

WidthFormer: Toward Efficient Transformer-based BEV View Transformation

Chenhongyi Yang¹ Tianwei Lin² Lichao Huang² Elliot J. Crowley¹
¹School of Engineering, University of Edinburgh ²Horizon Robotics

Abstract—We present *WidthFormer*, a novel transformer-based module to compute Bird’s-Eye-View (BEV) representations from multi-view cameras for real-time autonomous-driving applications. *WidthFormer* is computationally efficient, robust and does not require any special engineering effort to deploy. We first introduce a novel 3D positional encoding mechanism capable of accurately encapsulating 3D geometric information, which enables our model to compute high-quality BEV representations with only a single transformer decoder layer. This mechanism is also beneficial for existing sparse 3D object detectors. Inspired by the recently proposed works, we further improve our model’s efficiency by vertically compressing the image features when serving as attention keys and values, and then we develop two modules to compensate for potential information loss due to feature compression. Experimental evaluation on the widely-used nuScenes 3D object detection benchmark demonstrates that our method outperforms previous approaches across different 3D detection architectures. More importantly, our model is highly efficient. For example, when using 256×704 input images, it achieves 1.5 ms and 2.8 ms latency on NVIDIA 3090 GPU and Horizon Journey-5 computation solutions. Furthermore, *WidthFormer* also exhibits strong robustness to different degrees of camera perturbations. Our study offers valuable insights into the deployment of BEV transformation methods in real-world, complex road environments. Code is available at <https://github.com/ChenhongyiYang/WidthFormer>.

I. INTRODUCTION

In recent years, the field of vision-based Bird’s-Eye-View (BEV) 3D object detection has garnered significant interest and witnessed substantial advancements [21], [11], [16], [14], [5]. In contrast to directly detecting objects from image features, identifying 3D objects from a unified BEV representation aligns more intuitively with human perception and can be readily adapted to tasks such as 3D object tracking [34] for autonomous driving. Furthermore, BEV representations can be easily integrated with other modalities like LIDAR point clouds [14], [6].

The transformation of multi-view image features into a unified BEV representation, also referred as View Transformation (VT), plays a centre role in BEV-based 3D object detection. These can be broadly categorised into two streams: 1) Lift-Splat based approaches [21], [11], [10], [14], which first lift the image features to 3D space, and then gather the projected image features in the vertical direction to yield features for each BEV bin; and 2) Transformer-based methods [16], [12], which derive BEV representations by querying the image features via attention operations. While both types of methods have achieved great success [1], [24], their deployment to real-time autonomous-driving applica-

tions running on edge-computing devices [9], [22] presents considerable challenges. Here, we outline the three main obstacles that impede the successful deployment of these methods. (1) *Need for non-standard Operations*: Both Lift-Splat-based and Transformer-based methods necessitate special operations that require significant engineering effort for efficient implementation [11], [20], [22]. For example, in BEVFusion [20], a complex CUDA multi-threading mechanism was devised to gather point features in each spatial grid. The deformable attention operation [38] employed by BEVFormer [16] and PolarFormer [12] also necessitates specialised engineering effort when being deployed to edge computing devices. (2) *Heavy Computation*: Previous Transformer-based methods [16], [12] utilise multiple layers of transformer decoders to compute the BEV representation, which significantly impacts processing speed and impedes their deployability for real-time applications. Moreover, the stacking of deformable attention operations results in a large amount of random memory reads, which is often a bottleneck for edge-computing devices. (3) *Lack of Robustness*: The camera poses on a vehicle are usually perturbed due to a variety of factors, such as collisions or wear and tear. As we will demonstrate in Sec. IV, such perturbations substantially degrade the quality of BEV representations produced by prior Lift-Splat-based and Transformer-based methods. This makes these methods difficult to apply in real-world, complex driving environments. Given these considerations, we aim to develop a new BEV transformation method that is efficient, robust, and does not necessitate specialized engineering effort for deployment.

In this paper, we introduce *WidthFormer*: a new transformer-based BEV view transformation method. Our efficient VT module comprises a single layer of the transformer decoder [26] and does not involve any non-standard operations. Empowered by *Reference Positional Encoding (RefPE)*, a new method to compute positional encoding for transformer-based 3D object detection. *WidthFormer* computes BEV representation by directly querying the image features using the BEV queries computed from the coordinate on the BEV plane. To further mitigate the significant computational cost caused by the large amount BEV query vectors, we follow recently proposed approaches [36] to compress the image features in the vertical direction, significantly enhancing *WidthFormer*’s efficiency and scalability. In Fig. 1, we show a high-level comparison between *WidthFormer* and other BEV transformation methods. Moreover, to compensate for any potential information loss in

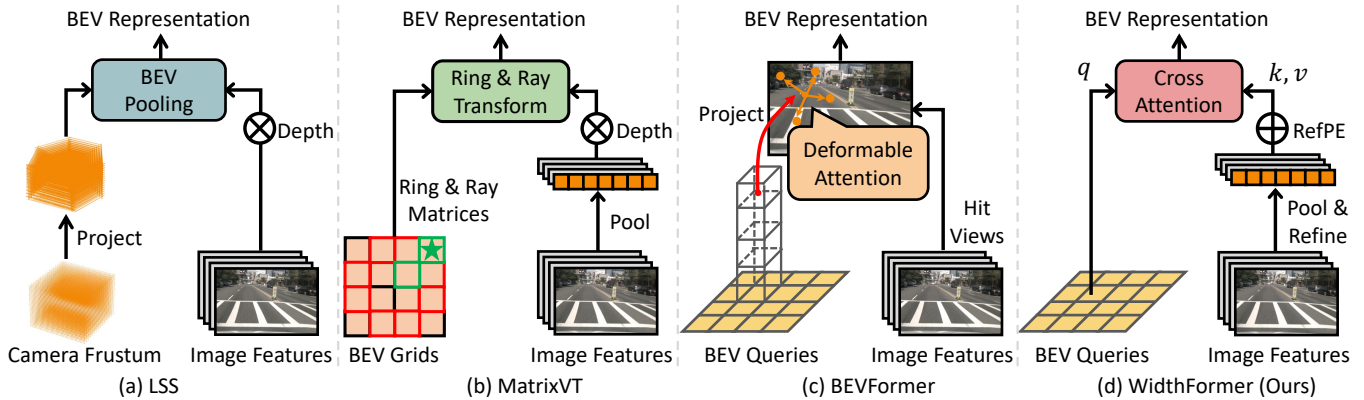


Fig. 1: Comparison on different BEV view transformation paradigms.

these compressed features, we designed two techniques: (1) We design an efficient Refine Transformer, where the compressed features attend to and extract valuable information from the original image features, and (2) We propose to train the model with complementary tasks that directly inject task-related knowledge into the compressed features.

