Transporting Matrix $M_{FT} \in \mathbb{R}^{(H_B \times W_B) \times (N_c \times H_I \times W_I \times N_d)}$ to describe this mapping. The FTM can be strictly equal to Pillar Pooling in LSS [17] or encodes different sampling rules (*i.e.*, Gaussian sampling). The FTM enables feature transportation with Matrix Multiplication (MatMul):

$$F_{BEV} = M_{FT} \cdot F_{inter} \quad (2)$$

where $F_{inter} \in \mathbb{R}^{(N_c \times H_I \times W_I \times N_d) \times C}$ is the “lifted” feature F_{inter} in matrix format, $F_{BEV} \in \mathbb{R}^{(H_B \times W_B) \times C}$ is the BEV feature F_{BEV} in matrix format.

Replacing the “splat” operation with FTM eliminates the need for customized operators. However, the sparse mapping between the 3D space and BEV grids can lead to a *massive and highly sparse* FTM, which harms the efficiency of matrix-based VT. To address the sparse mapping problem without customized operators, we propose two techniques to reduce the sparsity of FTM and speed up the transformation.

3.2. Prime Extraction for Autonomous Driving

The high-dimensional intermediate tensor is the primary cause of the sparse mapping and the low efficiency of Lift-Splat-like VT. An intuitive way to reduce sparsity is reducing the size of F_{inter} . In the context of autonomous driving, a notable observation can lead to feature compression. As can be seen in Fig. 2, the height dimension of camera images contains less informative data. The height dimension is filled with redundant and background features, making it

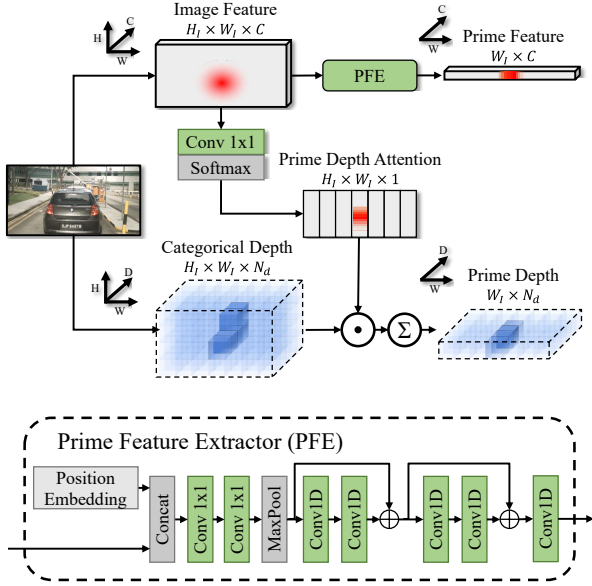


Figure 3. The Prime Extraction module. The categorical depth is reduced guided by Prime Depth Attention, while the image feature is reduced by Prime Feature Extractor (PFE), which consists of MaxPooling and convolutions for refinement.

ideal for feature compression. Therefore, we propose Prime Extraction, which is a compression technique that can be used in autonomous driving and other scenarios with information redundancy of this nature.

As shown in Fig. 3, the Prime Extraction module predicts Prime Depth Attention for each direction (a column of the tensor) guided by the image feature. It generates the Prime Depth as a *weighted sum* of categorical depth distributions. Meanwhile, the Prime Feature is obtained by the Prime Feature Extractor (PFE), which consists of position embedding, column-wise max-pooling, and convolutions for refinement. By applying Prime Extraction to matrix-based VT, we successfully reduce the FTM to $M'_{FT} \in \mathbb{R}^{(H_B \times W_B) \times (N_c \times W_I \times N_d)}$, which is H_I times smaller than the raw matrix. The pipeline of generating the BEV feature using the Prime Feature and the Prime Depth is demonstrated in Fig. 4 (yellow box).

The prime extraction is motivated by an observation: The image feature’s height dimension has a lower response variance than the width dimension. This observation indicates that this dimension contains less information than the width dimension. We thus propose to compress image features on the height dimension. Previous works have also exploited reducing the height dimension of image features [20, 1], but we firstly propose compressing both image features and corresponding depth to boost VT.

In Sec. 4.4, we will show that the extracted Prime Feature and Prime Depth effectively retain valuable information from the raw feature and produce BEV features of the

same high quality as the raw feature. Moreover, the Prime Extraction technique can be individually adapted to existing Lift-Splat-like VTs to enhance their efficiency at almost no performance cost.

