# DualBEV: Unifying Dual View Transformation with Probabilistic Correspondences

Peidong Li, Wancheng Shen, Qihao Huang, Dixiao Cui

Zhijia Technology {lipeidong, shenwancheng, huangqihao, cuidixiao}@smartxtruck.com

Abstract. Camera-based Bird's-Eye-View (BEV) perception often struggles between adopting 3D-to-2D or 2D-to-3D view transformation (VT). The 3D-to-2D VT typically employs resource-intensive Transformer to establish robust correspondences between 3D and 2D features, while the 2D-to-3D VT utilizes the Lift-Splat-Shoot (LSS) pipeline for real-time application, potentially missing distant information. To address these limitations, we propose DualBEV, a unified framework that utilizes a shared feature transformation incorporating three probabilistic measurements for both strategies. By considering dual-view correspondences in one stage, DualBEV effectively bridges the gap between these strategies, harnessing their individual strengths. Our method achieves state-of-theart performance without Transformer, delivering comparable efficiency to the LSS approach, with 55.2% mAP and 63.4% NDS on the nuScenes test set. Code is available at https://github.com/PeidongLi/DualBEV.

Keywords: View Transformation · BEV · Autonomous Driving



Fig. 1: Unified Feature Transformation: Our approach considers correspondences between BEV and image space utilizing image probability, projection probability and BEV probability. In the 3D-to-2D strategy, HeightTrans (HT) projects pre-defined 3D points to sample features, while in the 2D-to-3D strategy, LSS lifts image features to 3D space, both through image probability and projection probability from different directions. Finally, BEV probability is applied to enhance the representation of features.



**Fig. 2:** Comparison of Fusion Strategy.  $\oplus$  means sum function.  $\otimes$  denotes multiplication.  $\otimes$  denotes channel-attention-based fusion.

# 1 Introduction

Effective BEV object detection in autonomous driving relies on precise feature transformation from the perspective to the BEV space, facilitated by the VT module. Current methods predominantly employ either a 2D-to-3D or 3D-to-2D strategy. In 2D-to-3D methods [7, 8, 10, 13, 14, 34, 38], dense 2D features are elevated by predicting depth probabilities, but the inherent uncertainty in depth prediction can introduce inaccuracies, particularly in distant regions. Conversely, 3D-to-2D methods [9, 15, 17–19, 30, 32, 36] often use BEV queries to sample 2D features, leveraging Transformer [29] to learn attention weights for each 3D-2D correspondence, while introducing computational and deployment complexities.

HeightFormer [32] and FB-BEV [16] have explored integrating both VT. Typically, these methods employ a two-stage strategy due to different feature transformation of dual VT. Utilizing LSS features to initialize the Transformerbased VT, this strategy is constrained by the performance of the initial feature, hindering seamless fusion between the dual VT. Additionally, these methods still confront challenges in achieving real-time deployment in autonomous driving.

In this paper, we argue that both VT inherently establish correspondences between 3D and 2D features from different views, while LSS and Transformer serving as distinct methods for evaluating these correspondences. To unify dual VT, we propose a unified feature transformation (illustrated in Fig. 1) applicable to both 2D-to-3D and 3D-to-2D VT, evaluating the correspondences by three probabilistic measurements: 1) **BEV Probability**, aimed at mitigating the impact of blank BEV grids in feature construction; 2) **Projection Probability**, which distinguishes multiple correspondences, accounting for different 3D points projecting into the same 2D position; 3) **Image Probability**, aiding in excluding background features during feature transformation.

Applying this unified feature transformation, we shed light on an often overlooked approach: leveraging Convolutional Neural Network (CNN) for 3D-to-2D VT, with the introduction of **HeightTrans**. Beyond its impressive performance, we demonstrate its potential for acceleration through pre-computation, rendering it suitable for real-time autonomous driving applications. Meanwhile, we enhance the traditional LSS pipeline by integrating this feature transformation, denoted as **Prob-LSS**, showcasing its universality for current detectors.

Combining HeightTrans and Prob-LSS, our research introduces DualBEV (see Tab. 1), an innovative approach that incorporates and considers correspondences from both BEV and perspective view in the one-stage manner, eliminating the reliance on initial features. Furthermore, we propose a robust BEV **Table 1:** Comparison of different VT formulations for obtaining the BEV feature F at location (x, y). p means point in the BEV grid, with (u, v) representing the corresponding image position. I denotes the image features, D represents the depth maps, and M signifies the instance mask of the images. Q denotes BEV queries, where B represents the binary BEV mask, and P represents the predicted BEV probability. The notation  $\mathcal{F}_{da}$  refers to the deformable attention function, while  $\mathcal{G}_{2d}$  denotes the bilinear grid sampler. f denotes fusion function, and w represents depth consistency [16].  $N_z$  means number of pre-defined points in BEV grid,  $N_c$  means number of corresponding points in image and  $N_d$  is the number of image frustum points pooling into the grid.

