

Exploring Recurrent Long-term Temporal Fusion for Multi-view 3D Perception

Chunrui Han^{1*}, Jinrong Yang^{2*}, Jianjian Sun¹, Zheng Ge¹, Runpei Dong³
 Hongyu Zhou¹, Weixin Mao⁴, Yuang Peng⁵, Xiaoping Li², Xiangyu Zhang¹

Abstract—Long-term temporal fusion is a crucial but often overlooked technique in camera-based Bird’s-Eye-View (BEV) 3D perception. Existing methods are mostly in a parallel manner. While parallel fusion can benefit from long-term information, it suffers from increasing computational and memory overheads as the fusion window size grows. Alternatively, BEVFormer adopts a recurrent fusion pipeline so that history information can be efficiently integrated, yet it fails to benefit from longer temporal frames. In this paper, we explore an embarrassingly simple long-term recurrent fusion strategy built upon the LSS-based methods and find it already able to enjoy the merits from both sides, i.e., rich long-term information and efficient fusion pipeline. A temporal embedding module is further proposed to improve the model’s robustness against occasionally missed frames in practical scenarios. We name this simple but effective fusing pipeline VideoBEV. Experimental results on the nuScenes benchmark show that VideoBEV obtains strong performance on various camera-based 3D perception tasks, including object detection (55.4% mAP and 62.9% NDS), segmentation (48.6% vehicle mIoU), tracking (54.8% AMOTA), and motion prediction (0.80m minADE and 0.463 EPA).

Index Terms—Multi-view 3D object detection, recurrent network and long-term temporal fusion

I. INTRODUCTION

TEMPORAL fusion technique is crucial to autonomous driving systems and it has drawn growing attention in recent years. Many approaches for temporal feature fusion have been developed, and the existing research in camera-based Bird’s-Eye-View (BEV) 3D perception can be divided into two categories, i.e., *parallel* fusion and *recurrent* fusion.

Parallel fusion, popularized by [1], [2], [3], first aligns all history features within a fixed-length window to the current frame and then fuses them, see Fig. 1(a). This paradigm is conceptually simple but effective. A recent work [4] further showcases that *parallel* fusion benefits from increasing the history frame number up to 16. This covers around 8 seconds of temporal information on the nuScenes [5] benchmark, making *parallel* fusion the dominant method in this field. However, these advantages come at the cost of several issues. Firstly, *parallel* fusion typically requires a *fixed* window size [6], which impedes the utilization of longer history frames, but real-world driving usually involves long distances. Secondly,

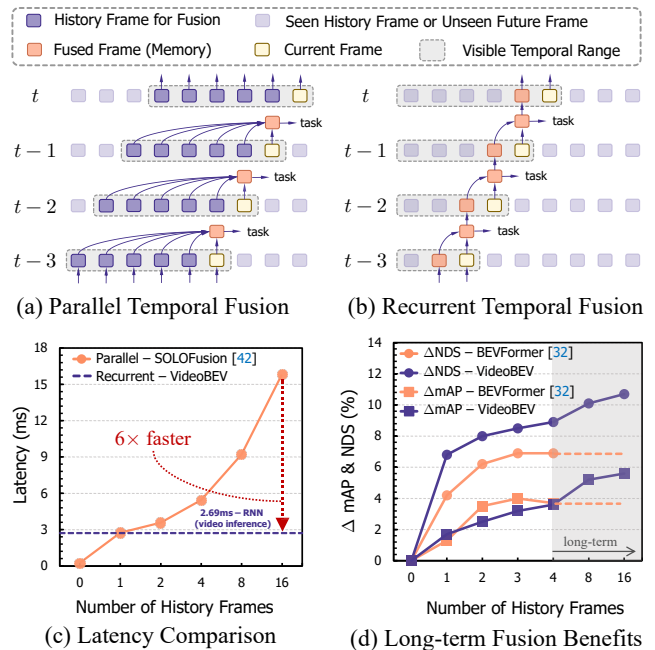


Fig. 1. **Conceptual comparison of two mainstream temporal feature fusion mechanisms.** (a) *Parallel* temporal propagation within fixed temporal segments of each time stamp [6], [7], [8], [2], [4], [9], [10]; (b) *Recurrent* temporal fusion with an iteratively updated long-term memory within the video sequence of any length [11], [12], [13], [14]. (c) Efficiency comparison between our recurrent style VideoBEV and parallel style SOLOFusion [4]. (d) Comparison of benefits (Δ mAP \uparrow and Δ NDS \uparrow) from long-term fusion between earlier recurrent style BEVFormer [14] and our VideoBEV, the numbers of BEVFormer are taken from [14].

this paradigm usually leads to a larger *computation budget* compared to the *recurrent* manner. As shown in Fig. 1(c), SOLOFusion [4] suffers from the growing latency when increasing the number of history frames. These issues hinder the application of the *parallel* temporal fusion technique.

Compared to *parallel* fusion, *recurrent* fusion is more feasible for *longer* history frames since it encodes all history information into a single memory feature (i.e., Fused Frame in Fig. 1(b)). However, the pioneering method BEVFormer [14] shows that *recurrent* feature fusion cannot benefit from longer history frames. See Fig. 1(d), both mAP and NDS stop improving when the number of history frames is more than 3. The reasons could be two-fold: (i) the temporal fusion is *intertwined* with the view transformation process of the current frame, making it more difficult to fuse temporal information, (ii) the spatial-temporal fusion network in BEVFormer is a Transformer [15] architecture that is *deep* and *cumbersome*, which may consequently lead to the typical *gradient vanishing*

*Equal Contribution

¹Chunrui Han, Jianjian Sun, Zheng Ge, Hongyu Zhou, Xiangyu Zhang are with the Megvii Technology. ²Jinrong Yang, Xiaoping Li are with the Huazhong University of Science and Technology. ³Runpei Dong is with the Xi’an Jiaotong University. ⁴Weixin Mao is with the Waseda University. ⁵Yuang Peng is with the Tsinghua University. Corresponding author: Chunrui Han. (e-mail: chunrui.han@vip1.ict.ac.cn)

issue in RNN when the sequence length is long [11], [16]. They [10] consequently turn back into the *parallel* manner. As a result, in the multi-view 3D perception field, *none of the existing methods can simultaneously enjoy an efficient fusing pipeline and the benefits carried by long-term information.*

Is it not feasible to apply efficient long-term temporal fusion to multi-view 3D perception tasks? The answer is *no*. By leveraging a *decoupled* view transformation and temporal fusion procedures on LSS-based detectors [17], [1], [18], we find an embarrassing fact that a simple temporal fusion strategy can facilitate our goal. During training, we sample BEV features within a sufficiently long window (*e.g.*, 16 frames) and fuse them sequentially. During inference, the sequential fusion mechanism is retained throughout the entire driving process with the sampling window strategy discarded. This methodology is similar to BEVFormer [14], despite that we sample more frames during training and use a framework with *decoupled* spatiotemporal fusion. As a result, we obtain a simple but effective multi-frame BEV framework, dubbed VideoBEV, which can be applied to diverse perception tasks in autonomous driving. To ensure stable and robust 3D motion perception when facing occasionally missed frames in real-world scenarios, we propose a temporal embedding module to encode timestamps, with which the dynamic temporal interval information can be effectively modeled.