3.3. “Ring and Ray” Decomposition

The sparsity of FTM can be further reduced by matrix decomposition. To this end, we propose the “Ring & Ray” Decomposition. Without loss of generality, we can set the N_c to 1 and view M'_{FT} as a tensor M'_{FT} . In that case, the shape of M'_{FT} would be $H_B \times W_B \times W_I \times N_d$. We note that its dimension of size W_I can be viewed as the *direction* in a *polar coordinate*, since each column of the image feature represents information of a specific direction. Likewise, the dimension of size N_d can be viewed as the *distance* in the polar coordinate. In other words, the image feature required for a specific BEV grid can be located by direction and distance. We thus propose to orthogonally decompose the M'_{FT} into two separate matrices, each encoding the directional or distance information. Specifically, we use a Ring Matrix $M_{\text{Ring}} \in \mathbb{R}^{N_d \times (H_B \times W_B)}$ to encode distance information and a Ray Matrix $M_{\text{Ray}} \in \mathbb{R}^{W_I \times (H_B \times W_B)}$ to encode directional information (see Appendix 1.1 for pseudo code). The Ring & Ray Decomposition effectively reduces the size of static parameters. The number of predefined parameters (size of FTM) is reduced from $W_I \times N_d \times H_B \times W_B$ to $(W_I + N_d) \times H_B \times W_B$, which is typically 30 to 50 times less¹.

Then we show how to use these two matrices for VT. Given the Prime Feature and Prime Depth, we first do the per-pixel outer product of them as in Lift-Splat to obtain the “lifted feature” $F_{\text{inter}} \in \mathbb{R}^{W_I \times N_d \times C}$.

$$F_{\text{inter}} = \{F_{\text{inter}}^w\}_{W_I} = \{F^w \otimes (D^w)^T\}_{W_I} \quad (3)$$

Then, as illustrated in Fig. 4, instead of using M'_{FT} to directly transform it into a BEV feature (yellow box), we transpose F_{inter} , view it as a matrix $F_{\text{inter}} \in \mathbb{R}^{N_d \times (W_I \times C)}$, and do MatMul between the Ring Matrix and F_{inter} to obtain an intermediate feature. The intermediate feature is then masked by doing a Hadamard Product with the Ray Matrix, summed on the dimension of W_I to obtain the BEV feature. (Fig. 4, blue box, summation omitted).

$$F_{\text{inter}}^* = M_{\text{Ray}} \odot (M_{\text{Ring}} \cdot F_{\text{inter}}) \quad (4)$$

$$F_{\text{BEV}} = \sum_w F_{\text{inter}}^{*w} \quad (5)$$

where $F_{\text{inter}}^* \in \mathbb{R}^{W_I \times C \times H_B \times W_B}$ is F_{inter}^* in tensor form.

However, this decomposition do not reduce the FLOPs during VT and introduces the Intermediate Feature in Fig. 4 (blue) whose size if huge and depends on feature channel

¹Under feature width 44 and 88, with 112 depth bins.

Pipelines of View Transformation Using Matrices

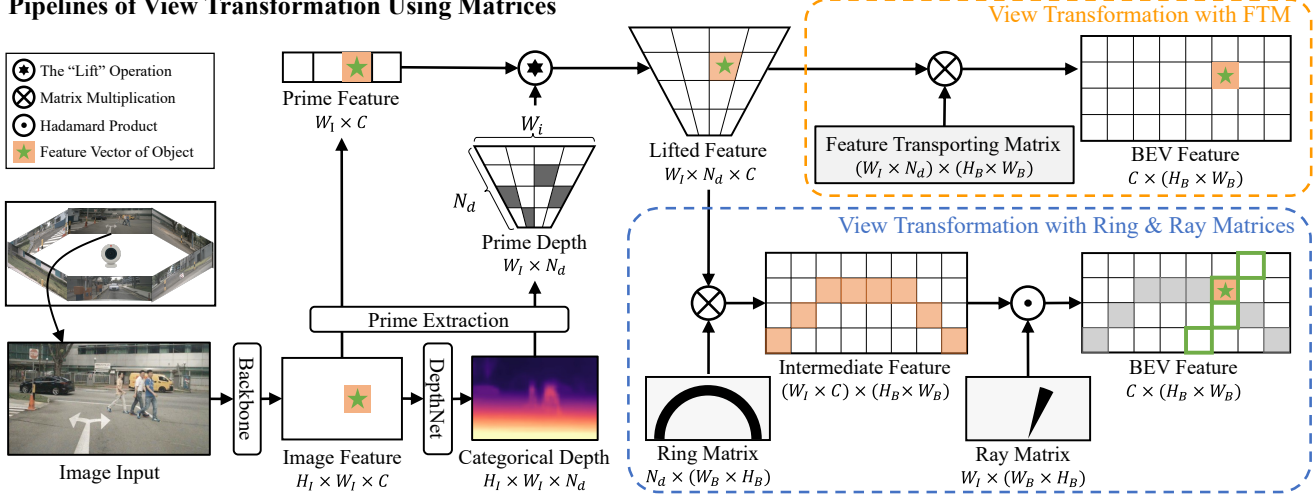


Figure 4. View Transformation using matrices. We first extract features from the image and predict categorical depth for each pixel. The obtained feature and depth are sent to the Prime Extraction module for compression. With the compressed Prime Feature and the Prime Depth, we “lift” the 1D Prime Feature into 2D space using Prime Depth. The “lifted feature” is then transformed into the BEV feature by MatMul with FTM or the decomposed Ring & Ray matrices.