Method	Strategy VT Formulation of F(x,y)
BEVDet [7]	$  2D \rightarrow 3D   \sum_{k=1}^{N_d} D_k \cdot I_k$
SA-BEV [10]	$2D \rightarrow 3D \left  \sum_{k=1}^{N_d} D_k \cdot M_k \cdot I_k \right $
Prob-LSS(ours)	$2D \rightarrow 3D   P(x,y) \cdot \sum_{k=1}^{N_d} D_k \cdot M_k \cdot I_k$
Simple-BEV [3]	$ $ 3D $\rightarrow$ 2D $ $ Conv(Concat( $\mathcal{G}_{2d}((u, v), I)))$
BEVFormer [15]	$3D \rightarrow 2D \left  \sum_{i=1}^{N_z} \sum_{j=1}^{N_c} \mathcal{F}_{da}(Q(x,y), p_{ij}, I) \right $
DA-BEVFormer [16]	$3D \rightarrow 2D \left  \sum_{i=1}^{N_z} \sum_{j=1}^{N_c} \mathcal{F}_{da}(Q(x,y), p_{ij}, I) \cdot w(D_{ij}) \right $
$\operatorname{HeightTrans}(\operatorname{ours})$	$3D \rightarrow 2D  P(x,y) \stackrel{\circ}{\cdot} \sum_{i=1}^{N_z} \sum_{j=1}^{N_c} D_{ij} \cdot M_{ij} \cdot I_{ij}$
FB-BEV [16]	$\frac{1}{3D\&2D} \sum_{k=1}^{N_d} D_k \cdot I_k + B \cdot \sum_{i=1}^{N_z} \sum_{j=1}^{N_c} \mathcal{F}_{da}(Q, p_{ij}, I) \cdot w(D_{ij})$
DualBEV(ours)	$3D\&2D   P(x,y) \cdot f(D \cdot M \cdot I)$

feature fusion module termed Dual Feature Fusion (DFF) module. This module enhances the integration of dual BEV features by leveraging channel attention module, while spatial attention module further helps to refine BEV probability prediction. DualBEV operates on the principle of "broad input, stringent output", leveraging precise dual-view probabilistic correspondences to comprehend and represent the probability distribution of the scene.

Our major contributions are summarized below:

- 1) We unveil the inherent similarity between 3D-to-2D and 2D-to-3D VT and identify a unified feature transformation covering dual VT. This transformation enables accurate correspondence establishment from both BEV and perspective view, significantly bridging the gap between the dual strategies.
- 2) We propose a novel CNN-based 3D-to-2D VT termed HeightTrans. Leveraging probabilistic sampling and pre-computation of lookup table, HeightTrans establishes precise 3D-2D correspondences effectively and efficiently.
- 3) We introduce the DFF for dual-view features fusion. This fusion strategy captures information from both close and distant regions in one-stage, culminating in the generation of comprehensive BEV features substantially.
- 4) Our efficient framework DualBEV achieves state-of-the-art performance with a remarkable 55.2% mAP and 63.4% NDS on the nuScenes test set even without Transformer, highlighting the significance of capturing precise dualview correspondences for view transformation.

# 2 Related Work

## 2.1 3D-to-2D View Transformation

OFT-Net [28] pioneered the integration of sampling methods into 3D-to-2D VT for monocular detection, populating voxel features by aggregating image features from corresponding projection regions. Recent BEV methods extend geometric projection with Transformer utilizing a cross-attention mechanism, broadly categorized into two streams: explicit dense BEV queries or implicit sparse object queries. The former [15, 32, 37] constructs a pre-defined BEV space covering a limited 3D range, with 3D-2D correspondences heavily reliant on attention mechanisms, incurring high computational costs. The latter [17, 18, 30] employs learnable object queries to cover all possible object proposals, a concept challenging to apply in dense tasks like lane segmentation and 3D occupancy prediction.

Contrarily, Simple-BEV [3] projects 3D voxels into images and samples features bilinearly, akin to OFT-Net's approach. It bypasses deformable attention [33] weights presented in Transformer-based methods, instead employing convolution to reduce concatenated channels in the height dimension. However, this simplistic sampling method without weights still lags behind 2D-to-3D CNNbased methods in latency. In this paper, we propose HeightTrans to evaluate correspondences and sum features directly in the BEV grid through a lookup table, improving the speed of this strategy significantly.

#### 2.2 2D-to-3D View Transformation

A prevalent approach for 2D-to-3D VT involves lifting multi-view 2D camera features into 3D via pixel-wise discrete depth estimation, followed by BEV feature extraction through pillar sum-pooling in 3D space. This paradigmatic approach was first introduced by LSS [26] and has been followed by many subsequent works [7, 13, 14, 27]. BEVDepth [14] and BEVStereo [13] highlight the critical role of accurate depth estimation, with explicit depth supervision enhancing performance. However, the efficiency of the subsequent *Splat* stage remains a notable challenge, addressed by innovations such as BEV Pooling [5, 7, 14]. BEV-SAN [34] proposed a Slice Attention Module to focus on the different height slices where the different categories are located. In the concurrent work, SA-BEV [10] suggested using SA-BEVPool to ignore points that are part of the background during BEV Pooling. Our work further extend the idea to the BEV space to ignore invalid features due to uncertainty in depth estimation.

#### 2.3 Fusion View Transformation

Recent methods, such as HeightFormer [32] and FB-BEV [16], attempt to fuse both approaches. HeightFormer introduces a height predictor on the initial BEV feature, using deformable attention for feature sampling to refine the initial feature. HeightFormer demonstrates the fusion ability by generating secondstage feature based on the first-stage feature and combining them. FB-BEV



Fig. 3: Overview of DualBEV: Initially, we employ SceneNet to predict the depth D (Projection probability) and instance mask M (Image probability) of input images. Subsequently, the Prob-LSS stream follows the BEVPoolv2 [5] to generate LSS feature. Concurrently, the HeightTrans stream utilizes the Prob-Sampling to project pre-defined 3D points onto the 2D space, retrieving corresponding image features. Throughout this process, all features are accompanied by probabilities derived from the depth map and instance mask. Finally, we fuse two streams and predict the BEV probability P by leveraging the DFF module, resulting in the final BEV feature F.

extends this concept further by introducing FRPN [16] to the first-stage BEV feature, selecting valid positions for Depth-Aware BEVFormer (DA-BEVFormer) to generate the second-stage feature. FB-BEV notes the differences between two paradigms, enabling a two-stage VT that leverages both 2D-to-3D and 3D-to-2D strategies. Our framework, DualBEV, further unveils the inherent sameness of dual VT, capturing information from each stream efficiently. Additionally, our method offers a more lightweight and deployment-friendly alternative without Transformer. Illustrated in Fig. 2, our approach fuses features sharing same transformation from different views in a one-stage manner, eliminating the reliance on initial BEV features and benefiting from more input information.