To summarise, we make the following four contributions: (1) We introduce WidthFormer, a light-weighted and deployment-friendly BEV transformation method that employs a single layer of a transformer decoder to compute BEV representations. (2) We propose Reference Positional Encoding (RefPE), a new positional encoding mechanism for 3D object detection, to help WidthFormer assist WidthFormer’s view transformation. It can also be used to boost the performance of sparse 3D object detectors in a plug-and-play way. (3) We evaluate the proposed modules on the widely-used nuScenes 3D object detection dataset [1]. The results show that RefPE can significantly improve the performance of sparse object detectors. Also, WidthFormer outperforms previous BEV transformation methods in both performance and efficiency across a range of 3D detection architectures.

II. RELATED WORK

Vision-based 3D Object Detection. Image-based 3D object detection is the fundamental module for many downstream applications. Monocular 3D detection [4], [8], [28] aims to detect 3D objects from a single input image. For example, FCOS3D [29] extends the 2D FCOS [25] detector for 3D detection by regressing 3D bounding boxes. Multi-view 3D object detection [5], [17], [30] incorporates multiple images for better geometric inference. For example, PETR [18] extend the sparse detector DETR [2] by introducing 3D positional encoding. PETRv2 improves PETR [19] by incorporating temporal modeling. StreamPETR [27] proposed a novel query propagation algorithm to better leverage temporal information in a long range. CAPE [32] improves the positional encoding of PETR by creating positional encodings under local camera coordinate frames. The recently proposed 3DDPE [23] incorporates multi-modal supervision to accurately depth, based on which the point-wise positional encodings are computed. BEV-based 3D object detectors [14], [11], [10], [15], [36] first transform multi-view images to a unified BEV representation from which 3D objects are

detected. BEVDet [11], [14] and its follow-up work use LSS [21] to compute BEV features and predict objects using convolutional heads [34]. BEVFormer [16] computes BEV features using deformable attention operations [38] and relies on a DETR-style head [2] for object detection.

Vision-based BEV Transformations. The intuitive IPM-based methods [13], [15] compute BEV features through 3D-to-2D projection and interpolation. A problem with these is that the quality of BEV features will be severely harmed when the flat-ground assumption does not hold. In Lift-Splat based methods [21], [11], [10], BEV features are computed by vertically pooling the projected point cloud features, weighting by their predicted depth. While being high-performing, the Lift-splat process [21] is far from efficient. BEVFusion [20] accelerates this pooling process with a multi-threading mechanism. M²BEV [31] saves memory usage by assuming a uniform depth distribution. MatrixVT [36] improves overall efficiency by compressing the visual features in the vertical dimension and then computes BEV features using an efficient polar coordinate transformation. BEVDepth [14] incorporates point clouds for improved depth estimation. Transformer-based VT methods directly output BEV representations through the attention mechanism. CVT [35] and PETR [18] rely on 3D positional encodings to provide the model with 3D geometric information. To improve efficiency, many recent approaches [16], [12] adopt deformable attention [38].

III. METHOD

In this section, we introduce our method in detail. We introduce the new 3D position encoding mechanism in Sec. III-A. Then we introduce our BEV transformation module in Sec. III-B. In Sec. III-C, we describe how we refine the compressed features against potential information loss.

A. Reference Positional Encoding (RefPE)

In this work, we design a new 3D positional encoding mechanism for both transformer-based 3D object detectors, e.g., PETR [18], and our BEV view transformation module. Intuitively, for each visual feature that serves as the keys and values in the transformer, we aim to learn its positional encoding by referring to a series of reference points; therefore, we call our method Reference Positional Encoding (RefPE).

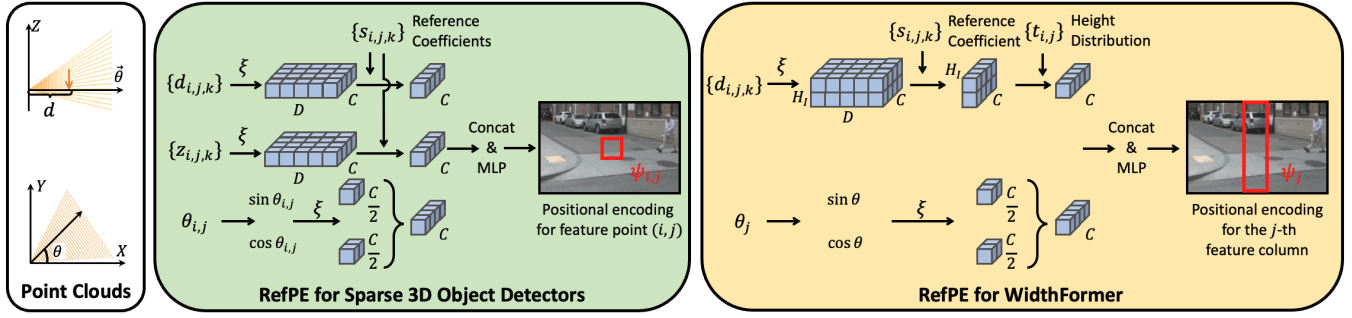


Fig. 2: Reference Positional Encoding (RefPE): RefPE has a rotation and a distance part. For sparse 3D detectors, it has another height part. The rotation encoding is computed using a camera ray’s rotation degree on the BEV plane. To compute the point-wise distance PE and height PE, we leverage the reference coefficients, predicted from the visual features, to aggregate the distance & height PEs of reference points on a camera ray. We remove the height PE for width features and compute their distance PE by aggregating all point-wise distance PE along an image column using a predicted height distribution.

