

EMIFF: Enhanced Multi-scale Image Feature Fusion for Vehicle-Infrastructure Cooperative 3D Object Detection

Zhe Wang¹, Siqi Fan¹, Xiaoliang Huo^{1,2}, Tongda Xu¹, Yan Wang^{1*}, Jingjing Liu¹, Yilun Chen¹, Ya-Qin Zhang^{1*}

Abstract—In autonomous driving, cooperative perception makes use of multi-view cameras from both vehicles and infrastructure, providing a global vantage point with rich semantic context of road conditions beyond a single vehicle viewpoint. Currently, two major challenges persist in vehicle-infrastructure cooperative 3D (VIC3D) object detection: 1) inherent pose errors when fusing multi-view images, caused by time asynchrony across cameras; 2) information loss in transmission process resulted from limited communication bandwidth. To address these issues, we propose a novel camera-based 3D detection framework for VIC3D task, *Enhanced Multi-scale Image Feature Fusion* (EMIFF). To fully exploit holistic perspectives from both vehicles and infrastructure, we propose *Multi-scale Cross Attention* (MCA) and *Camera-aware Channel Masking* (CCM) modules to enhance infrastructure and vehicle features at scale, spatial, and channel levels to correct the pose error introduced by camera asynchrony. We also introduce a *Feature Compression* (FC) module with channel and spatial compression blocks for transmission efficiency. Experiments show that EMIFF achieves SOTA on DAIR-V2X-C datasets, significantly outperforming previous early-fusion and late-fusion methods with comparable transmission costs.

I. INTRODUCTION

Subject to sensor limitations, autonomous vehicles lack a global perception capability for monitoring holistic road conditions and accurately detecting surrounding objects, which bears great safety risks [1], [2]. Vehicle-to-everything (V2X) [3], [4] aims to build a communication system between vehicles and other devices in a complex traffic environment. Vehicle and infrastructure cooperation can significantly expand the perception range and improve perception capability [5]. Cameras from both two sides provide a global vantage point with a rich semantic context of road conditions beyond a single-vehicle viewpoint [6]. Vehicle-infrastructure cooperative 3D object detection (VIC3D) from cameras is a significant task for autonomous driving.

Compared with vanilla single-vehicle 3D object detection, VIC3D tasks face more unique challenges. One challenge is inherent pose error when fusing multi-view images from vehicles and those from infrastructure [7], caused by time asynchrony across agents [8]. As shown in Figure 1, this




Fig. 1. Labels (3D bounding boxes) projected from 3D space to vehicle (a) and infrastructure (b) image planes using calibration parameters $P_{inf/veh}$ often suffer from misalignment between the ground truth and the projection position in 2D images (as illustrated by the misaligned green bounding boxes). The reason for this misalignment is that the camera’s capture time $T_{inf/veh}$ are different and the moving object captured from the vehicle camera (in green) and infrastructure camera (in red) will appear at different locations.

pose errors can result in inaccurate relative positions between objects and annotations. Another challenge is limited communication bandwidth between agents resulting in information loss between transmissions [3]. The raw sensor data possesses ample information required for fusion; however, it necessitates greater bandwidth, thus necessitating fusion methods to prioritize the balance between performance and transmission cost. Therefore, fusion methods to tackle such cross-agent perception challenges are the key to VIC3D.

Many fusion works on V2X are proposed based on simulated datasets, such as OPV2V [3], V2X-Sim [9] and V2XSet [10], which neglect above challenges and have a sim-to-real gap. Most existing research only focused on LiDAR-based methods due to the fusion convenience and the performance advantage, such as *early fusion* (EF) of raw signals [5], [11], [12], *intermediate fusion* (IF) of features [4], [3], [13], [14], and *late fusion* (LF) of prediction outputs [5], [15]. But due to the projection gap between 2D image plane and 3D space, image fusion can not be as direct as point clouds. In real scenarios, DAIR-V2X [5] adopts an LF method by combining prediction outputs from each camera, which is sensitive to calibration so that even when prediction from the infrastructure side is perfect, the vehicle will receive biased 3D detection.

In this paper, we propose a novel framework for VIC3D

¹Zhe Wang, Siqi Fan, Tongda Xu, Yan Wang*, Jingjing Liu, Yilun Chen, and Ya-Qin Zhang* are with the Institute for AI Industry Research (AIR), Tsinghua University, Beijing, China. {wangzhe, fansiqi, xutongda, wangyan }@air.tsinghua.edu.cn

²Xiaoliang Huo is with the School of Software, Beihang University, Beijing, China. huoxiaoliangchn@buaa.edu.cn

*Correspondence to Yan Wang, Ya-Qin Zhang.

†Code will be released at <https://github.com/Bosszhe/EMIFF>

task, *Enhanced Multi-scale Image Feature Fusion* (EMIFF). We choose intermediate fusion since it doesn't highly rely on accurate calibration parameters. For feature-level fusion, high-dimensional features extracted from raw data can be compressed, transmitted, and dynamically enhanced [8], which can be used to alleviate the negative effect of pose errors. We design modules to compress transmitted features to reduce transmission cost and enable feature enhancement in scale level, spatial level, and channel level.