## 3 Method

As illustrated in Fig. 3, the DualBEV pipeline commences with the extraction of image feature  $I \in \mathbb{R}^{N \times C_I \times H_I \times W_I}$  from N cameras by the image backbone, where  $H_I \times W_I$  is the shape of image feature. Subsequently, a SceneNet is employed to generate both the instance mask  $M \in \mathbb{R}^{N \times C_M \times H_I \times W_I}$  and depth map  $D \in \mathbb{R}^{N \times C_D \times H_I \times W_I}$ . The structure of SceneNet mirrors that of DepthNet [14] with just increasing output channels. Binary Cross Entropy (BCE) loss is applied for instance supervision same as depth supervision, following SA-BEV [10].

The HeightTrans module employs probabilistic sampling to acquire image features. Simultaneously, the Prob-LSS stream follows the approach outlined in BEVPoolv2 [5] by lifting image with instance segmentation via depth prediction. The features from these two streams are then fed into the DFF module for fusion and BEV probability prediction. Finally, the BEV probability  $P \in \mathbb{R}^{1 \times H_F \times W_F}$  is applied to the fused feature to obtain the final BEV feature  $F \in \mathbb{R}^{C_F \times H_F \times W_F}$  for downstream tasks, where  $H_F \times W_F$  is the shape of the BEV feature.

### 3.1 HeightTrans

The fundamental principle of 3D-to-2D VT revolves around selecting 3D positions for projection into the image space and assessing these 3D-2D correspondences. In methods focusing on explicit BEV feature generation, a pre-defined BEV map facilitates the derivation of a 2D position once the height of the BEV grid is determined. Building upon this concept, our approach begins by sampling sets of 3D points within a pre-defined BEV map. These sampled correspondences are then considered and filtered carefully to generate the BEV feature by summing within each BEV grid.

**BEV Height** While existing approaches often rely on sparse uniform sampling in the height range of the BEV grid to initialize 3D points, it's important to recognize that different heights encode distinct information in the 3D space. Inspired by BEV-SAN [34], HeightTrans introduces a multi-resolution sampling strategy covering the entire height range [-5m, 3m], with a resolution of 0.5m within the ROI (Region of Interest) of [-2m, 2m], and 1.0m outside this range. This sampling strategy enhances focus on small objects that might be easily missed with a coarser resolution. Unlike deformable attention methods where a set of offsets is typically predicted around the projected 2D position, our method, with its increased number of sampling points in 3D space, eliminates the need for predicting offsets in the image space, allowing for pre-computation.

**Prob-Sampling** With our pre-defined set of 3D sampling points  $p_{3d} \in \mathbb{R}^3$ , the subsequent task involves acquiring features for each position and weighing the various correspondences. Given a 3D point  $p_{3d} = (x, y, z)$  in the 3D space, the camera's extrinsic matrix T and intrinsic matrix K, the projection yields a corresponding 2D point  $p_{2d} = d \cdot (u, v, 1)$  in the image space, where d signifies the depth of the point.

$$p_{2d} = K \cdot T \cdot p_{3d} \tag{1}$$

A straightforward method to obtain the 3D features  $F_{ht}$  is by using a bilinear grid sampler  $\mathcal{G}_{2d}$  to sample the image feature I at the projected position  $p_{2d}$ :

$$F_{ht}(p_{3d}) = \mathcal{G}_{2d}(I, p_{2d})$$
 (2)

However, the projected position may land on a background pixel, which is not only useless but could also be misleading for detection. Instead, we utilize the instance mask M derived from SceneNet to represent the image probability  $P_{img}$ , which we apply to the image feature to mitigate this concern.

$$F_{ht}(p_{3d}) = \mathcal{G}_{2d}(M \cdot I, p_{2d}) \tag{3}$$

To distinguish multiple 3D points hitting the same 2D position, we further use a projection probability  $P_{proj}$  to evaluate these multiple correspondences.  $P_{proj}$ is obtained by a trilinear grid sampler  $\mathcal{G}_{3d}$  on the depth map D, treating the depth channel as the third dimension.

$$F_{ht}(p_{3d}) = \mathcal{G}_{3d}(D, p_{2d}) \cdot \mathcal{G}_{2d}(M \cdot I, p_{2d}) \tag{4}$$

Finally, to address the issue of blank BEV grids that offer no useful information for detection, we introduce a BEV probability  $P_{bev}$  to represent the occupied probability of the BEV grid, where (x, y) is the location in BEV space.