Given the multi-view image features $\mathbf{F}^I \in \mathbb{R}^{N_c \times H_I \times W_I \times C}$, where N_c is the number of cameras; H_I , W_I and C are the height and width, and channel dimension, the computed positional encodings $\psi \in \mathbb{R}^{N_c \times H_I \times W_I \times C}$ has a same shape with the input features. Specifically, given a feature pixel¹ whose coordinate on the image plane is $\mathbf{P}_{i,j} = [u_{i,j}, v_{i,j}]^\top$, we first lift the pixel by associating it to D discrete depth bins, resulting in D reference points whose homogeneous coordinates are $\{\hat{\mathbf{P}}_{i,j,k} = [u_{i,j} \times d_k, v_{i,j} \times d_k, d_k]^\top \mid k \in |D|\}$. We project all D points to a unified 3D space by:

$$\mathbf{C}_{i,j,k} = \mathbf{R}^n \cdot \mathbf{I}^{-1} \cdot \hat{\mathbf{P}}_{i,j,k} + \mathbf{T}^n, \quad (1)$$

where $\mathbf{C}_{i,j,k} = [x_{i,j,k}, y_{i,j,k}, z_{i,j,k}]^\top$ is the point’s 3D Cartesian coordinate; $\mathbf{I} \in \mathbb{R}^{3 \times 3}$ is the camera intrinsic matrix; $\mathbf{R}^n \in \mathbb{R}^{3 \times 3}$ and $\mathbf{T}^n \in \mathbb{R}^{3 \times 1}$ are the rotation matrix and translation vector that transform the coordinate in the n -th view to the unified LIDAR coordinate frame [1], [7].

We then compute positional encodings using polar coordinates where the ego vehicle is located at the origin: each 3D point’s positional information is encoded as the combination of three parts: 1) distance to the ego on the BEV plane, 2) rotation related to the ego on the BEV plane, and 3) height to the ground. Specifically, for the 2D point $\mathbf{P}_{i,j}$ and its reference points $\{\mathbf{C}_{i,j,k} \mid k \in |D|\}$, we first compute the positional encodings for all its reference points $\{\tilde{\psi}_{i,j,k}\}$:

$$\tilde{\psi}_{i,j,k} = \text{Concat}(\xi(d_{i,j,k}), \xi(\sin \theta_{i,j,k}), \xi(\cos \theta_{i,j,k}), \xi(z_{i,j,k})) \quad (2)$$

$$d_{i,j,k} = \sqrt{x_{i,j,k}^2 + y_{i,j,k}^2}, \quad \sin \theta_{i,j,k} = \frac{y_{i,j,k}}{d_{i,j,k}}, \quad \cos \theta_{i,j,k} = \frac{x_{i,j,k}}{d_{i,j,k}} \quad (3)$$

where ξ is the Fourier positional encoding [26]. The positional encoding of $\mathbf{P}_{i,j}$ is computed by aggregating $\{\tilde{\psi}_{i,j,k}\}$ with a series of reference coefficients and an MLP:

$$\psi_{i,j} = \text{MLP}\left(\sum_{k=1}^D s_{i,j,k} \cdot \tilde{\psi}_{i,j,k}\right) \quad (4)$$

where the reference coefficients $\{s_{*,*,k}\}$ are predicted by a light-weighted convolution head.

RefPE for Sparse 3D Detectors As shown in Fig. 2, RefPE can be easily plugged into any PETR-style

detectors [18], [19]. The only modification that we need to make is to compute the positional encoding of the query vectors in a similar way that we compute RefPE. Specifically, given a 3D query anchor $[x_q, y_q, z_q]^\top$ in the LIDAR coordinate frame, we compute its positional encoding by:

$$\psi_q = \text{MLP}\left(\text{Concat}(\xi(d_q), \xi(\sin \theta_q), \xi(\cos \theta_q), \xi(z_q))\right) \quad (5)$$

$$d_q = \sqrt{x_q^2 + y_q^2}, \quad \sin \theta_q = \frac{y_q}{d_q}, \quad \cos \theta_q = \frac{x_q}{d_q} \quad (6)$$

B. BEV Transformation with WidthFormer

Empowered by the aforementioned RefPE, we introduce WidthFormer, an efficient transformer module for BEV view transformation. Our model’s overview is shown in Fig. 3.

Problem Description. The input of WidthFormer is the multi-view image features $\mathbf{F}^I \in \mathbb{R}^{N_c \times H_I \times W_I \times C}$ and the camera parameters. The output is a unified BEV representation $\mathbf{F}^B \in \mathbb{R}^{H_B \times W_B \times C}$ where H_B and W_B are the height and width of the BEV features, which are used to solve different downstream tasks, e.g., 3D object detection.

BEV Queries. Similar to how we compute positional encodings for the image features, WidthFormer computes the BEV query vectors using the center coordinates of each BEV grid. Specifically, for the grid with center coordinates $[x, y]^\top$, we compute its BEV query vector by:

$$q^B = \text{MLP}\left(\text{Concat}(\xi(d), \xi(\sin \theta), \xi(\cos \theta))\right) \quad (7)$$

$$d_{i,j} = \sqrt{x^2 + y^2}, \quad \sin \theta = \frac{y}{d}, \quad \cos \theta = \frac{x}{d} \quad (8)$$

Different from sparse 3D detectors, here we remove the height component when encoding the BEV positional information. As a consequence, the features of different heights will be automatically aggregated.

Width Features as Keys & Values. Intuitively, the BEV features can be easily computed by querying the multi-view image features, added with 3D positional encodings, with the BEV query vectors. However, as the sizes of BEV feature maps are usually in high resolution, e.g., 128×128 , such a naive approach will post a significant computational overhead, severely limiting the efficiency and scalability of the model. To overcome this difficulty, we follow recent works [36] to pool the image features \mathbf{F}^I in the vertical direction, yielding width features $\mathbf{F}^W \in \mathbb{R}^{N_c \times W_I \times C}$ to serve

¹Here we omit the view index for simplicity.

