
Supplementary Materials: Leveraging Vision-Centric Multi-Modal Expertise for 3D Object Detection

Anonymous Author(s)

Affiliation

Address

email

1 In this supplementary document, we present a comprehensive account of the implementation and
2 training details in Section A. We delve into the analysis of misalignment resulting from temporal
3 fusion and discuss the effectiveness of our proposed methods for addressing this issue in Section B.
4 Moreover, we provide additional visualizations in Section C. Lastly, we explore the potential social
5 impact of our research in Section D.

6 A Experiment Details

7 A.1 Dataset and Evaluation Metrics

8 We conduct our experiments on the nuScenes dataset [1], a widely used benchmark for autonomous
9 driving tasks. The dataset encompasses diverse driving scenarios captured using cameras and LiDAR
10 sensors, offering rich information for both visual and LiDAR-based 3D object detection. The dataset
11 comprises 700 training scenes, 150 validation scenes, and 150 testing scenes. Each scene spans
12 approximately 20 seconds, with key frames annotated at a 2 Hz frequency.

13 The two dominant metrics for the nuScenes detection task are the nuScenes Detection Score (NDS)
14 and mean Average Precision (mAP). The mAP for nuScenes is computed based on the center distance
15 between predictions and ground truth annotations on the ground plane. Moreover, the nuScenes dataset
16 defines five true positive metrics (mATE, mASE, mAOE, mAVE, mAAE) for measuring translation,
17 scale, orientation, velocity, and attribute, respectively. The NDS for nuScenes is a weighted sum of
18 mAP and the five true positive metrics, defined as $NDS = \frac{1}{10}[5mAP + \sum_{mTP}(1 - \min(1, mTP))]$.

19 A.2 Implementation Details

20 We conduct experiments on BEVDepth [4]. The codebase is developed upon MMDetection3D [2].
21 Main experiments are trained on 8 NVIDIA A100 GPUs, while ablation experiments are conducted
22 on 8 NVIDIA V100 GPUS. For BEVDepth, the model is trained for 20 epochs with an initial learning
23 rate of $2e-4$. In the distillation process, the per-GPU batch size is set to 4, whereas during the training
24 of the baseline model, it is set to 8. Normal data augmentations are introduced in the training process
25 such as flip and rotate. In our apprentice models, future frames are not incorporated into the long-term
26 temporal fusion throughout the training phase to ensure a fair comparison.

27 In our research, we implement distinct temporal modeling strategies for both apprentice and expert
28 models. For the apprentice models, we incorporate a sequence of eight frames into the temporal
29 modeling process. In contrast, the expert models integrate four future frames into the temporal
30 modeling as demonstrated in our primary results. However, in our ablation study, we deviate from
31 this approach and instead employ eight historical frames for temporal modeling.

Table 1: Experiment settings. * denotes that the training schedule for VCD-E is approximately one-fourth of the original schedule. This reduction was implemented to expedite the training process during the ablation study. The first group is engaged in training on the main results, whereas the second group is utilized in the ablation study.

Method	Backbone	Image Size	mAP (%)	NDS (%)
VCD-E	ConvNext-B [5]	512 x 1408	67.7	71.1
VCD-A	Res-50 [3]	256 x 704	41.8	54.2
VCD-E*	ConvNext-B [5]	256 x 704	54.2	58.8
VCD-A	Res-50 [3]	256 x 704	29.7	40.9

32 A.3 Experiments Settings

33 The setting of adopted expert-apprentice pairs is depicted in Tab. 1. We categorize the distillation
 34 setting into two distinct groups. The primary group is engaged in training on the main results, whereas
 35 the second group is utilized for the ablation study.