$$F_{ht}(p_{3d}) = P(x,y) \cdot \mathcal{G}_{3d}(D,p_{2d}) \cdot \mathcal{G}_{2d}(M \cdot I,p_{2d})$$

$$\tag{5}$$

Acceleration BEVPoolv2 [5] utilizes pre-computation for the calculation of 3D points index in BEV space from the defined frustum, where both image feature index and depth map index remain fixed during inference. Similarly, we can accelerate our VT by building a lookup table after replacing the grid sampler with the round function. The BEV feature in Eq. (5) can be simplified as:

$$F_{ht}(x, y, z) = P(x, y) \cdot D(u, v, d) \cdot M(u, v) \cdot I(u, v)$$
(6)

Now we can sum the features of  $N_z$  pre-defined points in each BEV grid for  $N_c$  corresponding 2D positions to obtain the final HeightTrans feature as:

$$F_{ht}(x,y) = P(x,y) \cdot \sum_{i=1}^{N_z} \sum_{j=1}^{N_c} D(u_{ij}, v_{ij}, d_{ij}) \cdot M(u_{ij}, v_{ij}) \cdot I(u_{ij}, v_{ij})$$
(7)

As shown in Tab. 1, this representation bears similarity to BEV Pooling. Therefore, we can establish the lookup table by projecting pre-defined 3D points into the image space and subsequently calculating the index in the feature map and depth map. Unlike BEV Pooling, the index in BEV map is constant in our method. Then we can use the same CUDA operator as in BEVPoolv2 [5] with the lookup table to speed up the calculation of BEV features during inference.

#### 3.2 Prob-LSS

The traditional LSS pipeline begins by predicting the depth probability of each pixel to facilitate its elevation into frustum points, which is then projected into the BEV space through BEV Pooling. SA-BEV [10] introduced Semantic-Aware BEV Pooling to circumvent the need for lifting irrelevant pixels in image space. However, the inherent uncertainty in depth estimation may still result in extra-neous information being present in the BEV space.

To address this issue, we further integrate BEV probability into the LSS pipeline, referred to as Prob-LSS. The construction of LSS features at location (x, y) is similar to Eq. (7) and can be expressed as follows:

$$F_{lss}(x,y) = P(x,y) \cdot \sum_{k=1}^{N_d} D(u_k, v_k, d_k) \cdot M(u_k, v_k) \cdot I(u_k, v_k)$$
(8)



Fig. 4: BEV Feature Visualization with GT boxes. Prob-LSS pays more attention to close range while HeightTrans can also capture distant information. In the red rectangle (distant range), where our unified framework compensates weak detection on barriers of Prob-LSS with HeightTrans. In the orange rectangle (close range), BEV features are enhanced from dual streams. Ego locates in the center of BEV features.

where  $N_d$  represents the number of projected frustum points in the BEV grid at location (x, y). It's noteworthy that for each BEV grid, HeightTrans provides a constant number with  $N_z \times N_c$  transformed image features, whereas LSS offers a dynamic number with  $N_d$  features depending on the lifted frustum. This dynamic nature of LSS acts as a complementary aspect to the HeightTrans.

#### 3.3 Dual Feature Fusion

Unifying Eqs. (7) and (8), we observe that BEV probability can be obtained after the fusion of dual features. Inspired by the CBAM [31] and AFF [2] framework, we propose a Dual Feature Fusion (DFF) module to integrate these features and predict the BEV probability effectively. The DFF module comprises a fusion module f, which utilizes channel attention to predict weights for dual feature summation, and a Spatial Attention Enhanced ProbNet (SAE-ProbNet) to predict the BEV probability P, formulated as:

$$F(x,y) = P(x,y) \cdot f(F_{lss}(x,y), F_{ht}(x,y)) \tag{9}$$

**Fusion Module** As shown in Fig. 4,  $F_{lss}$  focuses more on close objects occupying most of the image pixels, while  $F_{ht}$  tends to contain information of distant objects. This observation suggests two-stage methods [16,32] initialized by LSS features, may not guide second-stage VT to extract features at distant regions.

To tackle these challenges and obtain a robust BEV representation, our fusion module f concatenates the dual features and passes them into a channelattention-based fusion (CAF) module to predict the affinity for feature selection.



Fig. 5: Dual Feature Fusion Module: Dual features are first concatenated and then passed into the CAF module for fusion. Subsequently, the SAE-ProbNet is utilized to obtain the BEV probability for the final BEV feature.

The fusion feature  $F_{channel}$  is obtained using the following equation:

$$F_{channel} = f(F_{lss}, F_{ht}) \tag{10}$$

$$= C(F_{lss} \ominus F_{ht}) \cdot F_{lss} + (1 - C(F_{lss} \ominus F_{ht})) \cdot F_{ht}$$
(11)

where  $\ominus$  denotes concatenation, and *C* denotes the CAF module, which is modified from the MS-CAM [2] as illustrated in Fig. 5. This fusion phase aims to select features from two streams softly with a learning weight, enhancing the representation in both close and distant regions.

**BEV Probability Prediction** We utilize the BEV probability P predicted by the SAE-ProbNet to aggregate  $F_{channel}$ , thereby mitigating the impact of blank BEV grids. We employ ProbNet to extract local information, serving as the local stream  $P_l$ . Further enhancement of the prediction is achieved by incorporating a spatial attention module to capture global information, serving as the global stream  $P_g$ . The entire module can be formulated as:

$$F_{spatial} = P \cdot F_{channel} \tag{12}$$

$$= \sigma(P_l(F_{channel}) + P_g(F_{channel})) \cdot F_{channel}$$
(13)

As depicted in Fig. 5, the local stream, ProbNet, employs a  $3 \times 3$  convolution kernel to reduce the channel dimensions. Subsequently, it undergoes processing through a ResBlock-CBAM [31] and a  $1 \times 1$  convolution operation to acquire

#### 10 P. Li et al.

**Table 2:** Comparison on the nuScenes val set. ResNet-50 is used as image backbone and resolution is set to  $704 \times 256$ . †: We apply BEVDepth [14] with BEVPoolv2 [5] as baseline method.  $\star$ : For fair comparison with LSS method, we disable SceneNet and instance supervision in single-frame. ‡: Features' operation is performed at 1/16 resolution without BEV-Paste [10] for fair comparison.