C. To reduce the calculation and memory footprint during VT, we combine Eq. 3 to Eq. 5 and rewrite them in a mathematically equivalent form (see Appendix 1.2 for proof):

$$F_{\text{BEV}} = F \cdot (M_{\text{Ray}} \odot (M_{\text{Ring}} \cdot D)) \quad (6)$$

With this reformulation, we reduce the calculation during VT from $2 \times W_I \times C \times N_d \times H_B \times W_B$ FLOPs to $2(C + N_d + 1) \times W_I \times H_B \times W_B$ FLOPs; the memory footprint is also reduced from $W_I \times N_d \times H_B \times W_B$ to $(W_I + N_d) \times H_B \times W_B$.

3.4. Overall Pipeline

With above techniques, we reduce the calculation and memory footprint using FTM by hundreds of times, making VT with FTM not only feasible but also efficient. Given the multi-view images for a specific scene, we conclude the overall pipeline of MatrixVT as follows:

1. We first use an image backbone to extract image features from each image.
2. Then, a depth predictor is adopted to predict categorical depth distribution for each feature pixel to obtain the depth prediction.
3. After that, we send each image feature and corresponding depth to the Prime Extraction module, obtaining the Prime Feature and the Prime Depth, which is the compressed feature and depth.
4. Finally, with the Prime Feature, Prime Depth, and the pre-defined Ring & Ray Matrices, we use Eq. 6 (see also Fig. 1, lower) to obtain the BEV feature.

4. Experiments

In this section, we compare the performance, latency, and memory footprint of MatrixVT and other existing VT methods on the nuScenes benchmark [2].

4.1. Implementation Details

We conduct our experiments based on BEVDepth [9], the current state-of-the-art detector on the nuScenes benchmark. In order to conduct a fair comparison of performance and efficiency, we re-implement BEVDepth according to their paper. Unless otherwise specified, we use ResNet-50 [3] and VoVNet2-99 [8] pre-trained on DD3D [16] as the image backbone and SECOND FPN [31] as the image neck and BEV neck. The input image adopts pre-processing and data augmentations same as in [9]. We use BEV feature size 128×128 for low input resolution and 256×256 for high resolution on detection. Segmentation experiments use 200×200 BEV resolution as in LSS [17]. We use the DepthNet [9] to predict categorical depth from 2m to 58m in nuScenes, with uniform 112 division. During training, CBGS [34] and model EMA are adopted. Models are trained to converge since MatrixVT converges a little slower than other methods, but no more than 30 epochs.

4.2. Comparison of Performances

4.2.1 Object Detection

We conduct experiments under several settings to evaluate the performance of MatrixVT in Tab. 1 and Tab. 2. On the nuScenes benchmark, we first adopt the ResNet family as the backbone without applying multi-frame fusion.

Method	Backbone	Resolution	MF	mAP \uparrow	mATE \downarrow	mASE \downarrow	mAOE \downarrow	mAVE \downarrow	mAAE \downarrow	NDS \uparrow
BEVDet [6]	Res-50	256 \times 704	\times	0.286	0.724	0.278	0.590	0.873	0.247	0.372
PETR [11]	Res-50	384 \times 1056	\times	0.313	0.768	0.278	0.564	0.923	0.225	0.381
BEVDepth [9]	Res-50	256 \times 704	\times	0.337	0.646	0.271	0.574	0.838	0.220	0.414
MatrixVT	Res-50	256 \times 704	\times	0.336	0.653	0.271	0.473	0.903	0.231	0.415
BEVDet [6]	Res-101	384 \times 1056	\times	0.330	0.702	0.272	0.534	0.932	0.251	0.396
FCOS3D [28]	Res-101	900 \times 1600	\times	0.343	0.725	0.263	0.422	1.292	0.153	0.415
PETR [11]	Res-101	512 \times 1408	\times	0.357	0.710	0.270	0.490	0.885	0.224	0.421
BEVFormer [10]	R101-DCN	900 \times 1600	\times	0.375	0.725	0.272	0.391	0.802	0.200	0.448
MatrixVT	Res-101	512 \times 1408	\times	0.396	0.577	0.261	0.397	0.870	0.207	0.467
BEVFormer [10]	R101-DCN	900 \times 1600	\checkmark	0.416	0.673	0.274	0.372	0.394	0.198	0.517
BEVDet4D [5]	Swin-B	640 \times 1600	\checkmark	0.421	0.579	0.258	0.329	0.301	0.191	0.545
BEVDepth [9]	V2-99	512 \times 1408	\checkmark	0.464	0.528	0.254	0.350	0.354	0.198	0.564
MatrixVT	V2-99	512 \times 1408	\checkmark	0.466	0.535	0.260	0.380	0.342	0.198	0.562

Table 1. Experimental results of object detection on the nuScenes val set. The “MF” indicates multi-frame fusion.