Extensive experiments are conducted on four 3D perception tasks, including 3D object detection, map segmentation, object tracking, and object motion prediction. For example, on the nuScenes benchmark, VideoBEV achieves **55.4%** mAP and **62.9%** NDS on the 3D detection task, which improves **+2.9%** mAP and **+1.9%** NDS over the single-frame baseline. On the 3D object tracking benchmark that models object motion states, VideoBEV achieves **54.8%** AMOTA, significantly outperforming the single-frame baseline by **+6.8%**. While obtaining strong performance on various tasks, VideoBEV is still far more efficient than its long-term counterpart SOLO-Fusion [4]. Although extremely simple, our VideoBEV is the first method that demonstrates the benefit of continuously increasing the number of history frames. Besides, our VideoBEV has provided an in-depth understanding of the significance of efficient long-term temporal fusion while also establishing a new baseline for spatiotemporal multi-view 3D perception.

II. RELATED WORKS

A. Camera-Based Single-Frame 3D Perception

The majority of camera-based single-frame 3D prediction techniques in the beginning simply predicted 3D boxes from images. By creating a 3D box with the anticipated properties of a 3D object using a 2D box, Mousavian *et al.* [19] pioneered this direction. FCOS3D [20] simply extends the 2D object detector [21] to a 3D object detector by decoupling the defined 7-DoF 3D targets as 2D and 3D attributes. PETR [22] encodes the position information of 3D coordinates into image features, producing the 3D position-aware feature. Inspired by LiDAR-based methods [23], [24], recent advances employ view transformation to transform the feature from perspective view to the Bird’s-Eye-View (BEV) for unified 3D detection.

LSS [17] proposes the LSS-based view transformation method, which first “lift”s each image individually into a frustum of feature for each camera, then “splat”s all frustums into a rasterized BEV grid. BEVDet [25] utilizes the LSS-based view transformation to extract BEV features and conducts 3D detection thereon. To achieve more trustworthy depth for LSS-based view transformation, BEVDepth [1] uses the depth from LiDAR as the supervision for precise depth estimation.

B. Camera-Based Multi-Frame 3D Perception

Multi-frame fusion for LiDAR-based 3D detectors is a widely used technology [26], [27], [28]. However, 3D perception from a single vision frame without LiDAR is an ill-posed problem due to the lack of accurate depth information. Recent works make efforts to multi-frame 3D perception since different frames generally offer different views of objects. Saha *et al.* [29] formulate BEV map construction from an image as a set of 1D sequence-to-sequence translations and propose a dynamic module incorporating temporal information from past estimation to build a spatiotemporal BEV representation. BEVDet4D [2] extends BEVDet [25] and fuses the history frame’s features with the current frame after removing ego-motion impact. PETRv2 [3] directly achieves the temporal alignment by simply aligning the 3D coordinates of the history and current frames. BEVFormer [14] designs a temporal self-attention to recurrently fuse the history BEV information for obtaining precise BEV features. BEVStereo [18] employs the temporal multi-view stereo (MVS) [30] to tackle the ill-posed issue of depth perception in camera-based 3D tasks. STS [31] leverages the geometry correspondence between frames across time to facilitate accurate depth learning. The above methods employ only limited history frames for temporal fusion. Differently, SOLOFusion [4] aligns the BEV feature from the previous timesteps of a long history to the current timestep and concatenates them for long-term temporal fusion. However, it suffers from high inference latency, memory, and module parameter bottleneck. Our proposed recurrent temporal fusion module can avoid these issues. Besides, for the first time, our VideoBEV successfully demonstrates the benefit of continuously increasing the number of history frames.

III. METHODOLOGY

VideoBEV employs a recurrent long-term fusion module that fuses a video stream sequentially. A temporal embedding module is further introduced to tackle the instability of perception caused by missed frames in real-world circumstances. The overall architecture is shown in Fig. 2. Sec. III-A first gives a brief overview of VideoBEV, then Sec. III-B introduces the recurrent style BEV fusion in detail. In the end, Sec. III-C introduces the temporal embedding modeling.

A. Overview of VideoBEV

The overall pipeline of VideoBEV is similar to that of existing LSS-based [17] 3D BEV detection, *e.g.*, BEVDet [25], BEVDepth [1] and BEVStereo [18], *etc.*, except the recurrent style temporal fusion and temporal embedding. Generally, it can be separated into three modules:

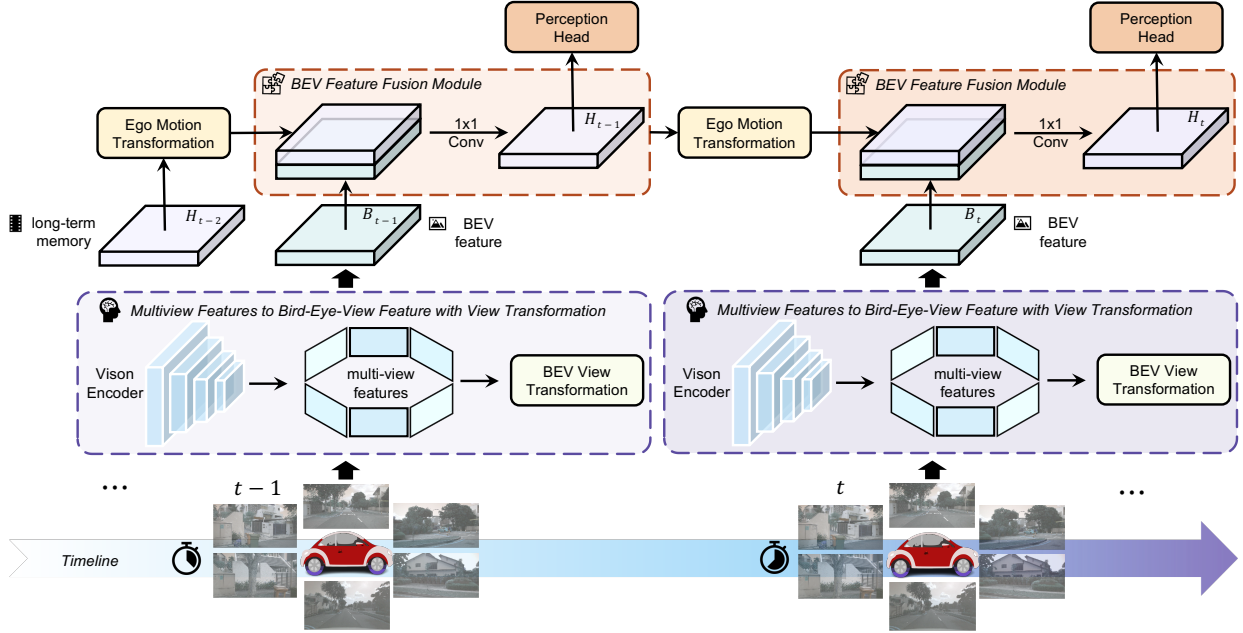


Fig. 2. **Overview of VideoBEV.** The backbone first extracts image features of different views of a frame, which are transformed to BEV from the image view to obtain the BEV feature. Then, the recurrent fusion module fuses the new BEV feature with the one of long-term memory, based on which the memory is updated and the 3D perception tasks are conducted.