Specifically, Feature Compression (FC) module compresses 2D features transmitted from the infrastructure to vehicle. Since the receptive field is larger in smaller-scale features, which theoretically has higher tolerance to slight location errors, Multi-scale Cross Attention (MCA) module aims to achieve attentive scale-wise feature selection between features. MCA also corrects features at the spatial level with attentive offset to overcome pixel-wise shift caused by pose errors. To correct location errors born from multiple cameras, features are further enhanced by a Camera-aware Channel Masking (CCM) module via a learned channel-wise mask following the guidance of camera parameters. Then, the enhanced features are transformed into voxel features leveraging calibration parameters. Finally splatted into BEV space, the fused feature is fed into detection heads for object detection. Experiments demonstrate the effectiveness of each EMIFF module in reducing pose errors and achieving better prediction accuracy than existing EF and LF methods. Our contributions can be summarized as follows:

- We propose EMIFF, a novel framework for camera-based VIC3D object detection, using an intermediate fusion method to tackle cross-agent perception challenges.
- We design MCA and CCM modules to dynamically enhance image features for better detection performance, with an additional FC module to reduce transmission costs in VIC3D system.
- We achieve state-of-the-art results on DAIR-V2X-C dataset, the latest VIC3D benchmark with real data, where EMIFF outperforms existing LF and EF methods with comparable transmission costs.

II. RELATED WORK

A. V2X Cooperative Perception

Current research on V2X cooperative perception mainly focuses on simulated datasets, such as OPV2V [3], V2X-Sim [9] and V2XSet [10]. Existing intermediate-fusion methods focused on simulated point clouds, such as V2VNet [13] which transmitted compressed features to nearby vehicles and generated joint perception/prediction. DiscoNet [4] introduced graphs into feature fusion and proposed edge weights to highlight different informative regions during feature propagation. Recent Where2comm [11] considered the spatial confidence of features and selected features with high confidence and complementary to others, which effectively saves transmission costs. FFNET [8] introduces the concept of feature flow to tackle the issue of temporal asynchrony in real-world scenarios, albeit its applicability is limited to point

cloud data. Different from point clouds, images from vehicle and infrastructure have a huge view gap, thus features need to be transformed into unified space for fusion. One direct way for fusing multi-view images is late fusion, such as DAIR-V2X [5], which proposed a result-level fusion model for cameras with separate detectors [16]. Few approaches have focused on IF methods for cameras, especially in real scenarios.

B. Camera-based Feature Fusion

Direct Prediction methods extract image features with object query [17], [18], [19], [20] or directly on front-view image [21]. DETR3D [17] used a sparse set of 3D object queries to sample 2D multi-view image features and predicted 3D bounding boxes with set-to-set loss. PETR [19], [20] transformed image features into 3D position-aware representation by encoding 3D coordinates into position embedding. FCOS3D [21] transformed 3D labels to front-view images and directly predicted 3D information by extending FCOS [22] to 3D detection.

Lift-based methods project features from image plane to BEV (bird's eye view) plane through depth estimation. Most methods [23], [24], [25], [26], [27] applied 2D-to-3D transformation following LSS [28], which predicted a depth distribution for each pixel and lifted image features into frustum features with camera parameters, then splatted all frustums into a rasterized BEV feature. BEVDepth [29] claimed the quality of intermediate depth estimation is the key to improving multi-view 3D object detection and added explicit depth supervision with groundtruth depth generated from point clouds. PON [30] learned the transformation leveraging geometry relationship between image locations and BEV locations in the horizontal direction.

Projection-based methods generate dense voxel or BEV representation from image features through 3D-to-2D projection [1]. ImVoxelNet [16] aggregated the projected features from several images via a simple element-wise averaging, where spatial information might not be exploited sufficiently. Transformer-based methods [31], [32] mapped perspective view to BEV with designed BEV queries and leveraged cross- and self-attention to aggregate spatial and temporal information into BEV queries. Since global attention needs huge memory with high time cost, deformable attention was adopted in BEVFormer [31].

III. METHOD

EMIFF aims to fuse vehicle and infrastructure features by utilizing V2X communication. It includes four main modules: Feature Compression (FC), Multi-scale Cross Attention (MCA), Camera-aware Channel Masking (CCM), and Point-Sampling Voxel Fusion, as illustrated in Figure 2.

A. Feature Compression

The images from vehicle and infrastructure are denoted as I_{veh} and I_{inf} , respectively, and the shape of both images are $[H \times W \times 3]$. Since infrastructure cameras are typically installed at a higher elevation than vehicles, resulting in a