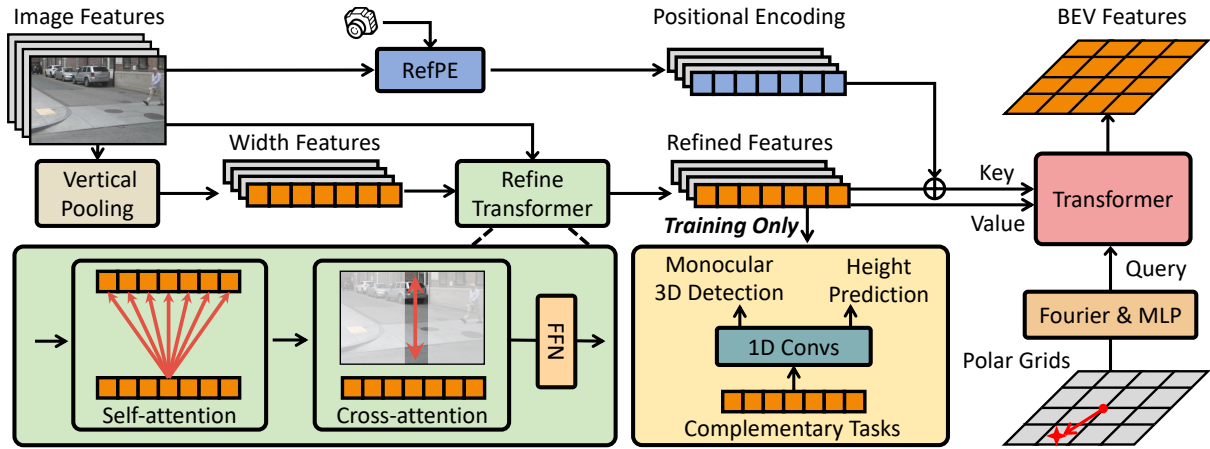


Fig. 3: WidthFormer takes multi-view images as input and outputs the converted BEV features. It first compresses the image features into width features by pooling their height dimension. Then the width features are refined with a Refine Transformer to compensate for potential information loss. After adding with our Reference Positional Encoding, the width features are fed into the transformer decoder to serve as keys & values, which are queried by the BEV query vectors computed from the pre-defined BEV polar coordinates.

as the attention keys and values. Compared with the full image features \mathbf{F}^I , the width features are H_I times smaller, greatly reducing the computational cost.

Adapting RefPE for Width Features. In Sec. III-A, we described how we compute RefPE for the full multi-view features. As shown in Fig. 2, to compute RefPE for the width features, we use a light-weighted *Conv head* to predict a discrete categorical height distribution [36] over each image column, based on which the positional encoding for a width feature is computed by aggregating the pixel-wise positional encodings on its corresponding column. Specifically, for the j -th width feature and the predicted height distribution $\{t_{i,j} | i \in [H_I]\}$, we compute its RefPE by:

$$\psi_j = \text{MLP}\left(\sum_{i=1}^{H_I} t_{i,j} \cdot \tilde{\psi}_{i,j}\right), \quad \text{where } \sum_{i=1}^{H_I} t_{i,j} = 1 \quad (9)$$

Similar to the BEV queries, we remove the height component when computing the pixel-wise positional encodings.

Computing BEV Representations. After acquiring the BEV queries \mathbf{Q}^B , width features \mathbf{F}^W and RefPE for the width features Ψ^W , WidthFormer uses a single transformer decoder layer to compute the BEV features \mathbf{F}^B :

$$\mathbf{U}^B = \mathbf{Q}^B + \text{MHA}(\mathbf{Q}^B, \mathbf{F}^W + \Psi^W, \mathbf{F}^W), \quad (10)$$

$$\mathbf{F}^B = \mathbf{U}^B + \text{FFN}(\mathbf{U}^B) \quad (11)$$

MHA and FFN are the multi-head attention operation and the feed-forward network in a transformer layer [26]. Note that we removed the self-attention module in the standard transformer decoder layer [26] as it is computationally expensive and has no impact on model performance.

C. Refining Width Features

Compressing 2D image features into 1D width features can greatly improve model efficiency and scalability. However, important information may also be lost during compression. Therefore, we introduce the following two techniques to compensate for any potential information loss.

Refine Transformer. Refine Transformer is a light-weight transformer decoder [26]. It refines an initial width feature by making it attend to and retrieve information from both other width features and the original image features. As illustrated in Fig. 3, the initial width feature is computed by MaxPooling the height dimension of the image features. In Refine Transformer, a width feature first retrieves information from other width features with a self-attention operation; it then retrieves information from its corresponding image column with a cross-attention operation. Finally, a feed-forward network is used to compute the final width feature. Refine Transformer is highly efficient because it has a linear complexity with respect to the input image size. Specifically, for an image feature map with size (H, W) , the computational complexity is $O(W^2) + O(WH)$. In Sec. IV-C, we show that adding Refine Transformer only incurs minimal latency to the model.

Complementary Tasks. To further improve the representation ability of width features, during training, we train the model with complimentary tasks to directly inject task-related information into the width features, which is motivated by BEVFormer v2 [33]. As shown in Fig. 3(a), the complementary tasks include a monocular 3D detection task and a height prediction task. Specifically, we append an FCOS3D-style head [29] for the two tasks. The head takes 1D width features as input and detects 3D objects in a monocular manner. To make it able to take 1D width features as input, we made two modifications: (1) We changed all 2D convolution operations to 1D convolutions; (2) we ignored the height range and only limited the width ranges during label encoding. To align the complementary tasks with WidthFormer, we changed the original regressional depth estimation to a categorical style. For height prediction, we appended an extra branch to the FCOS3D head to predict an object's height location in the original image features, which can complement the information lost in height pooling. Note that the head for the complementary tasks can be completely removed during model inference, so it will not

Detector	PE	mAP	NDS	mATE	mASE	mAOE	mAVE	mAAE
PETR-DN	PETR	34.3	37.2	77.0	27.7	59.3	112.6	35.2
	FPE	35.1	37.5	76.5	27.7	60.4	126.3	35.7
	CAPE	34.7	40.6	-	-	-	-	-
	Ours	37.1	41.2	72.3	27.5	54.8	94.1	24.3
StreamPETR	PETR	37.2	47.7	69.1	27.9	62.5	28.9	20.4
	FPE	38.6	48.7	66.1	27.4	63.0	28.3	20.5
	3DPPE	26.4	41.3	78.3	28.2	61.3	31.5	19.8
	Ours	40.2	49.9	64.5	27.8	60.8	27.6	20.5

TABLE I: Comparison of different positional encoding methods using PETR-DN, and StreamPETR detectors. ResNet-50-DCN is used as the default backbone. The input size is set to 512×1408 for PETR-DN and 256×704 for StreamPETR. All models are trained for 24 epochs without CBGS.