- 1) BEV feature extraction module: a backbone network extracts the per-frame image feature of different camera views, which is further translated from perspective view to BEV for obtaining the BEV feature.
- 2) Temporal fusion module: the recurrent style temporal fusion module fuses the BEV feature of the input frame with the stored long-term memory. Besides, a recurrent style temporal embedding module is employed to embed the sequence of time intervals between adjacent frames in the video sequence.
- 3) 3D perception module: a 3D perception head is applied to the fused BEV feature and temporal embedding to conduct 3D perception for the input frame.

In the BEV feature extraction module, the backbone can be any network, *e.g.*, ResNet-50 [32], ConvNeXt-B [33]; the view transformation (VT) can be generally LSS-based VT such as BEVDet [25], BEVDepth [1], and MatrixVT [34], *etc.*, or query-based VT such as BEVFormer [14]. We utilize the LSS-based VT due to its effectiveness. In the 3D perception module, the head can be any BEV-based task, *e.g.*, 3D objection detection, map segmentation, and tracking, *etc.* The temporal fusion module is newly proposed in this paper and will be introduced in the following subsections.

B. Recurrent Temporal BEV Feature Fusion

The BEV feature of a single frame generally describes objects from a single view (time step), which is inadequate for precise 3D perception. To obtain abundant features of objects, recent works, *e.g.*, SOLOFusion [4], explore the temporal context information as the substitute for multi-views since different frames often offer different views of subjects. However, as pointed out by BEVFormer [14] and BEVFormer V2 [10], existing *recurrent* style fusion fails to bring further performance gains with long-term sequence. In

contrast, the *parallel* temporal fusion is able to fuse long-term video sequences effectively. Hence, we motivate our long-term recurrent style temporal fusion model from that of the sliding-window methods, introduced next.

To better understand the recurrent style fusion, we first revisit the parallel fusion with a temporal window size k in SOLOFusion [4]. Suppose the BEV feature of a video sequence as $\{B_i\}_{i=1}^T$, and B_i is the BEV feature of a frame at time step t_i . The parallel fusion for the i -th frame in SOLOFusion [4] can be written as:

$$\hat{H}_i = \left[f_{\text{sample}}(B_{i-k+1}, P_{i,i-k+1}); \dots; f_{\text{sample}}(B_{i-1}, P_{i,i-1}); B_i \right] * \mathbf{U}, \quad (1)$$

where \hat{H}_i is the fused BEV feature for i -th frame, $P_{j,i}$ is the view transformation matrix from the ego coordinate of j -th frame to that of i -th frame considering the ego-motion, f_{sample} refers to the grid sampling operation proposed by Jaderberg *et al.* [35], \mathbf{U} is the convolution kernel, $[x;y]$ represents the concatenation of x and y along the channel dimension, and $*$ denotes the convolution operator. As can be seen, in the parallel temporal fusion, a concatenation operator is applied first to concatenate the aligned BEV feature in the window, on which the convolution operator is employed to fuse BEV features of different frames. The above formulation can be further expanded by splitting the kernel \mathbf{U} along the channel dimension, *i.e.*,

$$\begin{aligned} \hat{H}_i &= \left[f_{\text{sample}}(B_{i-k+1}, P_{i,i-k+1}); \dots; f_{\text{sample}}(B_{i-1}, P_{i,i-1}); B_i \right] * \\ &\quad \left[\mathbf{U}_1; \dots; \mathbf{U}_{k-1}; \mathbf{U}_k \right] \\ &= \sum_{j=1}^k f_{\text{sample}}(B_{i-k+j}, P_{i,i-k+j}) * \mathbf{U}_j, \end{aligned} \quad (2)$$

where \mathbf{U}_j is the j -th chunk by equally splitting \mathbf{U} along the channel dimension.

The formulation of recurrent fusion is similar to that of parallel fusion. The difference is that instead of storing all the

history BEV features in the temporal window and concatenating them, we store only the long-term memory of BEV feature and concatenate it with that of the current frame. Taking \bar{H}_i as long-term memory of BEV feature at time step t_i , the formulation of recurrent style fusion is:

$$\bar{H}_i = \left[f_{\text{sample}}(\bar{H}_{i-1}, P_{i,i-1}); B_i \right] * \mathbf{V}, \quad (3)$$

where \mathbf{V} is the convolution kernels. Considering the long-term BEV feature memory \bar{H}_{i-1} is obtained by fusing the \bar{H}_{i-2} to the BEV feature of $(i-1)$ -th frame, the above formulation can be further expanded by splitting the kernel \mathbf{V} into two chunks \mathbf{V}_{mem} and \mathbf{V}_{cur} along the channel dimension, respectively convolving the long-term memory and the current BEV features:

$$\begin{aligned} \bar{H}_i &= \left[f_{\text{sample}}(\bar{H}_{i-1}, P_{i,i-1}); B_i \right] * \left[\mathbf{V}_{\text{mem}}; \mathbf{V}_{\text{cur}} \right] & (4) \\ &= f_{\text{sample}}(\bar{H}_{i-1}, P_{i,i-1}) * \mathbf{V}_{\text{mem}} + B_i * \mathbf{V}_{\text{cur}} \\ &= f_{\text{sample}}(\bar{H}_{i-2}, P_{i,i-2}) * \mathbf{V}_{\text{mem}} * \mathbf{V}_{\text{mem}} + \\ &\quad f_{\text{sample}}(B_{i-1}, P_{i,i-1}) * \mathbf{V}_{\text{cur}} * \mathbf{V}_{\text{mem}} + B_i * \mathbf{V}_{\text{cur}} \\ &\triangleq f_{\text{sample}}(\bar{H}_{i-2}, P_{i,i-2}) * \mathbf{V}_{\text{mem}}^2 + f_{\text{sample}}(B_{i-1}, P_{i,i-1}) * \\ &\quad \mathbf{V}_{\text{cur}} * \mathbf{V}_{\text{mem}} + B_i * \mathbf{V}_{\text{cur}} \\ &= \sum_{j=1}^i f_{\text{sample}}(B_j, P_{i,j}) * \mathbf{V}_{\text{cur}} * \mathbf{V}_{\text{mem}}^{i-j}. & (5) \end{aligned}$$

Here, $B * \mathbf{V}^n$ denotes convolution of B with convolution kernel \mathbf{V} repeating n times.