36 B The Analysis of Temporal Fusion

37 B.1 The Misalignment of Motion Objects

38 As highlighted in preceding studies [6], long-term temporal fusion may face misalignment issues
 39 in motion estimation, which can be discerned through a reduction in performance on metrics like
 40 mATE. Let’s consider a moving object and analyze the impact of inaccurate motion estimation on its
 41 position in the fused frame. We will assume that the environment is static, except for the moving
 42 object. Let the position of the moving object in the world coordinate system be represented by
 43 $\mathbf{P}_i^w = (x_i^w, y_i^w, z_i^w, 1)^T$ in each of the N frames captured at times t_1, t_2, \dots, t_N . The actual motion
 44 of the moving object between frames is represented by \mathbf{M}_i^{obj} , and the estimated motion is represented
 45 by $\hat{\mathbf{M}}_i^{obj}$. The difference between the estimated and actual motion of the object can be denoted as:

$$\Delta \mathbf{M}_i^{obj} = \mathbf{M}_i^{obj} - \hat{\mathbf{M}}_i^{obj}. \quad (1)$$

46 As we have already computed the transformation matrix \mathbf{T}_i based on the estimated ego motion, we
 47 can calculate the transformed object position in the current frame, considering its actual motion, as:

$$\mathbf{P}_i^{w'} = \mathbf{T}_i \mathbf{M}_i^{obj} \mathbf{P}_i^w. \quad (2)$$

48 The error in the transformed object position can be computed as:

$$\mathbf{e}_i^{obj} = \mathbf{P}_i^{w'} - \hat{\mathbf{P}}_i^{w'}. \quad (3)$$

49 In the long-term fusion process, we integrate the information from all N frames. Assuming we use a
 50 fusion function F , the fused position in the current frame can be represented as:

$$\mathbf{P}_{fusion}^{obj} = F(\mathbf{P}_1^{w'}, \mathbf{P}_2^{w'} \dots, \mathbf{P}_N^{w'}). \quad (4)$$

51 The inaccuracies in the motion estimation of the moving object for each frame can propagate through
 52 the fusion function and result in a misaligned object in the fused frame. The overall error in the fused
 53 position can be represented as a function of the errors in each frame:

$$\mathbf{e}_{fusion}^{obj} = G(\mathbf{e}_1^{obj}, \mathbf{e}_2^{obj} \dots, \mathbf{e}_N^{obj}), \quad (5)$$

54 where G represents a function that combines the errors from each frame. The fused position of
 55 the moving object will be less accurate due to these motion estimation errors, leading to a decline
 56 in object detection performance in the long-term setting. To address the issue mentioned earlier,
 57 we introduce the trajectory-based distillation module, which compensates for the misalignment of
 58 moving objects. We will provide further details in the subsequent discussion.

Table 2: The performance gains of different trajectory length for trajectory-based distillation. As the trajectory length increases, the benefits derived from the distillation process become more pronounced.

Trajectory Length	Distill	mAP (%)	NDS (%)
-	✗	29.7	40.9
0	✓	31.8	42.1
1	✓	33.1	44.5
3	✓	34.6	45.6
5	✓	35.4	45.9
9	✓	33.9	44.7

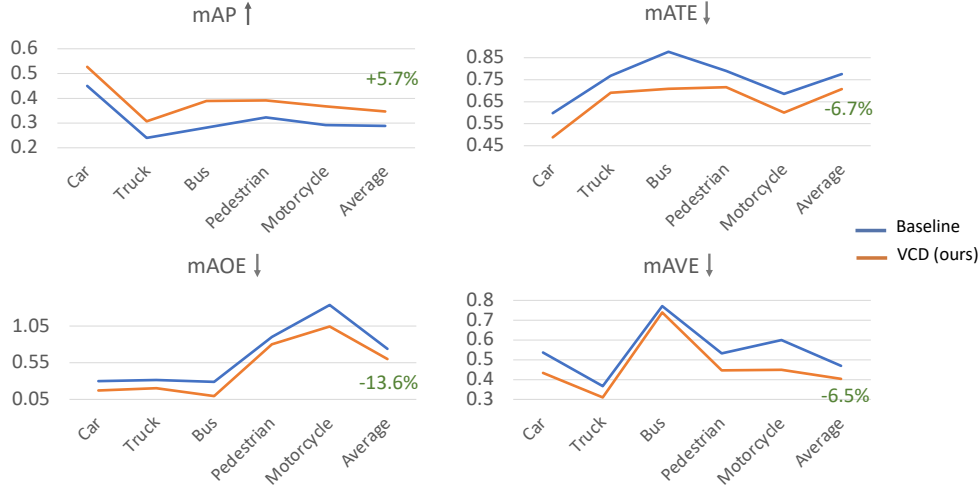


Figure 1: Effects of VCD on movable objects. Our distillation framework VCD consistently improves dynamic objects across a range of metrics.