Method	Frames	$ mAP\uparrow$	$NDS\uparrow$	mATE↓	mASE↓	mAOE↓	mAVE↓	$mAAE\downarrow$
BEVDet [7]	1	29.8	37.9	0.725	0.279	0.589	0.860	0.245
BEVDepth [14]	1	33.7	41.4	0.646	0.271	0.574	0.838	0.220
BEVDepth [5] †	1	34.2	40.7	0.645	0.273	0.599	0.890	0.240
DualBEV*	1	35.2	42.5	0.640	0.271	0.542	0.838	0.216
BEVDet4D [6]	2	32.2	45.7	0.511	0.241	0.386	0.301	0.121
BEVFormer [15]	2	33.0	45.9	0.686	0.272	0.482	0.417	0.201
BEVDepth [14]	2	35.1	47.5	0.639	0.267	0.479	0.428	0.198
BEVDepth <sup>†</sup> [5]	2	36.8	48.5	0.609	0.273	0.507	0.406	0.196
BEVStereo [13]	2	37.2	50.0	0.598	0.270	0.438	0.367	0.190
FB-BEV [16]	2	37.8	49.8	0.620	0.273	0.444	0.374	0.200
SA-BEV <sup>‡</sup> [10]	2	37.8	49.9	0.617	0.270	0.441	0.370	0.206
DualBEV	2	38.0	50.4	0.612	0.259	0.403	0.370	0.207

local attention. Supervised by a BEV mask, ProbNet utilizes BCE loss and Dice loss [24], with BEV Centerness [35] also incorporated into the loss to encourage the network to focus more on distant objects. On the other hand, the global stream utilizes a  $7 \times 7$  convolution kernel to compute the average and maximum values of the input feature, thus enlarging the perception field in BEV space. These streams are then combined before passing through the sigmoid function  $\sigma$ . This design aims to equip our module with the capability to capture both local and global attention for enhanced BEV probability prediction.

## 4 Experiments

#### 4.1 Datasets and Metrics

We conduct our experiments on the nuScenes dataset [1], a widely used benchmark for autonomous driving research. nuScenes offers comprehensive sensor data captured in urban driving scenarios, facilitating robust evaluation of object detection algorithms. NuScenes Detection Score (NDS) [1] serves as the official metric to measure the quality of 3D detection. NDS integrates mean Average Precision (mAP), 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), providing a holistic measure of detection quality across various aspects of performance.

## 4.2 Implementation Details

We employ ResNet-50 [4] operating at an image resolution of  $704 \times 256$ . During the 20-epoch training phase with the CBGS [39] and a batch size of 64, we

Method	Backbone	$mAP\uparrow$	$\mathrm{NDS}\uparrow$	mATE↓	mASE↓	mAOE↓	mAVE↓	$\mathrm{mAAE}{\downarrow}$
DETR3D [30]	V2-99	41.2	47.9	0.641	0.255	0.394	0.845	0.133
BEVDet4D [6]	Swin-B	45.1	56.9	0.511	0.241	0.386	0.301	0.121
UVTR [12]	V2-99	47.2	55.1	0.577	0.253	0.391	0.508	0.123
BEVFormer [15]	V2-99	48.1	56.9	0.582	0.256	0.375	0.378	0.126
BEVDepth [14]	V2-99	50.3	60.0	0.445	0.245	0.378	0.320	0.126
PETRv2 [20]	V2-99	50.8	59.1	0.543	0.241	0.360	0.367	0.118
Sparse4D [17]	V2-99	51.1	59.5	0.533	0.263	0.369	0.317	0.124
BEVStereo [13]	V2-99	52.5	61.0	0.431	0.246	0.358	0.357	0.138
SA-BEV [10]	V2-99	53.3	62.4	0.430	0.241	0.338	0.282	0.139
FB-BEV [16]	V2-99	53.7	62.4	0.439	0.250	0.358	0.270	0.128
SOLOFusion [25]	ConvNeXt-B	54.0	61.9	0.453	0.257	0.376	0.276	0.148
SparseBEV [18]	V2-99	54.3	62.7	0.502	0.244	0.324	0.251	0.126
DualBEV	V2-99	55.2	63.4	0.414	0.245	0.377	0.252	0.129

**Table 3:** Comparison with previous state-of-the-art BEV detectors on the nuScenes test set. Swin-B [21] and ConvNeXt-B [22] don't use extra data for depth training.

utilize the AdamW [23] optimizer with a learning rate set to  $2 \times 10^{-4}$ . Data augmentation techniques consistent with those of BEVDet [7] are applied. Our BEV grid size is set to  $128 \times 128$  without building voxel features. For the test set, we leverage VoVNet [11] at an image resolution of  $1600 \times 640$ , with the BEV grid size adjusted to  $256 \times 256$ . The model is trained for only 8 epochs with CBGS, and 8 previous keyframes are applied for temporal module following BEVDet4D [6]. All latency tests are conducted on a single NVIDIA 3090 GPU.