Comparing Eq. 4 to Eq. 2, it is seen that the two fusion styles have similar formulations of summing the convolved BEV feature of history frames. This may be the reason both the recurrent paradigm VideoBEV and the parallel paradigm SOLOFusion [4] can benefit from the long-term temporal information. However, from the final derivation, we can see that the convolution kernel \mathbf{V}_{mem} for the i -th frame in recurrent style fusion is computed repeatedly $i-j$ times (*i.e.*, the fusion time interval) when fusing with the j -th history frame. Thus, the recurrent style fusion is aware of the time interval for every history frame. In contrast, the sliding window fusion kernel \mathbf{U}_j for fusing the j -th history frame is computed once for all $j \in \{1, \dots, i\}$. As a result, it treats every history frame equally without the recurrent style explicit time interval modeling. Besides, the sliding window style fusion only fuses the history $k-1$ frames to current frames, while the recurrent style fusion fuses all the history frames, facilitating better 3D perception.

C. Temporal Embedding

Generally, the time interval between two adjacent frames in a video sequence is fixed, *e.g.*, 0.5s between two key-frames on nuScenes. However, this can not be guaranteed in the complex real scenes. We empirically find that the missed frames can dreadfully hurt motion estimation. As shown in Fig. 3, the average error of predicted velocity (AVE) becomes dramatically high when the frame missing rate is high. Thus, besides the BEV feature fusion, we propose a temporal embedding module to fuse the time interval between two adjacent frames for stable 3D perception. The temporal embedding module is also designed in a recurrent fashion.

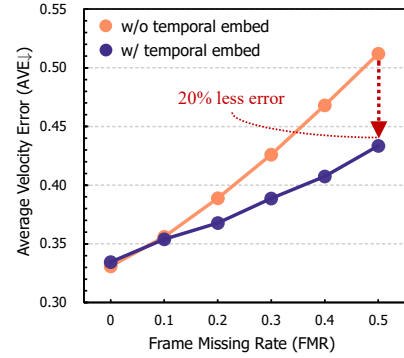


Fig. 3. **Average velocity error (AVE \downarrow) versus frame missing rate (FMR).** Without the proposed temporal embedding, the AVE is dramatically high when frames are missed, and this issue is substantially mitigated when using the proposed temporal embedding.

Taking $\Delta t_i = t_i - t_{i-1}$ as the time interval between the i -th frame and $i-1$ -th frame, the formulation of temporal embedding \bar{E}_i for i -th frame is as follows:

$$E_i = e(\Delta t_i \cdot \mathbf{1}), \quad (6)$$

$$\bar{E}_i = \left[\bar{E}_{i-1}; E_i \right] * K. \quad (7)$$

Here, $\mathbf{1}$ is the all-one matrix with the same spatial size as \bar{H}_i , $e(\cdot)$ is the temporal embedding function, which consists of two convolutional layers. K is the convolution kernel for the recurrent fusion. The fused temporal embedding is fed into the velocity head for robust velocity prediction.

D. Video Inference

During inference, each frame in the video sequence is evaluated chronologically. The long-term BEV memory feature is initialized with zero. When a new frame comes, the BEV feature is first fused with the memory, based on which the memory is updated and the 3D perceptron is conducted. As a result, the overhead of VideoBEV is consistently low with longer video inputs (see Fig. 5).

IV. EXPERIMENTS

A. Experimental Setting

a) Dataset: We use nuScenes dataset [5] for experimental evaluations. It contains 1,000 autonomous driving scenes with around 20 seconds per scene, which is split into 850 scenes for training (`train`) or validation (`val`) and 150 for testing (`test`). Six camera images from different perspectives are provided in each frame of the camera data.

b) Evaluation Metric: We use four commonly used tasks for autonomous driving systems, as stated below. We use the typical evaluation criteria for **3D object detection** and report the mean Average Precision (mAP) and nuScenes detection score (NDS). The 3D attributes of translation, scale, orientation, velocity, and attribute are evaluated using the 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), respectively. The Mean Intersection over Union (mIoU) of the drivable area, the lane, and the vehicle is reported following LSS [17] for the purpose of **map segmentation** evaluation. For object **tracking** evaluation, we

TABLE I
COMPARISON RESULTS ON 3D DETECTION ON THE NUSCENES **VAL** SET. ALL METHODS IN THE TABLE ARE TRAINED WITH CBGS. #FRAMES DENOTES THE NUMBER OF FRAMES USED DURING TRAINING.

Method	Backbone	Image Size	#Frames	mAP \uparrow	NDS \uparrow	mATE \downarrow	mASE \downarrow	mAOE \downarrow	mAVE \downarrow	mAAE \downarrow
BEVDet [25]	ResNet50	256 \times 704	1	0.298	0.379	0.725	0.279	0.589	0.860	0.245
PETR [22]	ResNet50	384 \times 1056	1	0.313	0.381	0.768	0.278	0.564	0.923	0.225
BEVDet4D [2]	ResNet50	256 \times 704	2	0.322	0.457	0.703	0.278	0.495	0.354	0.206
BEVDepth [1]	ResNet50	256 \times 704	2	0.351	0.475	0.639	0.267	0.479	0.428	0.198
STS [31]	ResNet50	256 \times 704	2	0.377	0.489	0.601	0.275	0.450	0.446	0.212
BEVStereo [18]	ResNet50	256 \times 704	2	0.372	0.500	0.598	0.270	0.438	0.367	0.190
AeDet [36]	ResNet50	256 \times 704	2	0.387	0.501	0.598	0.276	0.461	0.392	0.196
SOLOFusion [4]	ResNet50	256 \times 704	17	0.427	0.534	0.567	0.274	0.511	0.252	0.188
VideoBEV	ResNet50	256 \times 704	8	0.422	0.535	0.564	0.276	0.440	0.286	0.198

TABLE II
COMPARISON RESULTS ON 3D DETECTION ON THE NUSCENES **TEST** SET. TTA DENOTES TEST TIME AUGMENTATION STRATEGY. \dagger DENOTES RESULTS USING FUTURE FRAMES DURING TRAINING AND INFERENCE, AND \ddagger DENOTES RESULTS FROM THE OFFICIAL NUSCENES LEADERBOARD.