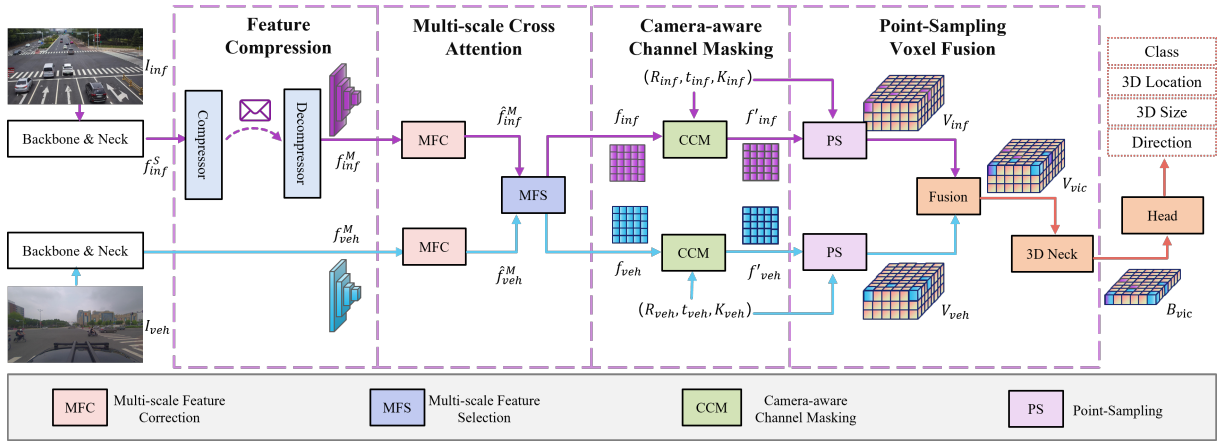


Fig. 2. The general framework of EMIFF. Separate image backbone and neck extract multi-scale image feature from vehicle and infrastructure images. FC module compresses source infrastructure feature f_{inf}^S and decompresses it to multi-scale ones f_{inf}^M . MCA module consisting of MFC and MFS blocks enhances multi-scale features $f_{veh/inf}^M$ by seeking the correlation between the two sides, and CCM takes camera parameters (R, t, K) as input to reweight features $f_{veh/inf}$ with channel relationship. Finally, Point-Sampling Voxel Fusion projects image features $f'_{veh/inf}$ into 3D space to generate a unified voxel feature V_{vic} , which can be applied to 3D neck and head in turn for detection prediction.

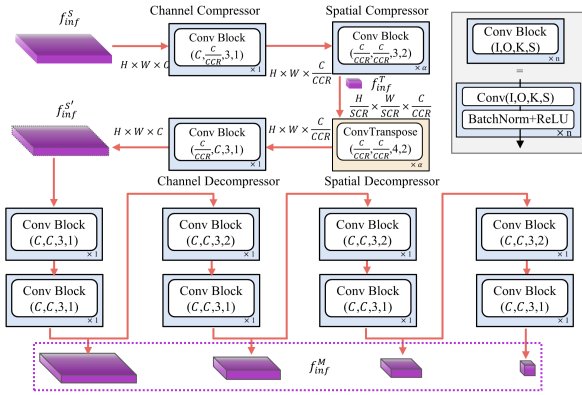


Fig. 3. Illustration of FC module. Feature f_{inf}^S is compressed into f_{inf}^T through the channel and spatial compressors, which is transmitted to vehicle and is decoded into f_{inf}^S through the channel and spatial decompressors. Finally, multi-scale infrastructure features f_{inf}^M can be recovered from f_{inf}^T with several Conv Blocks with stride 2.

huge view gap between images captured from vehicles and infrastructure, we use separate pre-trained backbones and necks on the vehicle and infrastructure respectively to extract multi-scale image features. The output multi-scale features can be denoted as $f_s^M, s = veh/inf$.

EMIFF transmits image features and camera parameters instead of voxel feature after projection because voxel feature is too large to be transmitted efficiently. The Feature Compression (FC) module (shown in Figure 3) compresses the largest infrastructure feature f_{inf}^0 (noted as f_{inf}^S) to f_{inf}^T , transmits f_{inf}^T to vehicle and regenerate multi-scale features f_{inf}^M through decompression.

B. Multi-scale Cross Attention

MCA module contains Multi-scale Feature Correction (MFC) and Multi-scale Feature Selection (MFS) blocks, as shown in Figure 4 and Figure 5. MFC block is designed

to select and integrate spatial-wise features with attentive offset, and thus pixel-wise shift caused by pose errors can be overcome to some extent. MFS block applies cross-attention between them to achieve attentive scale-wise feature selection.

The MFC module is first applied to multi-scale features. Since pose errors can cause a displacement between the projected and ground-truth positions on 2D plane, we apply DCN (deformable convolutional networks) [33] for each scale feature to allow every pixel to get spatial information surrounding it. Then, features at different scales are upsampled to the same size through UpConv blocks.

MFS applies MeanPooling operation to obtain the representation of different scales of infrastructure features, while vehicle features at different scales are first fused by mean operation and then refined by MeanPooling. To find the correlation between vehicle features and infrastructure features at different scales, cross attention is applied to infrastructure representations as Key and vehicle representation as Query, which generates attention weights ω_{inf}^m for each scale m . We calculate inter-product between features f_{inf}^M and weights ω_{inf}^m . The final outputs of MCA are augmented infrastructure image feature f_{inf} and vehicle image feature f_{veh} .

C. Camera-aware Channel Masking

Considering that objects closer to the camera are easier to detect and it's common for the same object to be closer to an infrastructure camera but far away from a vehicle. Since different channels represent object information at different distances, which is strongly correlated with camera parameters, it is intuitive to take camera parameters as priors to augment image features.