## 4.3 Main Results

**nuScenes val set** We adopt BEVDepth [14] with BEVPoolv2 [5] as baseline method. As illustrated in Tab. 2, DualBEV outperforms the baseline method with +1.0% mAP and +1.8% NDS in single frame. For multi-frame, DualBEV not only surpasses the baseline with +1.2% mAP and +1.9% NDS, but also outperforms BEVStereo [13] which uses temporal-stereo module for precise depth estimation. By leveraging HT features which capture the distant information, and efficient fusion with LSS feature, we also have a 0.5\% NDS increase compared to SA-BEV [10] and a 0.6\% NDS increase compared to FB-BEV [16].

**nuScenes test set** Tab. 3 demonstrates that DualBEV achieves state-of-the-art performance, yielding a remarkable 55.2% mAP and 63.4% NDS compared to prior works in VT. Notably, our approach exhibits a significant margin of 1.9% in mAP over SA-BEV [10] and 1.5% over FB-BEV [16], while also surpassing them by 1.0% in NDS. Additionally, when compared to SOLOFusion [25], which focus on temporal module, DualBEV maintains superiority with a lead of 1.2% in mAP and 1.5% in NDS even without temporal-stereo module. Furthermore, DualBEV outperforms the previous state-of-the-art Transformer-based VT approach SparseBEV [18], by a notable margin of 0.9% in mAP and 0.7% in NDS, proving the pivatol role of unveiling precise correspondences in VT.

#### 12 P. Li et al.

Table 4: Ablation study of different components on the nuScenes val set.

$Method \qquad  mAP\uparrow NDS\uparrow  latency\downarrow$					$P_{proj}$ .	$P_{img}$	$P_{bev} _{1}$	$mAP\uparrow$	$NDS\uparrow$	
Baseline [14]	34.2 40	.7   10	$0.2 \mathrm{ms}$		$\checkmark$			37.0	48.8	
+ ProbNet	34.6 41	.1 1	$11.1 \mathrm{ms}$		$\checkmark$	$\checkmark$		37.7	49.9	
+ HeightTrans	35.0 41	.5 1	$11.3 \mathrm{ms}$		$\checkmark$		$\checkmark$	37.4	49.6	
+ CAF	<b>35.3</b> 41	.6 1	11.5 ms			$\checkmark$	$\checkmark$	27.0	40.0	
+ SAE	35.2 <b>42</b>	<b>42.5</b> 11.7ms			$\checkmark$	$\checkmark$	$\checkmark$	<b>38.0</b>	50.4	
(a) Component Ablation.					(b) Probability Ablation.					
VT Operation	$ mAP\uparrow$	NDS↑	latency↓		Distrib	ution	Num	mAP	† NDS↑	
$SCA_{da}$ [16]	29.5	36.6	12.50ms		Unifo	$\mathbf{rm}$	4	35.0	41.9	
Bilinear-Sampling	[3] 30.5	37.6	7.15ms		Unifo	Uniform		35.3	42.2	
BEV Pooling [5]	30.7	38.2	0.32ms		Uniform		16	35.3	42.3	
Prob-Sampling	30.6	39.0	0.32ms		MF	8	13	35.2	42.5	
(c) VT Operation Comparison.					(d) Sa:	mplin	g Strat	egy in l	Height.	
Method $ P_{proj} P_{img} P_{bev}  \text{mAP} \uparrow \text{NDS} \uparrow$				-	Strategy	y Pr	ocess	mAP	$\uparrow \text{NDS}\uparrow$	
BEVDet [7] $\checkmark$		30.	7 38.2	-	Refine	two	-stage	e 35.1	41.7	
SA-BEV [10] ✓	$\checkmark$	30.9	9 38.7		Add	one	-stage	e 35.2	42.1	
Prob-LSS 🗸 🗸	$\checkmark$ $\checkmark$	31.	$5 \ 39.1$		DFF	one	-stage	35.2	42.5	
(e) Prob-LSS Ablation.				-		( <b>f</b> ) Fu	ision S	trategy.		

### 4.4 Ablation Study

Impact of Each Component We first explore each component's impact in Tab. 4(a). Starting with BEVDepth [14] as the baseline, the initial integration of ProbNet to provide BEV probability results in a 0.4% increase in both mAP and NDS. Combining HT features led to a 0.4% mAP and 0.4% NDS increase, show-casing compensation for Prob-LSS. Subsequently, we employ the CAF module to softly select features from the two streams, resulting in a further 0.3% mAP and 0.1% NDS increase. Finally, by incorporating the SAE module for improved BEV probability prediction with the addition of spatial attention module along-side ProbNet, our method achieves a substantial 0.9% improvement in NDS. The overall components gain a 1.0% mAP and 1.8% NDS improvement. Additionally, we also test the latency of VT from image feature to BEV feature including depth estimation. The entire proposed components reveal an extra latency of merely 1.5ms, with ProbNet introducing 0.9ms of this total.

**Impact of Each Probability** Tab. 4(b) demonstrates the impact of each probabilistic measurement for DualBEV. Alongside the projection probability  $P_{proj}$ , the inclusion of image probability  $P_{img}$  yields a notable improvement of 0.7% mAP and 1.1% NDS, while the introduction of BEV probability  $P_{bev}$  contributes

to a 0.4% mAP and 0.8% NDS enhancement. This observation implies that features of background pixels in the image space have a comparable detrimental effect to features of blank grids in the BEV space. An even more significant performance boost of 1.0% mAP and 1.6% NDS is achieved when both probabilities are applied simultaneously. However, disabling  $P_{proj}$  by adopting a uniform distribution for depth prediction per pixel leads to a significant decrease in performance, indicating the critical importance of  $P_{proj}$ , which serves as the cornerstone of the LSS pipeline.