Method	Backbone	Image Size	TTA	mAP \uparrow	NDS \uparrow	mATE \downarrow	mASE \downarrow	mAOE \downarrow	mAVE \downarrow	mAAE \downarrow
FCOS3D [37]	R101-DCN	900 \times 1600	✓	0.358	0.428	0.690	0.249	0.452	1.434	0.124
DETR3D [38]	V2-99	900 \times 1600	✓	0.412	0.479	0.641	0.255	0.394	0.845	0.133
UVTR [39]	V2-99	900 \times 1600	✗	0.472	0.551	0.577	0.253	0.391	0.508	0.123
BEVFormer [14]	V2-99	900 \times 1600	✗	0.481	0.569	0.582	0.256	0.375	0.378	0.126
BEVDet4D [2]	Swin-B	900 \times 1600	✓	0.451	0.569	0.511	0.241	0.386	0.301	0.121
PolarFormer [40]	V2-99	900 \times 1600	✗	0.493	0.572	0.556	0.256	0.364	0.439	0.127
PETRv2 [3]	RevCol	640 \times 1600	✗	0.512	0.592	0.547	0.242	0.360	0.367	0.126
HoP-BEVFormer [26]	V2-99	640 \times 1600	✗	0.517	0.603	0.501	0.245	0.346	0.362	0.105
BEVDepth [1]	ConvNeXt-B	640 \times 1600	✗	0.520	0.609	0.445	0.243	0.352	0.347	0.127
BEVStereo [18]	V2-99	640 \times 1600	✗	0.525	0.610	0.431	0.246	0.358	0.357	0.138
AeDet [36]	ConvNeXt-B	640 \times 1600	✓	0.531	0.620	0.439	0.247	0.344	0.292	0.130
SOLOFusion [4]	ConvNeXt-B	640 \times 1600	✗	0.540	0.619	0.453	0.257	0.376	0.276	0.148
VideoBEV	ConvNeXt-B	640 \times 1600	✗	0.554	0.629	0.457	0.249	0.381	0.266	0.132
BEVFormer V2 \dagger [10]	InternImage-B	640 \times 1600	✗	0.540	0.620	0.488	0.251	0.335	0.302	0.122
BEVFormer V2 \dagger [10]	InternImage-XL	640 \times 1600	✗	0.556	0.634	0.456	0.248	0.317	0.293	0.123
BEVFormer V2 Opt \ddagger	InternImage-XL	640 \times 1600	✗	0.580	0.648	0.448	0.262	0.342	0.238	0.128
BEVDet-Gamma \ddagger	Swin-B	640 \times 1600	✓	0.586	0.664	0.375	0.243	0.377	0.174	0.123
HoP-BEVDet4D \dagger [26]	ViT-L	640 \times 1600	✗	0.624	textbf{0.685}	0.367	0.249	0.354	0.171	0.131
VideoBEV \dagger	ConvNeXt-B	640 \times 1600	✗	0.592	0.670	0.385	0.246	0.323	0.174	0.137

report the average multi-object tracking accuracy (AMOTA), the average multi-object tracking precision (AMOTP), the recall (RECALL), and the multi-object tracking accuracy (MOTA) following the standard assessment metrics. For object **motion prediction** evaluation, we report the minimum Average Displacement Error (minADE), minimum Final Displacement Error (minFDE), Miss Rate (MR), and the End-to-end Prediction Accuracy (EPA) following ViP3D [41].

c) Implementation Details: We conduct our experiments based on the BEVDepth [1] and BEVStereo [18]. The learning rate, optimizer, and data augmentation are the same as BEVDepth [1]. Unless otherwise specified, we use ResNet50 [32] pre-trained on ImageNet [42] as the image backbone and SECOND FPN [43] as the image neck. The size of the BEV feature in all of our experiments is 128 \times 128. The perception ranges are [-51.2m, 51.2m] for the X and Y axis, and the resolution of each BEV grid is 0.8m.

B. Comparison to Prior Arts

3D Detection To fairly compare with existing SOTAs, we use the ResNet-50, ResNet-101, and ConvNext-Base as backbone respectively. The main results on Nuscenes **val** and **test** sets are shown in Tab. I and Tab. II. On **val** set, with the ResNet-50 backbone, the VideoBEV achieves comparable results to SOLOFusion [4] with fewer frames for training. On

TABLE III
COMPARISON RESULTS ON MAP SEGMENTATION ON THE NUSCENES **VAL** SET. \dagger DENOTES OUR BASELINE METHOD.

Method	mIoU-Drivable \uparrow	mIoU-Lane \uparrow	mIoU-Vehicle \uparrow
LSS [17]	0.729	0.200	0.321
FIERY [45]	-	-	0.382
M ² BEV [46]	0.759	0.380	-
BEVFormer [14]	0.775	0.239	0.467
UniAD [47]	0.691	0.313	-
BEVDepth \dagger [1]	0.816	0.453	0.460
VideoBEV	0.827	0.461	0.486

the **test** set, our VideoBEV achieves 55.4% mAP and 62.9% NDS without bells and whistles, outperforming all previous methods without the utilization of future frames. Furthermore, our VideoBEV can extend to future frames fusion in the offboard mode, where our method still surpasses most existing methods, including BEVFormer V2 which uses a heavier backbone network (*i.e.*, InternImage-XL [44]). These strong results clearly demonstrate the effectiveness of VideoBEV for fusing long-term temporal information.

Map Segmentation We evaluate VideoBEV on map segmentation task by simply adding a U-Net-like [62] network for the segmentation of the drivable area, the lane, and the vehicle in BEV. As shown in Tab. III, compared to our baseline (single-frame), VideoBEV improves the mIoUs of the three classes by +1.1%, +0.8%, and +2.6%, respectively. VideoBEV surpasses all existing SOTAs, including the BEVFormer [14]

TABLE IV
COMPARISON RESULTS ON 3D OBJECT TRACKING ON THE NUSCENES TEST SET. † DENOTES OUR BASELINE METHOD.

Method	AMOTA↑	AMOTP↓	RECALL↑	MOTA↑
CenterTrack [48]	0.046	1.543	23.3%	0.043
DEFT [49]	0.177	1.564	33.8%	0.156
Time3D [50]	0.210	1.360	-	0.173
QD3DT [51]	0.217	1.550	37.5%	0.198
TripletTrack [52]	0.268	1.504	40.0%	0.245
MUTR3D [53]	0.270	1.494	41.1%	0.245
PolarDETR [54]	0.273	1.185	40.4%	0.238
UniAD [47]	0.359	1.320	46.7%	-
SRCN3D [55]	0.398	1.317	53.8%	0.359
CC-3DT [56]	0.410	1.274	53.8%	0.357
PF-Track [57]	0.434	1.252	53.8%	0.378
QTrack† [58]	0.480	1.107	56.9%	0.431
UVTR [39]	0.519	1.125	59.9%	0.447
Sparse4D [59]	0.519	1.078	63.3%	0.459
VideoBEV	0.548	0.983	63.1%	0.475

TABLE V
COMPARISON TO EXISTING WORK ON PREDICTION ON THE NUSCENES VAL SET. † DENOTES OUR BASELINE METHOD.