Detector	VT Method	mAP	NDS	mATE	mASE	mAOE	mAVE	mAAE
BEVDet	IPM	25.3	34.5	78.5	27.6	62.5	85.9	26.6
	LSS	29.5	37.1	73.9	27.3	61.2	88.1	24.8
	MatrixVT	28.9	36.5	74.6	28.3	60.0	89.5	27.3
	FastBEV	28.9	37.1	73.3	28.1	62.6	82.6	27.1
	BEVFormer	29.1	34.1	76.1	28.3	71.8	97.2	30.0
	WidthFormer	30.7	37.3	72.8	27.6	63.7	89.6	26.8
BEVDet4D	IPM	27.1	41.0	77.8	28.6	57.9	39.7	21.5
	LSS	32.8	45.7	71.0	27.9	51.2	36.0	20.5
	MatrixVT	32.4	45.8	69.6	27.6	51.9	36.3	18.9
	FastBEV	30.8	42.4	73.7	28.1	53.7	51.6	22.4
	BEVFormer	31.1	41.1	74.9	28.2	63.7	52.6	24.0
	WidthFormer	34.0	46.3	70.4	27.9	52.9	35.4	19.7

TABLE II: Comparison of different BEV view transformation methods using BEVDet and BEVDet4D detectors. ResNet-50 is used as the default backbone network. The input size is set to 256×704 . All models are trained for epochs with CBGS.

harm inference efficiency. Also, training the auxiliary head only consumes $<10M$ extra GPU memory, so it has minimal effect on training efficiency.

IV. EXPERIMENTS

A. Experiment Settings

Dataset. We benchmark our method on the commonly used nuScenes dataset [1], which includes 700, 150 and 150 scenes for training, validation and testing. Each scene contains 6-view images that cover the whole surrounding environment. We follow the official evaluation protocol. Specifically, for the 3D object detection task, except for the commonly used mean average precision (mAP), the evaluation metrics also include the nuScenes true positive (TP) errors, which contain mean average translation error (mATE), mean Average Scale Error (mASE), mean Average Orientation Error (mAOE), mean average velocity error (mAVE) and mean average attribute error (mAAE). We also report the nuScenes detection score (NDS). We use the nuScenes training set for model training and evaluate the models on the nuScenes *val* set.

Implementation Details. We test our RefPE for sparse 3D object detectors following the open-sourced implementation in [27]. We test our proposed WidthFormer using two 3D detection architectures: BEVDet [11] and BEVDet4D [10], which covers single-frame and multi-frame settings. We adopt the implementation of all three detectors in the BEVDet [11] code base. Unless otherwise specified, we use BEVDet’s [11] default data pre-processing and augmentation settings. We also follow [11] to set the BEV feature size to 128×128 and BEV channel size to 64. For BEVDet4D [10] and BEVDepth4D [14] experiments, we follow the original BEVDet4D implementation where only one history frame is used. All models are trained for 24 epochs with CBGS [37]. An ImageNet pre-trained ResNet-50 is used as the default

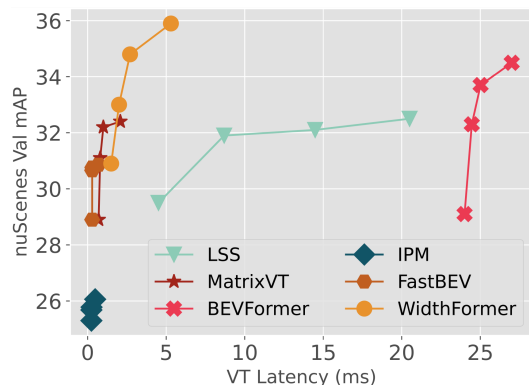


Fig. 4: CUDA Latency and mAP trade-off comparison of different VT methods using BEVDet on various size settings.

backbone network. All training and CUDA latency measurements are conducted using NVIDIA 3090 GPUs.

B. Main Results

RefPE vs Other Positional Encoding. In Tab. I, we compare our RefPE with other 3D positional encoding methods using a baseline detector PETR [18] and the state-of-the-art StreamPETR [27]. The competing PE approaches include the original 3D PE proposed in the PETR paper [19], the featurized 3D PE [19], CAPE [32] and the recently proposed 3DPPE [23]. For a fair comparison, all models are trained for 24 epochs without CBGS using ResNet-50-DCN as the backbone networks. The results show that our RefPE achieve the best performance on all three baselines. Specifically, on the PETR-DN detector, our method achieves significant 2.8 and 4.0 improvements on mAP and NDS over the baseline PETR, respectively. Also, on the high-performing StreamPETR, RefPE improves the mAP by 1.6 and NDS by 1.2 over the baseline StreamPETR that uses FPE as positional encoding. We also notice that, under our training recipe, the recently proposed 3DPPE does not achieve a satisfactory performance, showing its sensitivity to training setups. These results validate our RefPE’s effectiveness.

WidthFormer vs Other VT Methods. In Tab. II, we compare our proposed WidthFormer with other VT methods using the BEVDet [11] and BEVDet4D [10] 3D detectors. The competing methods include (1) Inverse Perspective Mapping (IPM) [13], (2) Lift-splat based LSS [21], MatrixVT [36] and FastBEV [15], and (3) transformer-based BEVFormer [16] (6-layers version). Note that the competing methods may use different backbone & head settings, multi-frame fusing strategies and training recipes in their original paper. Here, to ensure a fair comparison, we only use those VT methods to compute BEV representations and kept all other settings same following the original 3D detector [11], [10]. We also made sure the BEV representations computed from different VT methods have the same resolution and channel dimensions. Our method achieves better mAP and NDS performance than all competing methods on all three baseline architectures. Note that our method and BEVFormer [16] are both transformer-based, but our achieves better performance than BEVFormer in both single-frame and multi-frame settings. In addition, MatrixVT [36]

Setting	mAP	NDS	VT Latency
LSS Baseline	29.5	37.1	4.5 ms
+ Transformer	-	-	-
+ RefPE	31.0	38.0	4.6 ms
+ Width Feature	28.3	35.3	1.3 ms
+ Refine Transformer	29.3	36.2	1.5 ms
+ Auxiliary Head	30.7	37.3	1.5 ms

TABLE III: Step-by-step ablation studies.