Method	Backbone	mAP \uparrow	mAPH \uparrow	mAPL \uparrow
BEVDepth [9]	Res-50	0.290	0.208	0.253
MatrixVT	Res-50	0.304	0.215	0.259

Table 2. Performance of object detection on Waymo val set.

MatrixVT achieves 33.6% and 49.7% mAP with ResNet-50 and ResNet-101, which is comparable to the BEVDepth [9] and surpasses other methods by a large margin. We then test the upper bound of MatrixVT by replacing the backbone with V2-99 [8] pre-trained on external data and applying multi-frame fusion. Under this setting, MatrixVT achieves 46.6% mAP and 56.2% NDS, which is also comparable to the BEVDepth. Further, we validate if MatrixVT can be generalized to other benchmark. We conduct experiments on the Waymo benchmark [25], results show that MatrixVT still showing comparable or even better performance under this setting.

4.2.2 Map Segmentation

We also conduct experiments on map segmentation tasks to validate the quality of the BEV feature generated by MatrixVT. To achieve this, we simply put a U-Net-like [22] segmentation head on the BEV feature. For fair comparison, we put the same head on the BEV feature of BEVDepth for experiments, and results are reported as “BEVDepth-Seg” in Tab. 3. It is worth noting that previous works achieve the best segmentation performance under different settings (different resolution, head structure, *etc.*); thus, we report the highest performance of each method. As can be seen from Tab. 3, the map segmentation performance of MatrixVT surpasses most existing methods on all three sub-tasks and is comparable to our baseline, BEVDepth.

Method	IoU-Drive \uparrow	IoU-Lane \uparrow	IoU-Vehicle \uparrow
LSS [17]	0.729	0.200	0.321
FIERY [4]	-	-	0.382
M ² BEV [30]	0.759	0.380	-
BEVFormer [10]	0.775	0.239	0.467
BEVDepth-Seg	0.827	0.464	0.450
MatrixVT	0.835	0.448	0.462

Table 3. Experimental results of BEV segmentation on the nuScenes val set. We implement map segmentation on the BEVDepth by placing a segmentation head on the BEV feature.

4.3. Efficient Transformation

We compare the efficiency of View Transformation in two dimensions: Speed and Memory Consumption. Note that we measure the latency and memory footprint (using fp32) of *view transformers* since these metrics are affected by backbone and head design. We take the CPU as a representative general-purpose device where customized operators are unavailable. For a fair comparison, we measure and compare the other two Lift-Splat-like view transformers. The LS-BEVDet is the accelerated transformation used in BEVDet (with default parameters); the LS-BEVDepth uses the CUDA [14] operator proposed in BEVDepth [9] that is not available on CPU and other platforms. To demonstrate the characteristics of different methods, we define six transformation settings, namely S1~S6, varies in image feature size, BEV feature size, and feature channels that are closely related to model performance.

As can be seen in Fig. 5, the proposed MatrixVT boost the transformation significantly on the CPU, being 4 to 8 times faster than the LS-BEVDet[6]. On the CUDA platform [14], where customized operators enable faster transformation, MatrixVT still shows a much faster speed than