Method	minADE (m)↓	minFDE (m)↓	MR↓	EPA ↑
PnPNet-vision [60]	2.22	3.17	0.272	0.193
ViP3D [41]	2.05	2.84	0.246	0.226
PIP [61]	1.23	1.75	0.195	0.258
UniAD [47]	0.71	1.02	0.151	0.456
BEVDepth†[11]	1.19	1.62	0.133	0.386
VideoBEV	0.80	0.99	0.067	0.463

and UniAD [47]. This indicates the temporal information fused by our recurrent fusion module can improve the quality of BEV features for tasks that require dense spatial reasoning.

Object Tracking For 3D multi-object tracking (MOT) task, we employ QTrack [58] as our baseline method to generate the trajectories of all predicted 3D objects. As shown in Tab. IV, VideoBEV achieves the best performance on the nuScenes test set, which outperforms Sparse4D [59] and UVTR [39] by a clear margin of +2.9% AMOTA. Compared to our baseline method QTrack [58], a significant improvement of +6.8% is observed, demonstrating the superiority and consistent ability to identify objects moving over time.

Object Motion Prediction We also evaluate VideoBEV on the motion prediction task. Inspired by FutureDet [63], we first conduct future detection for all target agents in a finite future period (*i.e.*, 6 key-frames in 3s). Then, we simply utilize the velocity and time lag to associate the locations among current and future detection results. Finally, we take the detection confidence score from the last frame as the score of the corresponding associated motion trajectory. As shown in Tab. V, the single-frame BEVDepth baseline using the aforementioned strategy already yields a promising result, outperforming PIP [61] on all metrics. Further, by utilizing our efficient temporal fusion strategy of VideoBEV, the SOTA performance is achieved with only a ResNet-50 backbone and 256×704 input resolution. This demonstrates that our sequential modeling successfully captures the object motion states, which is conducive to future detection for further accurate object motion forecasting.

C. Ablation Study and Analysis

a) *Effectiveness of Recurrent Temporal Fusion:* To verify the effectiveness of recurrent temporal fusion, we use different numbers of history frames with ResNet-50 backbone for

TABLE VI
ABLATION STUDY ON HISTORY FRAMES NUMBER. #FRAMES DENOTES USED HISTORY FRAMES NUMBER FOR TRAINING.

#Frames	mAP↑	NDS↑	mATE↓	mAOE↓	mAVE↓
0	0.323	0.382	0.701	0.598	0.936
1	0.340	0.450	0.678	0.550	0.473
2	0.348	0.462	0.688	0.533	0.397
4	0.359	0.471	0.659	0.556	0.382
8	0.375	0.483	0.663	0.524	0.360
16	0.379	0.489	0.641	0.524	0.343
all	0.379	0.492	0.636	0.519	0.331

TABLE VII
ABLATION STUDY ON COMBINING SHALLOWER-LAYER FUSION. VIDEOBEV-D AND VIDEOBEV-S REPRESENT VIDEOBEV BASED ON BEVDEPTH [1] AND BEVSTEREO [18], RESPECTIVELY.

Method	mAP↑	NDS↑	mATE↓	mAOE↓	mAVE↓
BEVDepth [1]	0.323	0.382	0.701	0.598	0.936
VideoBEV-D	0.379	0.492	0.636	0.519	0.331
BEVStereo [18]	0.340	0.450	0.683	0.533	0.478
VideoBEV-S	0.395	0.502	0.606	0.511	0.344

training and testing. As shown in Fig. 1(d) and Tab. VI, with the increase of used history frames, the mAP and NDS are significantly improved. Specifically, the improvement of VideoBEV with 16 history frames is +5.6% mAP, +10.7% NDS over that without temporal fusion. When using video inference with all history frames, the performance is further improved compared to the 16 history frames counterpart. This demonstrates that though the frames sampled beyond 16-th history stamp are far from the current frame, they still contain temporal information useful to the current frame.

b) *Efficiency of Temporal Recurrent Fusion:* VideoBEV recurrently fuses the history BEV feature. Hence, only one BEV feature memory needs to be stored during inference. When a new frame comes, we only need to fuse its BEV feature to the stored one with a lightweight recurrent fusion module. This is efficient for both memory and computation. As shown in Fig. 5, when increasing the number of used history frames, the overhead of memory and latency is consistently lower compared to SOLOFusion [4], which is nearly the same as that without any history frame.

c) *Robustness to Missed Frames:* In practice, the frames could sometimes miss, resulting in different time intervals. However, we empirically find that varied time intervals between two adjacent frames can seriously hurt the velocity prediction. We use the frame missing rate (FMR) to study the influence of varied time intervals. As shown in Fig. 3, a higher FMR leads to a higher error in velocity prediction. Specifically, the error of velocity increases by 54.69% with 50% FLR compared to that without missed frames. However, when the temporal embedding is used, the error decreases significantly by 25.13%. This indicates that the temporal embedding correctly encodes the time information for velocity prediction, alleviating the missed frames issue.

d) *Shallower-Layer Fusion:* Short-term temporal fusion is as critical as long-term fusion since it provides more accurate depth estimation through stereo matching [4]. To study if combining VideoBEV with such short-term fusion could further bring benefits, we implement our VideoBEV based on recently proposed BEVStereo [18] that leverages




Fig. 4. **Visualization results of VideoBEV on the nuScenes val set.** We show the predicted 3D box results of single frame baseline and VideoBEV with ResNet-50 backbone in multi-camera images and bird’s-eye-view. The results of the baseline involving *false negative*, *incorrect object orientation*, and *inaccurate occluded object identifications* that are fixed by VideoBEV are highlighted with dashed circles in green, purple, and blue, respectively.

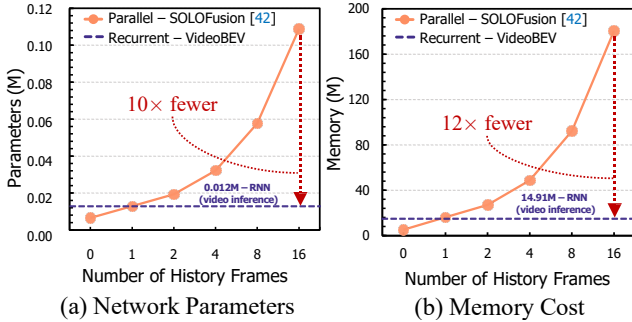


Fig. 5. **Efficiency comparison of two temporal feature fusion modules.** By comparing with parallel style SOLOFusion [4]: (a) Fusion module network parameter; (b) Memory cost of the fusion modules during inference.

short-term fusion. As shown in Tab. VII, with the shallower-layer temporal fusion (only one history frame is fused), the performance of VideoBEV is significantly improved on mAP and mATE. This implies more advanced temporal fusion on low-level features may further improve the 3D perception, leaving space for future investigation.

e) Visualization Analysis: In Fig. 4, we visualize the prediction results of VideoBEV and the single-frame baseline for a qualitative comparison. By leveraging and fusing long-term temporal information, VideoBEV successfully corrects the wrong predictions caused by commonly met issues like false negatives, wrong orientation estimation, and inaccurate identification for occluded objects. It demonstrates the superiority of long-term temporal fusion that may be essential for the understanding of comprehensive driving scenes.