Inspired by the decoupled nature of SENet [34] and LSS [28], CCM will learn a channel-wise mask to weigh the importance between the channels. First, camera intrinsic and extrinsic are stretched into one dimension and concatenated

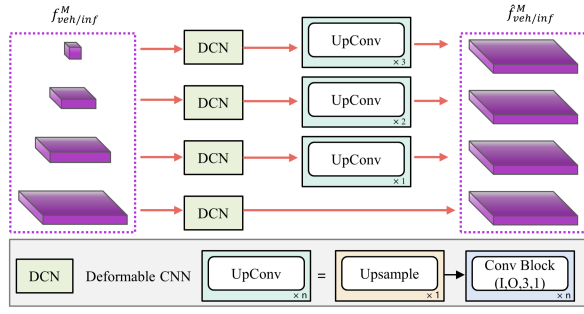


Fig. 4. Details of MFC. Every pixel-wise feature is integrated with the spatial information of surrounding pixels via DCN, and multi-scale features are scaled to the same size through UpConv blocks.




Fig. 5. Schema of MCA module. In the lower branch, vehicle feature f_{veh} is generated from f^M_{veh} through MFC Block and Mean. In the upper branch, f^M_{inf} is refined into ‘key’ through MFC Block and Mean Pooling, and queries are generated from f_{veh} through Mean Pooling. The output weights ω^m_{inf} of cross-attention are applied to f^M_{inf} with inner product to form infrastructure feature f_{inf} .

together. Then, they are scaled up to the feature’s dimension C using MLP to generate a channel mask $M_{veh/inf}$. Finally, $M_{veh/inf}$ is used to re-weight the image features $f_{veh/inf}$ in channel-wise and obtain results $f'_{veh/inf}$. The overall CCM module can be written as:

$$\begin{aligned} f'_s &= M_s \odot f_s, s = veh, inf \\ m_s &= \text{MLP}(\xi(R_s) \oplus \xi(t_s) \oplus \xi(K_s)) \end{aligned} \quad (1)$$

ξ denotes the flat operation and \oplus means concatenation. The input of MLP is the combination of camera rotation matrix $R_s \in \mathbb{R}^{3 \times 3}$, translation t_s and camera intrinsics K_s . M_s can be obtained from MLP’s output m_s through several Fully Connected (FC) layers and Activation (Relu, Sigmoid) layers.




Fig. 6. The schema of CCM module.

D. Point-Sampling Voxel Fusion

The augmented vehicle feature f'_{veh} and infrastructure feature f'_{inf} are projected into 3D space for fusion and

generate two voxel features, denoted as V_{veh} and V_{inf} , respectively. The details of projection principle are following ImVoxelNet [16].

We obtain the final voxel feature $V_{vic} \in N_x \times N_y \times N_z \times C_1$ by averaging sampled voxel features V_{veh} and V_{inf} . Then, the same 3D neck as [16], which is composed of 3D CNN and downsampling layers, transforms voxel feature V_{vic} into BEV feature $B_{vic} \in N_x \times N_y \times C_2$. BEV feature can be used as input of common 2D detection heads to predict 3D detection results. The loss of detection heads is similar to SECOND [35], which consists of smooth L1 Loss for bounding box L_{bbox} , focal loss for classification L_{cls} , and cross-entropy loss for direction L_{dir} .

IV. EXPERIMENTS

A. Implementation Details

Datasets. We conduct our experiments on a vehicle-infrastructure-cooperation dataset DAIR-V2X [5], in which all frames are captured from real scenarios. We utilize the VIC-Sync portion of DAIR-V2X-C dataset for training and evaluation, which is composed of 9,311 pairs of infrastructure and vehicle frames captured at the same time. Annotations of each pair frame are in world coordinate and need to be translated into vehicle coordinate system for training and evaluation.

Evaluation Metrics. Evaluation metrics are Average Precision (AP) [36] and Average Byte (AB) to measure detection performance and transmission cost, the same as [5]. The AP metrics are based on the detection range surrounding the vehicle, including Overall (0-100m), 0-30m, 30-50m, and 50-100m. All metrics are calculated with IoU = 0.5 and can be divided into 2 parts: AP_{3D} and AP_{BEV} . AB means the average size of transmitted data. It is the feature map f_{inf}^T in our method.

Training. We use ResNet-50 [37] as backbone and FPN [38] as 2D neck to extract image features. The channel number C of the neck’s output is 64. We set the channel of 3D voxel feature C_1 to 64 and the channel of BEV feature C_2 to 256 following [35], [39].

B. Object Detection Results

We compare the performance of baseline Late Fusion (LF) methods with ImVoxelNet and our proposed single-side model *EMIFF_Veh/Inf* on DAIR-V2X-C dataset. We also implement several multi-view camera-based methods that have been applied to nuScenes dataset [40], [36] (eg., BEVFormer [31], BEVDepth [29]) for VIC3D task. The evaluation results on VIC-Sync portion of DAIR-V2X-C dataset are shown in Table I and Figure 7. From the table, Intermediate Fusion (IF) method EMIFF has achieved state-of-the-art performance on the multi-view camera fusion benchmark, compared with other methods of Late Fusion (LF) and Early Fusion (EF). EMIFF obtains 15.61 AP_{3D} and 21.44 AP_{BEV} in overall setting.