Refinement	mAP	NDS	VT Latency
None	30.0	36.5	1.3 ms
Conv	30.2	36.3	1.4 ms
RF	30.7	37.3	1.5 ms
Conv+RF	30.9	37.2	1.6 ms

TABLE IV: Ablation study on feature refinement strategies.

and WidthFormer both compute BEV representations from width features, but ours is consistently better than MatrixVT on both settings. These results validate the effectiveness of our method as a general VT approach. In Fig. 4, we compare the CUDA latency and detection mAP curve of our method and other VT methods with different input resolutions and feature channels ($\{H_I, W_I, C\}$) on the BEVDet detector. Specifically, we tested four settings $\{256, 704, 64\}$, $\{384, 1056, 64\}$, $\{512, 1408, 64\}$ and $\{512, 1408, 128\}$. The result demonstrates the good speed and accuracy balance of our model: the detection mAP keeps improving as the inputs scale-up, while the VT latency keeps low. We do notice that IPM, MatrixVT and FastBEV achieve a better speed than ours, but this is at the expense of accuracy.

C. Ablation Studies and Discussions

We conduct ablation studies to test different design choices of our method. Unless otherwise specified, the experiments are conducted using a ResNet-50-based BEVDet.

Building WidthFormer step by step. We first examine the effect of our proposed modules by gradually adding them to the model and reporting accuracy & speed results in Tab. III. We start with the original LSS [21] baseline that achieves 29.3 mAP with 4.5 ms VT latency. We then replace LSS with a transformer layer, where BEV queries directly interact with the image features without 3D positional encoding (PE). However, the model fails to converge. We then add our RefPE encoding to the model, in which we compute PEs for every feature pixel without averaging them in the height dimension. The model achieves 31.0 mAP, which is better than the LSS baseline. This result validates the necessity of 3D PEs. However, querying the whole image features is inefficient (4.6 ms). By adopting width features as attention keys and values, the latency reduces to 1.3 ms. However, the accuracy falls to 28.3 mAP, implying that the feature compression is causing us to lose information. We mitigate this by adding the Refine Transformer to the model, which brings 1.0 mAP improvement with minor computation overhead (+0.2 ms). Finally, we add the complementary tasks that improve the final mAP to 30.7. These results show that the proposed Refine Transformer and complementary tasks can indeed mitigate the information loss problem.

Refine Transformer. We examine the design of our Refine Transformer and present the results in Tab. IV. The




Fig. 5: Latency comparison of different VT methods on the Horizon Journey-5 computation solutions.

Aux Loss	mAP	NDS
None	29.3	36.2
+ Width FCOS3D	29.7	36.4
+ Height Loss	30.5	36.9
+ CateDepth	30.7	37.3

TABLE V: Ablation Study on the auxiliary head and its losses.

baseline model that has no refinement to the width features achieves 30.0 mAP and 1.3 ms VT latency. When adding the convolutional-based refinement in MatrixVT [36], the accuracy is improved to 30.2 mAP. However, when adding our proposed Refine Transformer, the mAP is greatly improved to 30.7. We notice that our Refine Transformer is slightly slower than the convolution-based refinement but that this is negligible compared to the overall latency. We also tried to combine Refine Transformer with convolution-based refinement but found it only brought a marginal improvement in accuracy. This experiment validates the effectiveness of our Refine Transformer.

Complementary Tasks. We study our proposed complementary tasks and report the results in Tab. V. When training with complementary tasks is disabled, our model achieves 29.3 mAP. Then we add the 3D monocular detection task as in FCOS3D [29], the detection accuracy is improved to 29.7 mAP. Then we add the height estimation branch—where the model estimates an object’s height in the 2D image feature—the performance is further improved by 30.5 mAP. Finally, we modified the regression depth estimation loss into a categorical depth estimation loss, which gave us a final mAP to 30.7. This experiment validates the effectiveness of our training strategy.

Speed on Edge Computing Device. In Fig. 5, we report the speed-testing results comparing WidthFormer and other VT methods on the Horizon Journey-5 [22]. We compare our method with three other competing methods: GKT [3], LSS [21] and MatrixVT [36]. We report the VT latency with different size settings $\{H_I, W_I, C, H_B\}$ using square BEV grids. Specifically, the five size settings are: $S_1=(128, 352, 64, 128)$, $S_2=(256, 704, 64, 128)$, $S_3=(256, 704, 128, 128)$, $S_4=(256, 704, 128, 192)$ and $S_5=(304, 832, 128, 192)$. Note that the Journey-5 has special support for LSS, making it much faster than naive implementations. The results show that WidthFormer is nearly two times faster than LSS and several times faster




Fig. 6: Comparison of VT method’s robustness against different types of camera perturbations. We perturb each camera independently.

Method	Backbone	Resolution	Modality	MF	mAP \uparrow	NDS \uparrow	mATE \downarrow	mASE \downarrow	mAOE \downarrow	mAVE \downarrow	mAAE \downarrow
BEVDet [11]	Swin-B	512×1408	C		34.9	41.7	63.7	26.9	49.0	91.4	26.8
PETR [18]	Res-101	512×1408	C		35.7	42.1	71.0	27.0	49.0	88.5	22.4
BEVFormer-S [16]	Res-101	900×1600	C		37.5	44.8	-	-	-	-	-
BEVDet+WidthFormer	Res-101	512×1408	C		37.9	44.8	62.7	26.1	45.1	84.0	23.7
BEVDet4D [10]	Swin-B	640×1600	C	✓	39.6	51.5	61.9	26.0	36.1	39.9	18.9
BEVFormer [16]	Res-101	900×1600	C	✓	41.6	51.7	67.3	27.4	37.2	39.4	19.8
PolarFormer-T [12]	Res-101	900×1600	C	✓	38.3	48.8	70.7	26.9	34.4	51.8	19.6
Fast-BEV [15]	Res-101	900×1600	C	✓	40.2	53.1	58.2	27.8	30.4	32.8	20.9
PETRv2 [19]	Res-101	640×1600	C	✓	42.1	52.4	68.1	26.7	35.7	37.7	18.6
BEVDepth [14]	Res-101	640×1600	C+L	✓	41.2	53.5	56.5	26.6	35.8	33.1	19.0
BEVDistill [6]	Res-101	640×1600	C+L	✓	41.7	52.4	-	-	-	-	-
BEVDet4D+WidthFormer	Res-101	512×1408	C	✓	42.3	53.1	60.9	26.9	41.2	30.2	21.0
CAPE-T [32]	V2-99	320×800	C	✓	44.0	53.6	67.5	26.7	39.6	32.3	18.5
StreamPETR [27]	V2-99	320×800	C	✓	48.2	57.1	-	-	-	-	-
StreamPETR+3DPPE [23]	V2-99	320×800	C+L	✓	49.9	58.4	-	-	-	-	-
StreamPETR+RefPE	V2-99	320×800	C	✓	49.9	58.5	58.1	25.8	35.8	25.7	19.0

TABLE VI: Scaling-up detection results and comparison with other state-of-the-art 3D object detectors on nuScenes *val* set. ‘MF’ stands for multi-frame fusion; ‘C’ stands for camera; ‘L’ stands for LIDAR.

than GKT. We also notice that MatrixVT has a similar latency to ours, which differs from the results in Fig. 4. After investigation, we found that MatrixVT’s onboard speed is severely limited by the two very large ray and ring matrices. On the other hand, as shown in other experiments, WidthFormer has a better accuracy (Tab. II) than MatrixVT. Therefore, we argue that ours is a better choice for the actual autonomous driving applications.