V. CONCLUSIONS

This study investigates a simple recurrent long-term temporal fusion framework based on LSS-based methods for camera-based Bird’s-Eye-View 3D perception, dubbed VideoBEV. Unlike previous works, VideoBEV decouples the recurrent spatiotemporal fusion with a lightweight fusion process. Compared to parallel temporal fusion, VideoBEV’s resource-efficient recurrent fashion yields a superior computation budget, while enjoying the merits of parallel temporal fusion for long-term information modeling. In addition, a dedicated temporal embedding is proposed, which alleviates the frame missing issue in real-world scenarios. Extensive experiments on diverse BEV 3D perception tasks, including 3D object detection, map segmentation, 3D object tracking, and 3D

object motion prediction, are conducted, demonstrating the leading performance of our VideoBEV. This study reveals that long-term temporal information is essential for comprehensive scene understanding in 3D BEV perception. For the first time, we show that a longer-term (*e.g.*, 16 frames in 8s) recurrent temporal fusion brings further benefits for perception accuracy. This study establishes a new baseline for spatiotemporal BEV 3D perception, and we believe our findings will inspire future research into long-term temporal information fusion for autonomous driving.

REFERENCES

- [1] 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,” in *AAAI Conf. Artif. Intell. (AAAI)*, 2023.
- [2] J. Huang and G. Huang, “Bevdet4d: Exploit temporal cues in multi-camera 3d object detection,” *CoRR*, vol. abs/2203.17054, 2022.
- [3] 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,” *CoRR*, vol. abs/2206.01256, 2022.
- [4] J. Park, C. Xu, S. Yang, K. Keutzer, K. Kitani, M. Tomizuka, and W. Zhan, “Time will tell: New outlooks and A baseline for temporal multi-view 3d object detection,” in *Int. Conf. Learn. Represent. (ICLR)*, 2023.
- [5] H. Caesar, V. Bankiti, A. H. Lang, S. Vora, V. E. Liong, Q. Xu, A. Krishnan, Y. Pan, G. Baldan, and O. Beijbom, “nuscnets: A multimodal dataset for autonomous driving,” in *IEEE/CVF Conf. Comput. Vis. Pattern Recog. (CVPR)*, 2020.
- [6] A. Karpathy, G. Toderici, S. Shetty, T. Leung, R. Sukthankar, and L. Fei-Fei, “Large-scale video classification with convolutional neural networks,” in *IEEE/CVF Conf. Comput. Vis. Pattern Recog. (CVPR)*, 2014.
- [7] K. Simonyan and A. Zisserman, “Two-stream convolutional networks for action recognition in videos,” in *Adv. Neural Inform. Process. Syst. (NIPS)*, 2014.
- [8] W. Luo, B. Yang, and R. Urtasun, “Fast and furious: Real time end-to-end 3d detection, tracking and motion forecasting with a single convolutional net,” in *IEEE/CVF Conf. Comput. Vis. Pattern Recog. (CVPR)*, 2018.
- [9] B. Huang, Y. Li, E. Xie, F. Liang, L. Wang, M. Shen, F. Liu, T. Wang, P. Luo, and J. Shao, “Fast-bev: Towards real-time on-vehicle bird’s-eye view perception,” in *Adv. Neural Inform. Process. Syst. Workshop. (NeurIPS Workshop)*, 2022.
- [10] C. Yang, Y. Chen, H. Tian, C. Tao, X. Zhu, Z. Zhang, G. Huang, H. Li, Y. Qiao, L. Lu, J. Zhou, and J. Dai, “Bevformer v2: Adapting modern image backbones to bird’s-eye-view recognition via perspective supervision,” in *IEEE/CVF Conf. Comput. Vis. Pattern Recog. (CVPR)*, 2023.
- [11] S. Hochreiter and J. Schmidhuber, “Long short-term memory,” *Neural Comput.*, vol. 9, no. 8, pp. 1735–1780, 1997.
- [12] K. Cho, B. van Merriënboer, D. Bahdanau, and Y. Bengio, “On the properties of neural machine translation: Encoder-decoder approaches,” in *Empir. Method. Nat. Lang. Process. Worksh. (EMNLP Workshop)*, 2014.