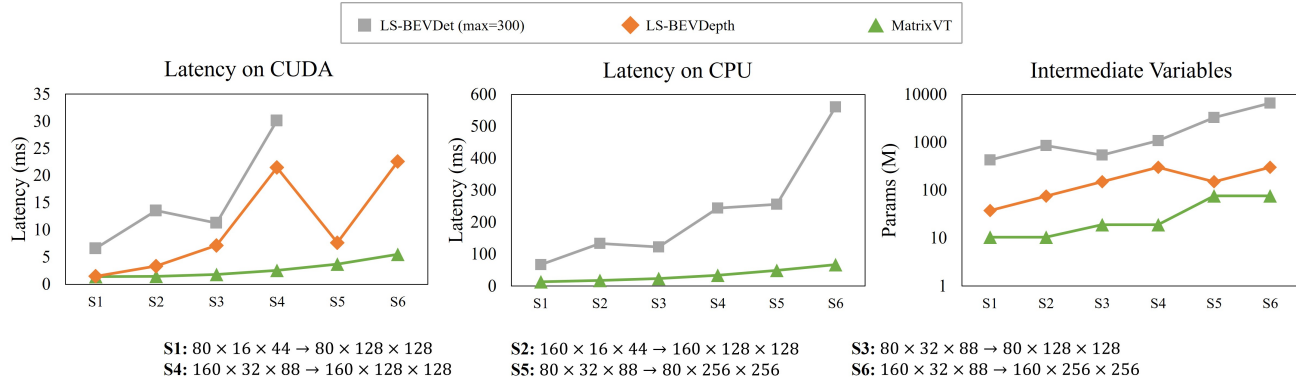


Figure 5. Latency (left, middle) and the number of intermediate parameters (right) of MatrixVT and two implementations of Lift-Splat, the S1~S6 denote transformation settings represented by image feature size ($C \times H_I \times W_I$) to BEV feature size ($C \times H_B \times W_B$). Note that the LS-BEVDet raises Out Of Memory error under S5 and S6 on CUDA. The LS-BEVDet is not available without the CUDA platform.

the LS-BEVDet under most settings. Besides, we calculate the number of intermediate variables during VT as an indicator of extra memory footprint. For MatrixVT, these variables include the Ring Matrix, Ray Matrix, and intermediate matrix; for Lift-Splat, intermediate variables include the intermediate tensor and the predefined BEV grids. As illustrated in Fig. 5 (right), MatrixVT consumes 2 to 15 times less memory than LS-BEVDet and 40 to 80 times less memory than LS-BEVDet.

In Tab. 4, we adopt our proposed methods onto the BEVDet step-by-step and evaluate improvements in memory footprint and latency. It can be seen that Prime Extraction effectively reduced memory footprint and latency by a large margin. For the low-resolution setting (Res-50 & 256), the Ring & Ray FTM further reduces the resource consumption. For the high-resolution setting (V2-99 & 512×1408), the application of Ring & Ray FTM enabled VT without custom operators at cost of extra computation.

4.4. Effectiveness

In this section, we validate the effectiveness of the proposed methods by performance comparison and visualization.

4.4.1 Effects on Performance

As mentioned in Sec. 3.2, we argue that the features and corresponding depths can be compressed with little or no information loss. We thus individually adopt the Prime Extraction and the Ring & Ray FTM onto the BEVDet [9] to verify their effectiveness. To adopt Prime Extraction individually, we compress image features and the depths before the common Lift-Splat. As shown in Tab. 4, the BEVDet with our improvements achieves 41.5% NDS with Res-50 and 56.2% NDS with VovNetv2-99 [8], which is comparable to the baseline. Thus, we argue that Prime Extraction effectively retained key information from raw features, and

Backbone	PE	RR	NDS	Mem.	Lat. (ms)
Res-50			0.416	528M	1.79
Res-50	✓		0.416	94M	1.42
Res-50	✓	✓	0.415	83M	1.33
V2-99			0.564	2.14G	8.85
V2-99	✓		0.563	405M	2.04
V2-99	✓	✓	0.562	498M	3.19

Table 4. Effects of using Prime Extraction (PE) and Ring & Ray FTM (RR) on BEVDet. The memory consumption and latency reported is the memory footprint and time duration on PyTorch during View Transformation (VT).

the application of Ring & Ray FTM would not harm the final performance.

4.4.2 Prime Information in Object Detection

We then delve into the mechanism of Prime Extraction by visualization. Fig. 6 shows the inputs and outputs of the Prime Extraction module. It can be seen that the Prime Extraction module trained on different tasks focusing on different information. The Prime Depth Attention in object detection focuses on foreground objects. Thus the Prime Depth retained the depth of objects while ignoring the depth of the background. Also, it can be seen from the yellow car in the second column, which is obscured by three pedestrians. Prime Extraction effectively distinguishes these objects at different depths.

4.4.3 Prime Information in Map Segmentation

For the map segmentation task, we take lane segmentation as an example — the Prime Extraction module focus on lane and road that is closely related to this task. However, since