EMIFF_Veh and *EMIFF_Inf* remove the MCA module but preserve CCM and FC modules so that models can be applied to the vehicle side and infrastructure side respectively

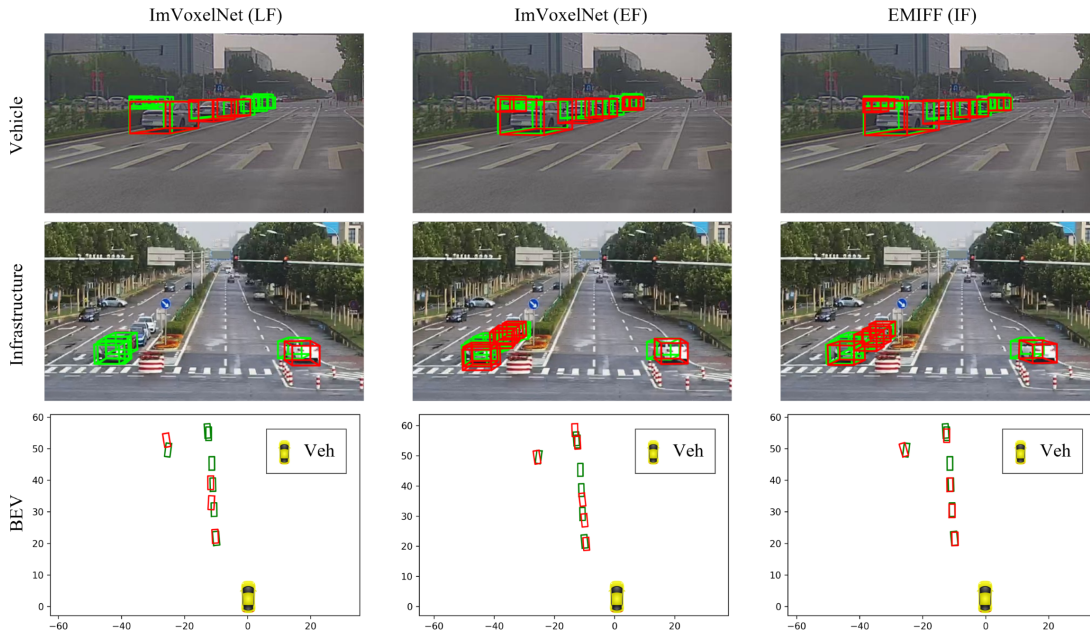


Fig. 7. Visualization results of ImVoxelNet (LF) (left column), ImVoxelNet_M (EF) (middle column), and EMIFF (IF) (right column). Bounding boxes in BEV (bottom row) are projected to vehicle and infrastructure image planes (top two rows). Groundtruth are in green and predictions in red. From BEV, it is clear that red and green bounding boxes from EMIFF are better aligned than LF and EF methods. This shows that ImVoxelNet (LF) and ImVoxelNet_M (EF) have detected more false positive objects and fewer true positive objects than EMIFF (IF).

without interaction between them, and predictions can be used for Late Fusion. EMIFF achieves higher AP_{3D} and AP_{BEV} compared with ImVoxelNet [16] under the setting of Only-Veh, Only-Inf, and LF. This indicates that EMIFF’s single-side model has a stronger feature extraction ability.

What is interesting is that Only-Inf methods achieve the best scores in 50-100m AP_{3D} and AP_{BEV} and this phenomenon can also be seen in [5]. As mentioned before, these metrics are related to detecting objects far from the ego vehicle. We count 16,934 objects within the distance range of 50-100m from vehicle, which are used to calculate the metric of 50-100m AP_{3D} . Among these objects, almost three-quarters (12,651) objects are closer to infrastructure camera, which are easier to be detected by Only-Inf models.

We also compared EMIFF with some representative cooperative perception models on DAIR-V2X-C datasets and experimental results are reported in Table II. The performance advantage of EMIFF is significant and it outperforms DiscoNet [4] by 23.69%. (* means results are from paper [41]).

C. Ablation Study

We remove MCA, CCM, and FC modules in EMIFF and regard it as baseline in the ablation study. We also conduct experiments to investigate when to fuse information from vehicle and infrastructure.

Effect of Each Component. The ablation results on MCA, CCM, and FC modules are summarized in Table III. The 1st row model can be denoted as *EMIFF_B*, which removes MCA, FC, and CCM modules and only keeps the fusion methodology at feature level. Comparing the 2nd and 3rd rows with the 1st row, both MCA and CCM can improve

performance over baseline, and MCA has increased AP_{3D} and AP_{BEV} by 1.05 and 0.59, better than 0.38 and 0.18 increase induced by CCM module. These results demonstrate the validity of MCA, which selects more useful infrastructure features at different scales based on vehicle features with a cross-attention mechanism. FC is designed to eliminate redundant information included in features, while it can also improve detection performance. This is because FC module increases the depth of the whole network and introduces extra computation, which can be regarded as feature refinement.




Fig. 8. AP_{3D} ($IoU=0.5$) with respect to Compression Rate (shown as number \times). CCR is changed from $\times 1$ to $\times 64$ and SCR is set from $\times 1$ to $\times 256$ with CCR set to $\times 64$.