Effect of Prob-Sampling In Tab. 4(c), we compare Prob-Sampling with other methods using BEVDet [7] as the base detector, focusing solely on VT operation. Regarding accuracy, Prob-Sampling outperforms 1-layer depth-aware spatial cross-attention (SCA<sub>da</sub>) proposed by FB-BEV [16] by 2.4% in NDS, while achieving a 1.4% NDS enhancement compared to bilinear sampling used in Simple-BEV [3]. Additionally, Prob-Sampling shows a 0.8% NDS improvement over BEV Pooling [5]. In terms of inference latency, Prob-Sampling's adoption of pre-computation for acceleration allows it to achieve comparable speed to BEV Pooling, representing a remarkable more than 20 times speed improvement over bilinear sampling and 40 times improvement over SCA<sub>da</sub>. This highlights Prob-Sampling's competitiveness in both accuracy and latency for BEV detection.

Effect of Sampling Strategy in Height In HeightTrans, we apply a multiresolution (MR) sampling strategy within the height range. We evaluate different sampling strategies as shown in Tab. 4(d). The performance almost reaches saturation after 8 uniform points. Notably, our proposed approach achieves a 0.6% increase in NDS compared to 4 uniform points utilized in BEVFormer [15], and even a 0.2% enhancement compared to 16 uniform points along the height.

Effect of Prob-LSS Our proposed Prob-LSS extends the representation of LSS approach as illustrated in Tab. 1. We further examine the evolution of LSS methods starting from BEVDet [7] without any auxiliary loss in Tab. 4(e). SA-BEV [10] expands upon the base version by introducing instance segmentation to filter out irrelevant information in the image space. This addition results in a 0.2% increase in mAP and a 0.5% increase in NDS. Our method further extends the idea to BEV space to mitigate the impact of irrelevant BEV features, particularly those generated by inaccurate depth estimation. This extension leads to a notable improvement of 0.6% in mAP and 0.4% in NDS.

Effect of Fusion Strategy Tab. 4(f) compares different fusion strategy for HT feature and Prob-LSS feature. We first explore the refinement strategy employed in FB-BEV [16], which utilizes Prob-LSS features to predict the BEV probability for HeightTrans. However, the refinement strategy yields even worse results than directly summing the both features for BEV probability prediction, with a 0.4% NDS decrease. Our proposed DFF strategy achieves a notable improvement, surpassing the refinement strategy by 0.8% in NDS.



Fig. 6: Visualization of detection results on image and BEV space.

## 4.5 Visualization

In Fig. 6, we present a qualitative comparison between BEVDepth [14] and Dual-BEV. Our approach exhibits superior performance in close and medium ranges, effectively eliminating false detections and accurately captures the curve of barriers even break (blue dashed rectangles). In the distant range, our method also recalls missed objects (orange dashed rectangles), benefiting from the compensation of HeightTrans. Additionally, our method provides precise information about small objects (purple dashed circles), which is not evident from BEV.

# 5 Conclusion and Limitation

In this work, we present a novel approach to unify feature transformation suitable for both 3D-to-2D and 2D-to-3D VT, coupled with pre-computation for speed enhancement. Leveraging CNN-based probabilistic correspondences, HeightTrans and Prob-LSS extend the capabilities of respective VT methods effectively. Through one-stage fusion of dual features using DFF, DualBEV captures the essence of VT and demonstrates the effectiveness of unveiling precise correspondences for BEV representation. Furthermore, our approach is versatile and applicable to tasks such as BEV segmentation or 3D occupancy prediction.

However, our framework's current design exclusively derives all probabilities from the current frame, neglecting historical information and underutilizing the temporal module. Furthermore, our framework heavily depends on depth estimation, as demonstrated by a significant performance drop observed after switching to a uniform distribution, as shown in Tab. 4(c).

# Acknowledgements

The authors would like to thank reviewers for their detailed comments and instructive suggestions. This research is supported by National Key R&D Program of China (2022YFB4300300).