Robustness under camera perturbation. In this section, we study the robustness of our method and other VT approaches under 6DoF camera perturbation: unexpected rotation and translation w.r.t. X , Y and Z axis in the camera coordinate system. Specifically, the perturbations are implemented by modifying the camera extrinsics with zero-mean Gaussian noises, whose magnitude is controlled by the standard deviation of the added Gaussian. Specifically, for the n -th camera, the noised rotation matrix is computed by $\hat{\mathbf{R}}^n = \mathbf{R}^n * \mathbf{R}'$, and the noised translation vector is computed by $\hat{\mathbf{T}}^n = \mathbf{T}^n + \mathbf{T}'$, where $*$ denotes matrix multiplication; \mathbf{R}' and \mathbf{T}' are the noisy rotation matrix and translation vector. We then use $\hat{\mathbf{R}}^n$ or $\hat{\mathbf{T}}^n$ to compute Eq. 1. We show the different VT methods’ mAP curves w.r.t. different types of perturbation in

Fig. 6. Interestingly, we find that the different VT methods are sensitive to different types of perturbations. For example, while the Y -axis translation causes performance decrease to IPM and BEVFormer, our method and Lift-splat based LSS and MatrixVT are completely insensitive to this perturbation because the height dimension is omitted in these approaches. In general, we observe that the Lift-splat based approaches are more robust than IPM and BEVFormer. The reason is that IPM and BEVFormer rely on projecting 3D space coordinates to 2D images, and a small amount of perturbation may cause a huge difference in the projected location, therefore harming performance. Particularly, our method is very robust to the translation perturbations with minor performance degrades compared to all competing approaches. Our method is also robust against X and Z axis rotation. However, we notice that all VT methods, including ours, have similar strong sensitivities to the Y -axis rotation, with the mAPs rapidly decreasing to zero as the noise magnitudes rise. The reason is that the Y -axis rotation can severely damages the predefined multi-view arrangement [1].

Scaling-up Experiments. In Tab. VI, we scale up our model and compare it with other state-of-the-art 3D detection ap-

proaches on the nuScenes *val* set. Specifically, for BEV-based detectors, we use ResNet-101 as the image encoder; we also scale up the input resolution and BEV channel dimensions to 512×1408 and 128, respectively. The results show that our method achieves on-par performance with other state-of-the-art methods. For example, BEVFormer [16] achieves 37.5 mAP & 44.8 NDS and 41.6 mAP & 51.7 NDS in single-frame and multi-frame settings, while ours achieves 37.9 mAP & 44.8 NDS and 40.1 mAP & 52.7 NDS. Additionally, the superior performance of StreamPETR+RefPE further validates the effectiveness of our design.

V. CONCLUSION

We introduce WidthFormer, a new transformer-based BEV transformation method. With the help of a new 3D positional encoding RefPE, which can also be used to boost a sparse 3D detector's performance, WidthFormer use a single layer of transformer decoder to compute the BEV representations from the vertically compressed visual features. Compared with previous BEV transformation approaches, our model is more efficient and easy to deploy to edge-computing devices. We also show that WidthFormer holds a good robustness against camera perturbations.