Voxel or BEV Fusion? To investigate when to fuse features in IF method (at voxel or BEV level), we compare the performance of EMIFF with *EMIFF_{BEV}*. The former belongs to the IF-Voxel pipeline while the latter belongs to the IF-BEV fusion pipeline, which condenses voxel features

Fusion	Model	AP_{3D} (IoU=0.5)				AP_{BEV} (IoU=0.5)				AB (Byte)
		Overall	0-30m	30-50m	50-100m	Overall	0-30m	30-50m	50-100m	
Only-Veh	ImVoxelNet [16]	7.29	16.98	2.35	0.13	8.85	19.89	3.44	0.28	\
	EMIFF_Veh	8.65	19.11	4.33	0.20	10.46	22.42	5.57	0.42	
Only-Inf	ImVoxelNet [16]	8.66	13.05	5.79	<u>5.50</u>	14.41	17.98	10.34	<u>11.19</u>	\
	EMIFF_Inf	9.76	13.59	6.90	6.63	14.81	18.78	<u>11.50</u>	11.43	
LF	ImVoxelNet [5]	11.08	22.27	4.40	2.33	14.76	27.02	7.13	4.73	0.28K
	EMIFF_Veh/Inf	11.99	<u>24.79</u>	6.08	2.30	15.79	30.39	8.50	4.84	0.28K
EF	BEVDepth [29]	7.36	16.23	1.79	0.18	13.17	26.42	5.00	4.82	550.84K
	BEVFormer_S [31]	8.80	18.07	3.71	1.76	13.45	24.76	6.46	4.63	
	ImVoxelNet [16]	<u>12.72</u>	23.63	<u>7.38</u>	3.11	<u>18.17</u>	<u>30.54</u>	11.39	7.00	
IF	EMIFF	15.61	29.12	9.07	4.01	21.44	36.24	13.51	8.28	32.64K

TABLE I

QUANTITATIVE EVALUATION ON DAIR-V2X-C. BEST VALUES ARE MARKED BY BOLD, AND THE SECOND BEST IS UNDERLINED. ALL SCORES IN %.

Fusion Modality	Model	AP_{3D}
Bounding Box	LateFusion [5] (CVPR'22)	11.08
	CoBEVT [6] (CoRL'22)	4.80
BEV	V2VNet* [13] (ECCV'20)	8.47
	When2com* [42] (CVPR'20)	9.84
	Where2comm* [11] (NeurIPS'22)	10.25
	V2X-ViT* [10] (ECCV'22)	10.75
	DiscoNet* [4] (NeurIPS'21)	12.62
Voxel	EMIFF	15.61

TABLE II

COMPARISON RESULTS OF REPRESENTATIVE COOPERATIVE METHODS ON DAIR-V2X-C. ALL SCORES IN %.

MCA	CCM	FC	AP_{3D}	AP_{BEV}
			13.60	20.05
	✓		13.98	20.23
✓			14.65	20.64
✓	✓		<u>15.27</u>	<u>21.03</u>
✓	✓	✓	15.61	21.44

TABLE III

ABLATION STUDY ON EMIFF.

V_{veh} and V_{inf} into BEV feature respectively with two 3D necks, and then two BEV features are averaged for fusion. Results (Table IV) show that fusion at the voxel level has better performance, which indicates that the transformation from voxel to BEV feature can cause higher information loss.

Fusion	Model	AP_{3D}	AP_{BEV}
LF	ImVoxelNet	11.08	14.76
EF	ImVoxelNet_M	<u>12.72</u>	<u>18.17</u>
IF (BEV)	EMIFF_BEV	11.50	16.23
IF (Voxel)	EMIFF	13.37	19.66

TABLE IV

ANALYSIS ON CHOICE OF FEATURE FUSION.

D. Influence of Model Capacities

To further explain the effect of the proposed modules, we replace the image backbone of EMIFF and the baseline model *EMIFF_B* from ResNet-50 to ResNet-101 to assess the extent of performance improvement that can be obtained by increasing the parameters and capacities. Results in Table V show that a more complex network with higher capacity

Backbone	Model	AP_{3D}	AP_{BEV}	Param	FLOPs
ResNet-50	EMIFF_B	13.60	20.05	47.82M	94.01G
ResNet-50	EMIFF	15.61	21.44	49.32M	123.76G
ResNet-101	EMIFF_B	14.06	20.56	85.79M	171.10G
ResNet-101	EMIFF	16.46	22.32	87.31M	201.46G

TABLE V

COMPARISON RESULTS OF MODEL CAPACITIES.

has better performance, while EMIFF with fewer parameters and lower FLOPs still outperforms baseline *EMIFF_B* with ResNet-101. This observation provides additional evidence supporting the effectiveness of the proposed modules.

E. Impact of Feature Compression

As seen in Figure 8, We investigate the effect of Channel Compressor and Spatial Compressor. First, we change Channel Compression Rate (CCR) from $\times 1$ to $\times 64$, and the model performance is almost stable at low compression rates, which indicates that channel compression can extract more useful information and remove redundancy. After CCR reaches the maximum, we continue to compress features with Spatial Compressor. The compression rate ranges from $\times 64$ to $\times 16384$. With compressed feature shapes getting smaller, the AP_{3D} declines from 15.33 to 12.63 but is still higher than LF, and the transmission cost has fallen to 0.51KB which is comparable to LF's cost.

V. CONCLUSIONS

EMIFF is a novel multi-view intermediate-fusion framework for camera-based VIC3D task. To correct the negative effect of pose errors and time asynchrony, we design a Multi-scale Cross-Attention module and Camera-aware Channel Masking module to fuse and augment multi-view features. EMIFF also effectively reduces transmission cost via Feature Compression, and has achieved state-of-the-art results on DAIR-V2X-C benchmark, significantly outperforming previous EF and LF methods. Future study points to extension of the framework to more data modalities.