## References

- 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: CVPR (2020)
- Dai, Y., Gieseke, F., Oehmcke, S., Wu, Y., Barnard, K.: Attentional feature fusion. In: WACV (2021)
- 3. Harley, A.W., Fang, Z., Li, J., Ambrus, R., Fragkiadaki, K.: Simple-BEV: What really matters for multi-sensor bev perception? In: ICRA (2023)
- He, K., Zhang, X., Ren, S., Sun, J.: Deep residual learning for image recognition. In: CVPR (2016)
- 5. Huang, J., Huang, G.: Bevpoolv2: A cutting-edge implementation of bevdet toward deployment. arXiv:2211.17111 (2022)
- Huang, J., Huang, G.: Bevdet4d: Exploit temporal cues in multi-camera 3d object detection. arXiv:2307.11477 (2023)
- Huang, J., Huang, G., Zhu, Z., Ye, Y., Du, D.: Bevdet: High-performance multicamera 3d object detection in bird-eye-view. arXiv:2112.11790 (2021)
- Huang, P., Liu, L., Zhang, R., Zhang, S., Xu, X., Wang, B., Liu, G.: Tig-bev: Multiview bev 3d object detection via target inner-geometry learning. arXiv:2212.13979 (2022)
- Jiang, Y., Zhang, L., Miao, Z., Zhu, X., Gao, J., Hu, W., Jiang, Y.G.: Polarformer: Multi-camera 3d object detection with polar transformers. In: AAAI (2023)
- 10. Jinqing, Z., Yanan, Z., Qingjie, L., Yunhong, W.: Sa-bev: Generating semanticaware bird's-eye-view feature for multi-view 3d object detection. In: ICCV (2023)
- 11. Lee, Y., Park, J.: Centermask: Real-time anchor-free instance segmentation. In: CVPR (2020)
- 12. Li, Y., Chen, Y., Qi, X., Li, Z., Sun, J., Jia, J.: Unifying voxel-based representation with transformer for 3d object detection. In: NeurIPS (2022)
- Li, Y., Bao, H., Ge, Z., Yang, J., Sun, J., Li, Z.: Bevstereo: Enhancing depth estimation in multi-view 3d object detection with dynamic temporal stereo. In: AAAI (2023)
- 14. Li, Y., Ge, Z., Yu, G., Yang, J., Wang, Z., Shi, Y., Sun, J., Li, Z.: Bevdepth: Acquisition of reliable depth for multi-view 3d object detection. In: AAAI (2023)
- 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: ECCV (2022)
- 16. Li, Z., Yu, Z., Wang, W., Anandkumar, A., Lu, T., Alvarez, J.M.: FB-BEV: BEV representation from forward-backward view transformations. In: ICCV (2023)
- Lin, X., Lin, T., Pei, Z., Huang, L., Su, Z.: Sparse4d: Multi-view 3d object detection with sparse spatial-temporal fusion. arXiv:2211.10581 (2022)
- Liu, H., Teng, Y., Lu, T., Wang, H., Wang, L.: Sparsebev: High-performance sparse 3d object detection from multi-camera videos. In: ICCV (2023)

- 16 P. Li et al.
- 19. Liu, Y., Wang, T., Zhang, X., Sun, J.: Petr: Position embedding transformation for multi-view 3d object detection. In: ECCV (2022)
- Liu, Y., Yan, J., Jia, F., Li, S., Gao, Q., Wang, T., Zhang, X., Sun, J.: Petrv2: A unified framework for 3d perception from multi-camera images. In: ICCV (2023)
- 21. 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: ICCV (2021)
- Liu, Z., Mao, H., Wu, C.Y., Feichtenhofer, C., Darrell, T., Xie, S.: A convnet for the 2020s. In: CVPR (2022)
- 23. Loshchilov, I., Hutter, F.: Decoupled weight decay regularization. In: ICLR (2019)
- 24. Milletari, F., Navab, N., Ahmadi, S.A.: V-net: Fully convolutional neural networks for volumetric medical image segmentation. In: 3DV (2016)
- Park, J., Xu, C., Yang, S., Keutzer, K., Kitani, K., Tomizuka, M., Zhan, W.: Time will tell: New outlooks and a baseline for temporal multi-view 3d object detection. In: ICLR (2023)
- Philion, J., Fidler, S.: Lift, splat, shoot: Encoding images from arbitrary camera rigs by implicitly unprojecting to 3d. In: ECCV (2020)
- 27. Reading, C., Harakeh, A., Chae, J., Waslander, S.L.: Categorical depth distributionnetwork for monocular 3d object detection. In: CVPR (2021)
- 28. Roddick, T., Kendall, A., Cipolla, R.: Orthographic feature transform for monocular 3d object detection. In: BMVC (2019)
- Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A.N., Kaiser, L., Polosukhin, I.: Attention is all you need. In: NeurIPS (2017)
- Wang, Y., Guizilini, V.C., Zhang, T., Wang, Y., Zhao, H., Solomon, J.: Detr3d: 3d object detection from multi-view images via 3d-to-2d queries. In: CoRL (2022)
- Woo, S., Park, J., Lee, J.Y., Kweon, I.S.: Cbam: Convolutional block attention module. In: ECCV (2018)
- Wu, Y., Li, R., Qin, Z., Zhao, X., Li, X.: Heightformer: Explicit height modeling without extra data for camera-only 3d object detection in bird's eye view. arXiv:2307.13510 (2023)
- 33. Xia, Z., Pan, X., Song, S., Li, L.E., Huang, G.: Vision transformer with deformable attention. In: CVPR (2022)
- 34. Xiaowei, C., Jiaming, L., Lu, M., Zhang, R., Wang, Z., Guo, Y., Zhang, S.: Bev-san: Accurate bev 3d object detection via slice attention networks. In: CVPR (2023)
- Xie, E., Yu, Z., Zhou, D., Philion, J., Anandkumar, A., Fidler, S., Luo, P., Alvarez, J.M.: M<sup>2</sup>BEV: Multi-camera joint 3D detection and segmentation with unified birds-eye view representation. arXiv:2204.05088 (2022)
- Xiong, K., Gong, S., Ye, X., Tan, X., Wan, J., Ding, E., Wang, J., Bai, X.: Cape: Camera view position embedding for multi-view 3d object detection. In: CVPR (2023)
- 37. Yang, C., Chen, Y., Tian, H., Tao, C., Zhu, X., Zhang, Z., Huang, G., Li, H., Qiao, Y., Lu, L., Zhou, J., Dai, J.: Bevformer v2: Adapting modern image backbones to bird's-eye-view recognition via perspective supervision. In: CVPR (2023)
- Zhou, H., Ge, Z., Li, Z., Zhang, X.: Matrixvt: Efficient multi-camera to bev transformation for 3d perception. In: ICCV (2023)
- Zhu, B., Jiang, Z., Zhou, X., Li, Z., Yu, G.: Class-balanced grouping and sampling for point cloud 3d object detection. arXiv:1908.09492 (2019)