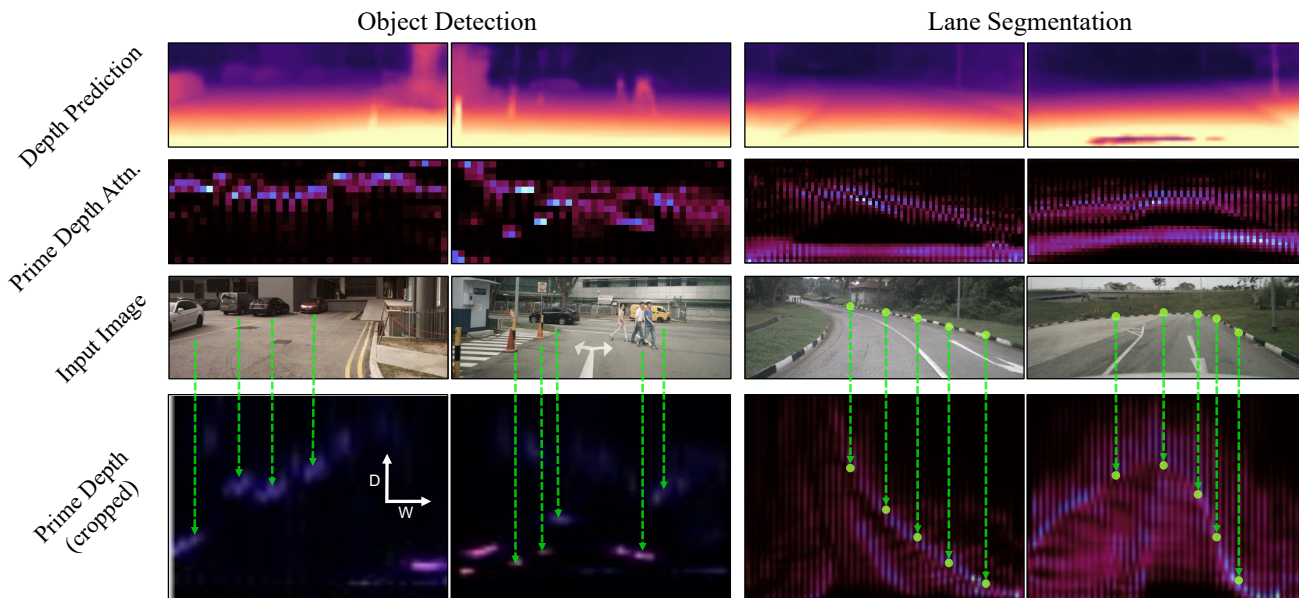


Figure 6. Depth Prediction, Prime Depth Attention, and Prime Depth for object detection and lane segmentation task. The categorical depth prediction is illustrated using distances with largest probability, while the Prime Depth is illustrated by raw probability values.

the area of the target category is a wide range covering several depth bins, the distribution of Prime Depth is uniform in the target area (see Fig. 6, 3rd and 4th columns). The observation indicates that Prime Extraction generates a new form of depth distribution that fits the map segmentation task. With the Prime Depth that is rather uniform, the same feature can be projected to multiple depth bins since they are occupied by the same category.

4.5. Extraction of Prime Feature

In the Prime Feature Extractor, we propose using max pooling followed by several 1D convolutions to reduce and refine the image feature. Before reduction, the coordinate of each pixel is embedded into the feature as position embedding. We conduct experiments to validate the contribution of each design.

A possible alternative of the max pooling is the CollapseConv as in [18] and [20], which merges the height dimension into the channel dimension and reduces the merged channel by linear projection. However, the design of CollapseConv brings some disadvantages. For example, the merged dimension is of size $C \times H$, which is high and requires extra memory transportation. To address these problems, we propose using max pooling to reduce the image feature in Prime Extraction. Tab. 5 shows that reduction using max pooling achieves even better performance than CollapseConv while eliminating these shortcomings.

We also conduct experiments to show the effectiveness of the refine sub-net after reducing in Tab. 5. The results indicate that the refine sub-net plays a vital role in adapting

Reduction	Refine	Pos. Emb.	mAP	NDS
Max Pooling			0.328	0.393
Max Pooling	✓		0.332	0.414
Max Pooling	✓	✓	0.336	0.415
CollapseConv	✓	✓	0.334	0.404

Table 5. Performance of MatrixVT with different feature reduction methods. The “Refine” indicates the refine convs, the “Pos.Emb.” indicates using front-view position embedding.

the reduced feature to BEV space, without which a performance drop of 0.9% mAP will occur. Finally, the experimental results in Tab. 5 have shown that the position embedding brings an improvement of 0.4% mAP, which is also preferable.

5. Conclusion

This paper proposes a new paradigm of View Transformation (VT) from multi-camera to Bird’s-Eye-View. The proposed method, MatrixVT, generalizes VT into a feature transportation matrix. We then propose Prime Extraction, which eliminates the redundancy during transformation, and the Ring & Ray Decomposition, which simplifies and boost the transformation. While being faster and more efficient on both specialized devices like GPU and general-purpose devices like CPU, our extensive experiments on the nuScenes benchmark indicate that the MatrixVT offers comparable performance to the state-of-the-art method.