ACKNOWLEDGMENT

This work is funded by the National Key R&D Program of China (2022ZD0115502) and Lenovo Research.

REFERENCES

- [1] Y. Ma, T. Wang, X. Bai, H. Yang, Y. Hou, Y. Wang, Y. Qiao, R. Yang, D. Manocha, and X. Zhu, "Vision-centric bev perception: A survey," *arXiv preprint arXiv:2208.02797*, 2022.
- [2] Y. Han, H. Zhang, H. Li, Y. Jin, C. Lang, and Y. Li, "Collaborative perception in autonomous driving: Methods, datasets and challenges," *arXiv preprint arXiv:2301.06262*, 2023.
- [3] R. Xu, H. Xiang, X. Xia, X. Han, J. Li, and J. Ma, "Opv2v: An open benchmark dataset and fusion pipeline for perception with vehicle-to-vehicle communication," in *2022 International Conference on Robotics and Automation (ICRA)*. IEEE, 2022, pp. 2583–2589.
- [4] E. Mehr, A. Jourdan, N. Thome, M. Cord, and V. Guittény, "Disconet: Shapes learning on disconnected manifolds for 3d editing," in *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 2019, pp. 3474–3483.
- [5] H. Yu, Y. Luo, M. Shu, Y. Huo, Z. Yang, Y. Shi, Z. Guo, H. Li, X. Hu, J. Yuan, and Z. Nie, "Dair-v2x: A large-scale dataset for vehicle-infrastructure cooperative 3d object detection," in *IEEE/CVF Conf. on Computer Vision and Pattern Recognition (CVPR)*, June 2022.
- [6] R. Xu, Z. Tu, H. Xiang, W. Shao, B. Zhou, and J. Ma, "Cobevt: Cooperative bird's eye view semantic segmentation with sparse transformers," *arXiv preprint arXiv:2207.02202*, 2022.
- [7] N. Vadivelu, M. Ren, J. Tu, J. Wang, and R. Urtasun, "Learning to communicate and correct pose errors," in *Proceedings of the 2020 Conference on Robot Learning*, ser. Proceedings of Machine Learning Research, J. Kober, F. Ramos, and C. Tomlin, Eds., vol. 155. PMLR, 16–18 Nov 2021, pp. 1195–1210.
- [8] H. Yu, Y. Tang, E. Xie, J. Mao, J. Yuan, P. Luo, and Z. Nie, "Vehicle-infrastructure cooperative 3d object detection via feature flow prediction," *arXiv preprint arXiv:2303.10552*, 2023.
- [9] Y. Li, D. Ma, Z. An, Z. Wang, Y. Zhong, S. Chen, and C. Feng, "V2x-sim: Multi-agent collaborative perception dataset and benchmark for autonomous driving," *IEEE Robotics and Automation Letters*, vol. 7, no. 4, pp. 10914–10921, 2022.
- [10] R. Xu, H. Xiang, Z. Tu, X. Xia, M.-H. Yang, and J. Ma, "V2x-vit: Vehicle-to-everything cooperative perception with vision transformer," in *Computer Vision—ECCV 2022: 17th European Conference, Tel Aviv, Israel, October 23–27, 2022, Proceedings, Part XXXIX*. Springer, 2022, pp. 107–124.
- [11] Y. Hu, S. Fang, Z. Lei, Y. Zhong, and S. Chen, "Where2comm: Communication-efficient collaborative perception via spatial confidence maps," *arXiv preprint arXiv:2209.12836*, 2022.
- [12] R. Chen, Y. Mu, R. Xu, W. Shao, C. Jiang, H. Xu, Z. Li, and P. Luo, "Co³: Cooperative unsupervised 3d representation learning for autonomous driving," *arXiv preprint arXiv:2206.04028*, 2022.
- [13] T.-H. Wang, S. Manivasagam, M. Liang, B. Yang, W. Zeng, and R. Urtasun, "V2vnnet: Vehicle-to-vehicle communication for joint perception and prediction," in *Computer Vision—ECCV 2020: 16th European Conference, Glasgow, UK, August 23–28, 2020, Proceedings, Part II 16*. Springer, 2020, pp. 605–621.
- [14] S. Fan, H. Yu, W. Yang, J. Yuan, and Z. Nie, "Quest: Query stream for vehicle-infrastructure cooperative perception," *arXiv preprint arXiv:2308.01804*, 2023.
- [15] W. Chen, R. Xu, H. Xiang, L. Liu, and J. Ma, "Model-agnostic multi-agent perception framework," *arXiv preprint arXiv:2203.13168*, 2022.
- [16] D. Rukhovich, A. Vorontsova, and A. Konushin, "Imvoxelnet: Image to voxels projection for monocular and multi-view general-purpose 3d object detection," in *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision*, 2022, pp. 2397–2406.
- [17] Y. Wang, V. C. Guizilini, T. Zhang, Y. Wang, H. Zhao, and J. Solomon, "Detr3d: 3d object detection from multi-view images via 3d-to-2d queries," in *Conference on Robot Learning*. PMLR, 2022, pp. 180–191.
- [18] X. Chen, T. Zhang, Y. Wang, Y. Wang, and H. Zhao, "Futr3d: A unified sensor fusion framework for 3d detection," *arXiv preprint arXiv:2203.10642*, 2022.
- [19] Y. Liu, T. Wang, X. Zhang, and J. Sun, "Petr: Position embedding transformation for multi-view 3d object detection," *arXiv preprint arXiv:2203.05625*, 2022.