REFERENCES

- [1] Holger Caesar, Varun Bankiti, Alex H Lang, Sourabh Vora, Venice Erin Liong, Qiang Xu, Anush Krishnan, Yu Pan, Giancarlo Baldan, and Oscar Beijbom. nuscenes: A multimodal dataset for autonomous driving. In *CVPR*, 2020.
- [2] Nicolas Carion, Francisco Massa, Gabriel Synnaeve, Nicolas Usunier, Alexander Kirillov, and Sergey Zagoruyko. End-to-End Object Detection with Transformers. In *ECCV*, 2020.
- [3] Shaoyu Chen, Tianheng Cheng, Xinggang Wang, Wenming Meng, Qian Zhang, and Wenyu Liu. Efficient and robust 2d-to-bev representation learning via geometry-guided kernel transformer. *arXiv preprint arXiv:2206.04584*, 2022.
- [4] Xiaozhi Chen, Kaustav Kundu, Ziyu Zhang, Huimin Ma, Sanja Fidler, and Raquel Urtasun. Monocular 3d object detection for autonomous driving. In *CVPR*, 2016.
- [5] Zehui Chen, Zhenyu Li, Shiquan Zhang, Liangji Fang, Qinhong Jiang, and Feng Zhao. Graph-detr3d: rethinking overlapping regions for multi-view 3d object detection. In *Proceedings of the 30th ACM International Conference on Multimedia*, 2022.
- [6] Zehui Chen, Zhenyu Li, Shiquan Zhang, Liangji Fang, Qinhong Jiang, and Feng Zhao. BEVDistill: Cross-modal BEV distillation for multi-view 3d object detection. In *ICLR*, 2023.
- [7] MMDetection3D Contributors. MMDetection3D: OpenMMLab next-generation platform for general 3D object detection. <https://github.com/open-mmlab/mmdetection3d>, 2020.
- [8] Mingyu Ding, Yuqi Huo, Hongwei Yi, Zhe Wang, Jianping Shi, Zhiwu Lu, and Ping Luo. Learning depth-guided convolutions for monocular 3d object detection. In *CVPR workshops*, 2020.
- [9] Jon Gutiérrez-Zaballa, Koldo Basterretxea, Javier Echanobe, M. Victoria Martínez, Unai Martínez-Corral, Óscar Mata-Carballeira, and Inés del Campo. On-chip hyperspectral image segmentation with fully convolutional networks for scene understanding in autonomous driving. *Journal of Systems Architecture*, 2023.
- [10] Junjie Huang and Guan Huang. Bevdet4d: Exploit temporal cues in multi-camera 3d object detection. *arXiv preprint arXiv:2203.17054*, 2022.
- [11] Junjie Huang, Guan Huang, Zheng Zhu, and Dalong Du. Bevdet: High-performance multi-camera 3d object detection in bird-eye-view. *arXiv preprint arXiv:2112.11790*, 2021.
- [12] Yanqin Jiang, Li Zhang, Zhenwei Miao, Xiatian Zhu, Jin Gao, Weiming Hu, and Yu-Gang Jiang. Polarformer: Multi-camera 3d object detection with polar transformers. In *AAAI*, 2023.
- [13] Youngseok Kim and Dongsuk Kum. Deep learning based vehicle position and orientation estimation via inverse perspective mapping image. In *2019 IEEE Intelligent Vehicles Symposium (IV)*, 2019.
- [14] 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. *AAAI*, 2022.
- [15] Yangguang Li, Bin Huang, Zeren Chen, Yufeng Cui, Feng Liang, Mingzhu Shen, Fenggang Liu, Enze Xie, Lu Sheng, Wanli Ouyang, et al. Fast-bev: A fast and strong bird's-eye view perception baseline. *arXiv preprint arXiv:2301.12511*, 2023.
- [16] 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. In *ECCV*, 2022.
- [17] Xuewu Lin, Tianwei Lin, Zixiang Pei, Lichao Huang, and Zhizhong Su. Sparse4d: Multi-view 3d object detection with sparse spatial-temporal fusion. *arXiv preprint arXiv:2211.10581*, 2022.
- [18] Yingfei Liu, Tiancai Wang, Xiangyu Zhang, and Jian Sun. Petr: Position embedding transformation for multi-view 3d object detection. *ECCV*, 2022.
- [19] Yingfei Liu, Junjie Yan, Fan Jia, Shuailin Li, Qi Gao, Tiancai Wang, Xiangyu Zhang, and Jian Sun. Petr2: A unified framework for 3d perception from multi-camera images. In *ICCV*, 2023.
- [20] Zhijian Liu, Haotian Tang, Alexander Amini, Xingyu Yang, Huizi Mao, Daniela Rus, and Song Han. Bevfusion: Multi-task multi-sensor fusion with unified bird's-eye view representation. In *ICRA*, 2023.
- [21] Jonah Philion and Sanja Fidler. Lift, splat, shoot: Encoding images from arbitrary camera rigs by implicitly unprojecting to 3d. In *ECCV*, 2020.
- [22] Horizon Robotics. Horizon Robotics Journey™ 5. <https://en.horizonrobotics.com/products/horizon-journey-5/>, 2023.
- [23] Changyong Shu, Fisher Yu, and Yifan Liu. 3d point positional encoding for multi-camera 3d object detection transformers. In *ICCV*, 2023.
- [24] Pei Sun, Henrik Kretzschmar, Xerxes Dotiwalla, Aurelien Chouard, Vijaysai Patnaik, Paul Tsui, James Guo, Yin Zhou, Yuning Chai, Benjamin Caine, et al. Scalability in perception for autonomous driving: Waymo open dataset. In *CVPR*, 2020.
- [25] Zhi Tian, Chunhua Shen, Hao Chen, and Tong He. Fcos: Fully convolutional one-stage object detection. In *CVPR*, 2019.
- [26] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, ukasz Kaiser, and Illia Polosukhin. Attention is All you Need. In I Guyon, U V Luxburg, S Bengio, H Wallach, R Fergus, S Vishwanathan, and R Garnett, editors, *NeurIPS*, 2017.
- [27] Shihao Wang, Yingfei Liu, Tiancai Wang, Ying Li, and Xiangyu Zhang. Exploring object-centric temporal modeling for efficient multi-view 3d object detection. In *ICCV*, 2023.
- [28] Tai Wang, ZHU Xinge, Jiangmiao Pang, and Dahua Lin. Probabilistic and geometric depth: Detecting objects in perspective. In *CoRL*, 2022.
- [29] Tai Wang, Xinge Zhu, Jiangmiao Pang, and Dahua Lin. Fcos3d: Fully convolutional one-stage monocular 3d object detection. In *ICCV*, 2021.
- [30] Yue Wang, Vitor Campagnolo Guizilini, Tianyuan Zhang, Yilun Wang, Hang Zhao, and Justin Solomon. Detr3d: 3d object detection from multi-view images via 3d-to-2d queries. In *CoRL*, 2022.
- [31] 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 arXiv:2204.05088*, 2022.
- [32] Kaixin Xiong, Shi Gong, Xiaoqing Ye, Xiao Tan, Ji Wan, Errui Ding, Jingdong Wang, and Xiang Bai. Cape: Camera view position embedding for multi-view 3d object detection. *CVPR*, 2023.
- [33] Chenyu Yang, Yuntao Chen, Hao Tian, Chenxin Tao, Xizhou Zhu, Zhaoxiang Zhang, Gao Huang, Hongyang Li, Yu Qiao, Lewei Lu, et al. Bevformer v2: Adapting modern image backbones to bird's-eye-view recognition via perspective supervision. In *CVPR*, 2023.
- [34] Tianwei Yin, Xingyi Zhou, and Philipp Krahenbühl. Center-based 3d object detection and tracking. In *CVPR*, 2021.
- [35] Brady Zhou and Philipp Krähenbühl. Cross-view transformers for real-time map-view semantic segmentation. In *CVPR*, 2022.
- [36] Hongyu Zhou, Zheng Ge, Zeming Li, and Xiangyu Zhang. Matrixvt: Efficient multi-camera to bev transformation for 3d perception. In *ICCV*, 2023.
- [37] Benjin Zhu, Zhengkai Jiang, Xiangxin Zhou, Zeming Li, and Gang Yu. Class-balanced grouping and sampling for point cloud 3d object detection. *arXiv preprint arXiv:1908.09492*, 2019.
- [38] Xizhou Zhu, Weiye Su, Lewei Lu, Bin Li, Xiaogang Wang, and Jifeng Dai. Deformable detr: Deformable transformers for end-to-end object detection. *ICLR*, 2020.