References

- [1] Xuyang Bai, Zeyu Hu, Xinge Zhu, Qingqiu Huang, Yilun Chen, Hongbo Fu, and Chiew-Lan Tai. Transfusion: Robust lidar-camera fusion for 3d object detection with transformers. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 1090–1099, 2022. 1, 4
- [2] Holger Caesar, Varun Bankiti, Alex H. Lang, Sourabh Vora, Venice Erin Liong, Qiang Xu, Anush Krishnan, Yu Pan, Giancarlo Baldan, and Oscar Beijbom. nuscenes: A multi-modal dataset for autonomous driving. In *2020 IEEE/CVF Conference on Computer Vision and Pattern Recognition, CVPR 2020, Seattle, WA, USA, June 13-19, 2020*, pages 11618–11628. IEEE, 2020. 2, 5
- [3] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *2016 IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2016, Las Vegas, NV, USA, June 27-30, 2016*, pages 770–778. IEEE Computer Society, 2016. 5
- [4] Anthony Hu, Zak Murez, Nikhil Mohan, Sofia Dudas, Jeffrey Hawke, Vijay Badrinarayanan, Roberto Cipolla, and Alex Kendall. Fiery: Future instance prediction in bird’s-eye view from surround monocular cameras. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 15273–15282, 2021. 6
- [5] Junjie Huang and Guan Huang. Bevdet4d: Exploit temporal cues in multi-camera 3d object detection. *ArXiv preprint*, abs/2203.17054, 2022. 1, 6
- [6] Junjie Huang, Guan Huang, Zheng Zhu, and Dalong Du. Bevdet: High-performance multi-camera 3d object detection in bird-eye-view. *ArXiv preprint*, abs/2112.11790, 2021. 1, 2, 3, 6
- [7] Alex H. Lang, Sourabh Vora, Holger Caesar, Lubing Zhou, Jiong Yang, and Oscar Beijbom. Pointpillars: Fast encoders for object detection from point clouds. In *IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2019, Long Beach, CA, USA, June 16-20, 2019*, pages 12697–12705. Computer Vision Foundation / IEEE, 2019. 2, 3
- [8] Youngwan Lee and Jongyoul Park. Centermask: Real-time anchor-free instance segmentation. In *2020 IEEE/CVF Conference on Computer Vision and Pattern Recognition, CVPR 2020, Seattle, WA, USA, June 13-19, 2020*, pages 13903–13912. IEEE, 2020. 5, 6, 7
- [9] Yin hao Li, Zheng Ge, Guanyi Yu, Jinrong Yang, Zengran Wang, Yukang Shi, Jianjian Sun, and Zeming Li. Bevdepth: Acquisition of reliable depth for multi-view 3d object detection. *ArXiv preprint*, abs/2206.10092, 2022. 1, 2, 3, 5, 6, 7
- [10] Zhiqi Li, Wenhai Wang, Hongyang Li, Enze Xie, Chonghao Sima, Tong Lu, Qiao Yu, and Jifeng Dai. BEVFormer: Learning Bird’s-Eye-View Representation from Multi-Camera Images via Spatiotemporal Transformers. *ArXiv preprint*, abs/2203.17270, 2022. 1, 2, 6
- [11] Yingfei Liu, Tiancai Wang, Xiangyu Zhang, and Jian Sun. Petr: Position embedding transformation for multi-view 3d object detection. *ArXiv preprint*, abs/2203.05625, 2022. 2, 6
- [12] Zhijian Liu, Haotian Tang, Alexander Amini, Xinyu Yang, Huizi Mao, Daniela Rus, and Song Han. Bevfusion: Multi-task multi-sensor fusion with unified bird’s-eye view representation. *arXiv preprint arXiv:2205.13542*, 2022. 1
- [13] Hanspeter A Mallot, Heinrich H Bülthoff, JJ Little, and Stefan Bohrer. Inverse perspective mapping simplifies optical flow computation and obstacle detection. *Biological cybernetics*, 64(3):177–185, 1991. 1
- [14] NVIDIA, Péter Vingelmann, and Frank H.P. Fitzek. Cuda, release: 10.2.89, 2020. 6
- [15] Bowen Pan, Jiankai Sun, Ho Yin Tiga Leung, Alex Andonian, and Bolei Zhou. Cross-view semantic segmentation for sensing surroundings. *IEEE Robotics and Automation Letters*, 5(3):4867–4873, 2020. 1
- [16] Dennis Park, Rares Ambrus, Vitor Guizilini, Jie Li, and Adrien Gaidon. Is pseudo-lidar needed for monocular 3d object detection? In *IEEE/CVF International Conference on Computer Vision (ICCV)*, 2021. 5
- [17] Jonah Philion and Sanja Fidler. Lift, splat, shoot: Encoding images from arbitrary camera rigs by implicitly unprojecting to 3d. In *European Conference on Computer Vision*, pages 194–210. Springer, 2020. 1, 2, 3, 5, 6
- [18] Cody Reading, Ali Harakeh, Julia Chae, and Steven L. Waslander. Categorical depth distribution network for monocular 3d object detection. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 8555–8564, 2021. 1, 2, 8
- [19] Lennart Reiher, Bastian Lampe, and Lutz Eckstein. A sim2real deep learning approach for the transformation of images from multiple vehicle-mounted cameras to a semantically segmented image in bird’s eye view. In *2020 IEEE 23rd International Conference on Intelligent Transportation Systems (ITSC)*, pages 1–7. IEEE, 2020. 1
- [20] Thomas Roddick and Roberto Cipolla. Predicting semantic map representations from images using pyramid occupancy networks. In *2020 IEEE/CVF Conference on Computer Vision and Pattern Recognition, CVPR 2020, Seattle, WA, USA, June 13-19, 2020*, pages 11135–11144. IEEE, 2020. 1, 2, 4, 8
- [21] Thomas Roddick, Alex Kendall, and Roberto Cipolla. Orthographic feature transform for monocular 3d object detection. In *30th British Machine Vision Conference 2019, BMVC 2019, Cardiff, UK, September 9-12, 2019*, page 285. BMVA Press, 2019. 1, 2
- [22] Olaf Ronneberger, Philipp Fischer, and Thomas Brox. U-net: Convolutional networks for biomedical image segmentation. In *International Conference on Medical image computing and computer-assisted intervention*, pages 234–241. Springer, 2015. 2, 6
- [23] Daeyun Shin, Zhile Ren, Erik B Sudderth, and Charles C Fowlkes. 3d scene reconstruction with multi-layer depth and epipolar transformers. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 2172–2182, 2019. 2
- [24] Siddharth Srivastava, Frederic Jurie, and Gaurav Sharma. Learning 2d to 3d lifting for object detection in 3d for autonomous vehicles. In *2019 IEEE/RSJ International Confer-*