- [20] Y. Liu, J. Yan, F. Jia, S. Li, Q. Gao, T. Wang, X. Zhang, and J. Sun, "Petrv2: A unified framework for 3d perception from multi-camera images," *arXiv preprint arXiv:2206.01256*, 2022.
- [21] T. Wang, X. Zhu, J. Pang, and D. Lin, "Fcos3d: Fully convolutional one-stage monocular 3d object detection," in *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 2021, pp. 913–922.
- [22] Z. Tian, C. Shen, H. Chen, and T. He, "Fcos: Fully convolutional one-stage object detection," in *Proceedings of the IEEE/CVF international conference on computer vision*, 2019, pp. 9627–9636.
- [23] J. Huang, G. Huang, Z. Zhu, and D. Du, "Bevdet: High-performance multi-camera 3d object detection in bird-eye-view," *arXiv preprint arXiv:2112.11790*, 2021.
- [24] J. Huang and G. Huang, "Bevdet4d: Exploit temporal cues in multi-camera 3d object detection," *arXiv preprint arXiv:2203.17054*, 2022.
- [25] E. Xie, Z. Yu, D. Zhou, J. Phillion, A. Anandkumar, S. Fidler, P. Luo, and J. M. Alvarez, "M² 2bev: Multi-camera joint 3d detection and segmentation with unified birds-eye view representation," *arXiv preprint arXiv:2204.05088*, 2022.
- [26] Y. Zhang, Z. Zhu, W. Zheng, J. Huang, G. Huang, J. Zhou, and J. Lu, "Beverse: Unified perception and prediction in birds-eye-view for vision-centric autonomous driving," *arXiv preprint arXiv:2205.09743*, 2022.
- [27] C. Reading, A. Harakeh, J. Chae, and S. L. Waslander, "Categorical depth distribution network for monocular 3d object detection," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2021, pp. 8555–8564.
- [28] J. Phillion and S. Fidler, "Lift, splat, shoot: Encoding images from arbitrary camera rigs by implicitly unprojecting to 3d," in *European Conference on Computer Vision*. Springer, 2020, pp. 194–210.
- [29] Y. Li, Z. Ge, G. Yu, J. Yang, Z. Wang, Y. Shi, J. Sun, and Z. Li, "Bevdepth: Acquisition of reliable depth for multi-view 3d object detection," *arXiv preprint arXiv:2206.10092*, 2022.
- [30] T. Roddick and R. Cipolla, "Predicting semantic map representations from images using pyramid occupancy networks," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2020, pp. 11 138–11 147.
- [31] Z. Li, W. Wang, H. Li, E. Xie, C. Sima, T. Lu, Q. Yu, and J. Dai, "Bevformer: Learning bird's-eye-view representation from multi-camera images via spatiotemporal transformers," *arXiv preprint arXiv:2203.17270*, 2022.
- [32] L. Peng, Z. Chen, Z. Fu, P. Liang, and E. Cheng, "Bevsegformer: Bird's eye view semantic segmentation from arbitrary camera rigs," *arXiv preprint arXiv:2203.04050*, 2022.
- [33] J. Dai, H. Qi, Y. Xiong, Y. Li, G. Zhang, H. Hu, and Y. Wei, "Deformable convolutional networks," in *Proceedings of the IEEE International Conference on Computer Vision (ICCV)*, Oct 2017.
- [34] J. Hu, L. Shen, and G. Sun, "Squeeze-and-excitation networks," in *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2018, pp. 7132–7141.
- [35] Y. Yan, Y. Mao, and B. Li, "Second: Sparsely embedded convolutional detection," *Sensors*, vol. 18, no. 10, p. 3337, 2018.
- [36] A. Geiger, P. Lenz, C. Stiller, and R. Urtasun, "Vision meets robotics: The kitti dataset," *International Journal of Robotics Research (IJRR)*, 2013.
- [37] K. He, X. Zhang, S. Ren, and J. Sun, "Deep residual learning for image recognition," in *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2016, pp. 770–778.
- [38] T.-Y. Lin, P. Dollár, R. Girshick, K. He, B. Hariharan, and S. Belongie, "Feature pyramid networks for object detection," in *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2017, pp. 2117–2125.
- [39] A. H. Lang, S. Vora, H. Caesar, L. Zhou, J. Yang, and O. Beijbom, "Pointpillars: Fast encoders for object detection from point clouds," in *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 2019, pp. 12 697–12 705.
- [40] H. Caesar, V. Bankiti, A. H. Lang, S. Vora, V. E. Liong, Q. Xu, A. Krishnan, Y. Pan, G. Baldan, and O. Beijbom, "nuscenes: A multimodal dataset for autonomous driving," *arXiv preprint arXiv:1903.11027*, 2019.
- [41] Y. Hu, Y. Lu, R. Xu, W. Xie, S. Chen, and Y. Wang, "Collaboration helps camera overtake lidar in 3d detection," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2023, pp. 9243–9252.
- [42] Y.-C. Liu, J. Tian, N. Glaser, and Z. Kira, "When2com: Multi-agent perception via communication graph grouping," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2020.