59 B.2 The Effectiveness of Trajectory-based Distillation

60 The results presented in Table 2 indicate that as the trajectory length increases, the benefits derived
 61 from the distillation process become more pronounced. The temporal fusion length for this experiment
 62 is set at eight. However, when the trajectory length exceeds five, there is a noticeable decrease in
 63 accuracy. We hypothesize that this decrease may be attributed to the model’s distracted attention
 64 towards distant motions. The density of traffic can lead to distant motion locations being occupied
 65 by other objects, which may not necessarily require additional trajectory supervision. This suggests
 66 that the application of excessive trajectory supervision in such scenarios could be unnecessary and
 67 inefficient.

68 B.3 The Improvements of Dynamic Objects

69 In this section, we present visualizations to demonstrate the improvements achieved in dynamic
 70 objects. Particularly noteworthy is the significant enhancement in the representation of dynamic
 71 objects through trajectory-based distillation, thereby highlighting the effectiveness of the trajectory-
 72 based module. As depicted in Fig. 1, our distillation framework consistently enhances dynamic
 73 objects across various metrics.

74 C Visualization

75 We have performed several visualizations in Fig. 2 to showcase the advancements achieved by our
 76 distillation framework. Our findings indicate that our models excel in accurately predicting 3D
 77 bounding boxes for the target objects.




Figure 2: Visualization of the predictions for 3D object detection generated by the VCD-A.

78 D Broader Impact

79 Our research introduces a novel perspective for multi-modal methodologies and a fresh distillation
80 paradigm for camera-only techniques. We believe that it can establish a robust baseline for the broader
81 scientific community. However, while our methods contribute to the enhancement of autonomous
82 driving, they are not yet capable of addressing more complex corner cases. Consequently, these
83 limitations could potentially introduce risks in real-world autonomous systems.

84 References

- 85 [1] Holger Caesar, Varun Bankiti, Alex H Lang, Sourabh Vora, Venice Erin Liong, Qiang Xu, Anush
86 Krishnan, Yu Pan, Giancarlo Baldan, and Oscar Beijbom. nuScenes: A multimodal dataset for
87 autonomous driving. In *Proc. IEEE Conf. on Computer Vision and Pattern Recognition (CVPR)*,
88 2020. 1
- 89 [2] MMDetection3D Contributors. MMDetection3D: OpenMMLab next-generation platform for
90 general 3D object detection. <https://github.com/open-mmlab/mmdetection3d>, 2020. 1
- 91 [3] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image
92 recognition. In *Proc. IEEE Conf. on Computer Vision and Pattern Recognition (CVPR)*, 2016. 2
- 93 [4] Yinhao Li, Zheng Ge, Guanyi Yu, Jinrong Yang, Zengran Wang, Yukang Shi, Jianjian Sun, and
94 Zeming Li. BEVDepth: Acquisition of reliable depth for multi-view 3d object detection. *arXiv*
95 *preprint arXiv:2206.10092*, 2022. 1
- 96 [5] Zhuang Liu, Hanzi Mao, Chao-Yuan Wu, Christoph Feichtenhofer, Trevor Darrell, and Saining
97 Xie. A convnet for the 2020s. In *Proc. IEEE Conf. on Computer Vision and Pattern Recognition*
98 *(CVPR)*, pages 11976–11986, 2022. 2
- 99 [6] Jinyung Park, Chenfeng Xu, Shijia Yang, Kurt Keutzer, Kris Kitani, Masayoshi Tomizuka,
100 and Wei Zhan. Time will tell: New outlooks and a baseline for temporal multi-view 3d object
101 detection. *arXiv preprint arXiv:2210.02443*, 2022. 2