- ence on *Intelligent Robots and Systems (IROS)*, pages 4504–4511. IEEE, 2019. [1](#), [2](#)
- [25] Pei Sun, Henrik Kretzschmar, Xerxes Dotiwalla, Aurelien Chouard, Vijaysai Patnaik, Paul Tsui, James Guo, Yin Zhou, Yuning Chai, Benjamin Caine, Vijay Vasudevan, Wei Han, Jiquan Ngiam, Hang Zhao, Aleksei Timofeev, Scott Ettinger, Maxim Krivokon, Amy Gao, Aditya Joshi, Yu Zhang, Jonathon Shlens, Zhifeng Chen, and Dragomir Anguelov. Scalability in perception for autonomous driving: Waymo open dataset. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2020. [6](#)
- [26] Zhi Tian, Chunhua Shen, Hao Chen, and Tong He. FCOS: fully convolutional one-stage object detection. In *2019 IEEE/CVF International Conference on Computer Vision, ICCV 2019, Seoul, Korea (South), October 27 - November 2, 2019*, pages 9626–9635. IEEE, 2019. [2](#)
- [27] Tai Wang, ZHU Xinge, Jiangmiao Pang, and Dahua Lin. Probabilistic and geometric depth: Detecting objects in perspective. In *Conference on Robot Learning*, pages 1475–1485. PMLR, 2022. [2](#)
- [28] Tai Wang, Xinge Zhu, Jiangmiao Pang, and Dahua Lin. Fcos3d: Fully convolutional one-stage monocular 3d object detection. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 913–922, 2021. [2](#), [6](#)
- [29] Yan Wang, Wei-Lun Chao, Divyansh Garg, Bharath Hariharan, Mark E. Campbell, and Kilian Q. Weinberger. Pseudolidar from visual depth estimation: Bridging the gap in 3d object detection for autonomous driving. In *IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2019, Long Beach, CA, USA, June 16-20, 2019*, pages 8445–8453. Computer Vision Foundation / IEEE, 2019. [3](#)
- [30] Enze Xie, Zhiding Yu, Daquan Zhou, Jonah Philion, Anima Anandkumar, Sanja Fidler, Ping Luo, and Jose M Alvarez. M²bev: Multi-camera joint 3d detection and segmentation with unified birds-eye view representation. *ArXiv preprint*, abs/2204.05088, 2022. [1](#), [6](#)
- [31] Yan Yan, Yuxing Mao, and Bo Li. Second: Sparsely embedded convolutional detection. *Sensors*, 18(10):3337, 2018. [2](#), [5](#)
- [32] Tianwei Yin, Xingyi Zhou, and Philipp Krahenbuhl. Center-based 3d object detection and tracking. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 11784–11793, 2021. [2](#)
- [33] Yin Zhou and Oncel Tuzel. Voxelnet: End-to-end learning for point cloud based 3d object detection. In *2018 IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2018, Salt Lake City, UT, USA, June 18-22, 2018*, pages 4490–4499. IEEE Computer Society, 2018. [2](#), [3](#)
- [34] Benjin Zhu, Zhengkai Jiang, Xiangxin Zhou, Zeming Li, and Gang Yu. Class-balanced grouping and sampling for point cloud 3d object detection. *ArXiv preprint*, abs/1908.09492, 2019. [5](#)