- [13] I. Sutskever, O. Vinyals, and Q. V. Le, “Sequence to sequence learning with neural networks,” in *Adv. Neural Inform. Process. Syst. (NIPS)*, 2014.
- [14] Z. Li, W. Wang, H. Li, E. Xie, C. Sima, T. Lu, Y. Qiao, and J. Dai, “Bevformer: Learning bird’s-eye-view representation from multi-camera images via spatiotemporal transformers,” in *Eur. Conf. Comput. Vis. (ECCV)*, 2022.
- [15] A. Vaswani, N. Shazeer, N. Parmar, J. Uszkoreit, L. Jones, A. N. Gomez, L. Kaiser, and I. Polosukhin, “Attention is all you need,” in *Adv. Neural Inform. Process. Syst. (NIPS)*, 2017.
- [16] J. Chung, Ç. Gülçehre, K. Cho, and Y. Bengio, “Empirical evaluation of gated recurrent neural networks on sequence modeling,” *CoRR*, vol. abs/1412.3555, 2014.
- [17] J. Philion and S. Fidler, “Lift, splat, shoot: Encoding images from arbitrary camera rigs by implicitly unprojecting to 3d,” in *Eur. Conf. Comput. Vis. (ECCV)*, 2020.
- [18] Y. Li, H. Bao, Z. Ge, J. Yang, J. Sun, and Z. Li, “Bevstereo: Enhancing depth estimation in multi-view 3d object detection with dynamic temporal stereo,” in *AAAI Conf. Artif. Intell. (AAAI)*, 2023.
- [19] A. Mousavian, D. Anguelov, J. Flynn, and J. Kosecka, “3d bounding box estimation using deep learning and geometry,” in *IEEE/CVF Conf. Comput. Vis. Pattern Recog. (CVPR)*, 2017.
- [20] T. Wang, X. Zhu, J. Pang, and D. Lin, “FCOS3D: fully convolutional one-stage monocular 3d object detection,” in *Int. Conf. Comput. Vis. Worksh. (ICCV Workshop)*, 2021.
- [21] Z. Tian, C. Shen, H. Chen, and T. He, “FCOS: fully convolutional one-stage object detection,” in *Int. Conf. Comput. Vis. (ICCV)*, 2019.
- [22] Y. Liu, T. Wang, X. Zhang, and J. Sun, “PETR: position embedding transformation for multi-view 3d object detection,” in *Eur. Conf. Comput. Vis. (ECCV)*, 2021.
- [23] X. Chen, H. Ma, J. Wan, B. Li, and T. Xia, “Multi-view 3d object detection network for autonomous driving,” in *IEEE/CVF Conf. Comput. Vis. Pattern Recog. (CVPR)*, 2017.
- [24] A. H. Lang, S. Vora, H. Caesar, L. Zhou, J. Yang, and O. Beijbom, “Pointpillars: Fast encoders for object detection from point clouds,” in *IEEE/CVF Conf. Comput. Vis. Pattern Recog. (CVPR)*, 2019.
- [25] J. Huang, G. Huang, Z. Zhu, and D. Du, “Bevdet: High-performance multi-camera 3d object detection in bird-eye-view,” *CoRR*, vol. abs/2112.11790, 2021.
- [26] Z. Zong, D. Jiang, G. Song, Z. Xue, J. Su, H. Li, and Y. Liu, “Temporal enhanced training of multi-view 3d object detector via historical object prediction,” in *Int. Conf. Comput. Vis. (ICCV)*, pp. 3781–3790, 2023.
- [27] A. Laddha, S. Gautam, G. P. Meyer, C. Vallespi-Gonzalez, and C. K. Wellington, “Rv-fusenet: Range view based fusion of time-series lidar data for joint 3d object detection and motion forecasting,” in *IEEE/RSJ Int. Conf. Intell. Robot. and Syst. (IROS)*, pp. 7060–7066, IEEE, 2021.
- [28] T. Yin, X. Zhou, and P. Krährenbühl, “Center-based 3d object detection and tracking,” in *IEEE/CVF Conf. Comput. Vis. Pattern Recog. (CVPR)*, 2021.
- [29] A. Saha, O. Mendez, C. Russell, and R. Bowden, “Translating images into maps,” in *IEEE Int. Conf. Robot. Autom. (ICRA)*, 2022.
- [30] T. Kanade and M. Okutomi, “A stereo matching algorithm with an adaptive window: Theory and experiment,” *IEEE Trans. Pattern Anal. Mach. Intell. (TPAMI)*, vol. 16, no. 9, pp. 920–932, 1994.
- [31] Z. Wang, C. Min, Z. Ge, Y. Li, Z. Li, H. Yang, and D. Huang, “STS: surround-view temporal stereo for multi-view 3d detection,” in *AAAI Conf. Artif. Intell. (AAAI)*, 2023.
- [32] K. He, X. Zhang, S. Ren, and J. Sun, “Deep residual learning for image recognition,” in *IEEE/CVF Conf. Comput. Vis. Pattern Recog. (CVPR)*, 2016.
- [33] Z. Liu, H. Mao, C. Wu, C. Feichtenhofer, T. Darrell, and S. Xie, “A convnet for the 2020s,” in *IEEE/CVF Conf. Comput. Vis. Pattern Recog. (CVPR)*, 2022.
- [34] H. Zhou, Z. Ge, Z. Li, and X. Zhang, “Matrixvt: Efficient multi-camera to BEV transformation for 3d perception,” *CoRR*, vol. abs/2211.10593, 2022.
- [35] M. Jaderberg, K. Simonyan, A. Zisserman, and K. Kavukcuoglu, “Spatial transformer networks,” in *Adv. Neural Inform. Process. Syst. (NIPS)*, 2015.
- [36] C. Feng, Z. Jie, Y. Zhong, X. Chu, and L. Ma, “Aedet: Azimuth-invariant multi-view 3d object detection,” *CoRR*, vol. abs/2211.12501, 2022.
- [37] T. Wang, X. Zhu, J. Pang, and D. Lin, “FCOS3D: fully convolutional one-stage monocular 3d object detection,” in *Int. Conf. Comput. Vis. Worksh. (ICCV Workshop)*, 2021.
- [38] Y. Wang, V. Guizilini, T. Zhang, Y. Wang, H. Zhao, and J. Solomon, “DETR3D: 3d object detection from multi-view images via 3d-to-2d queries,” in *Annu. Conf. Robot. Learn. (CoRL)*, 2021.
- [39] Y. Li, Y. Chen, X. Qi, Z. Li, J. Sun, and J. Jia, “Unifying voxel-based representation with transformer for 3d object detection,” in *Adv. Neural Inform. Process. Syst. (NeurIPS)*, 2022.
- [40] Y. Jiang, L. Zhang, Z. Miao, X. Zhu, J. Gao, W. Hu, and Y. Jiang, “Polarformer: Multi-camera 3d object detection with polar transformers,” in *AAAI Conf. Artif. Intell. (AAAI)*, 2023.
- [41] J. Gu, C. Hu, T. Zhang, X. Chen, Y. Wang, Y. Wang, and H. Zhao, “Vip3d: End-to-end visual trajectory prediction via 3d agent queries,” *CoRR*, vol. abs/2208.01582, 2022.
- [42] J. Deng, W. Dong, R. Socher, L. Li, K. Li, and L. Fei-Fei, “Imagenet: A large-scale hierarchical image database,” in *IEEE/CVF Conf. Comput. Vis. Pattern Recog. (CVPR)*, 2009.
- [43] Y. Yan, Y. Mao, and B. Li, “SECOND: sparsely embedded convolutional detection,” *Sensors*, vol. 18, no. 10, p. 3337, 2018.
- [44] W. Wang, J. Dai, Z. Chen, Z. Huang, Z. Li, X. Zhu, X. Hu, T. Lu, L. Lu, H. Li, X. Wang, and Y. Qiao, “Internimage: Exploring large-scale vision foundation models with deformable convolutions,” in *IEEE/CVF Conf. Comput. Vis. Pattern Recog. (CVPR)*, 2023.
- [45] A. Hu, Z. Murez, N. Mohan, S. Dudas, J. Hawke, V. Badrinarayanan, R. Cipolla, and A. Kendall, “FIERY: future instance prediction in bird’s-eye view from surround monocular cameras,” in *Int. Conf. Comput. Vis. (ICCV)*, 2021.
- [46] E. Xie, Z. Yu, D. Zhou, J. Philion, A. Anandkumar, S. Fidler, P. Luo, and J. M. Alvarez, “M²bev: Multi-camera joint 3d detection and segmentation with unified birds-eye view representation,” *CoRR*, vol. abs/2204.05088, 2022.
- [47] Y. Hu, J. Yang, L. Chen, K. Li, C. Sima, X. Zhu, S. Chai, S. Du, T. Lin, W. Wang, L. Lu, X. Jia, Q. Liu, J. Dai, Y. Qiao, and H. Li, “Goal-oriented autonomous driving,” *CoRR*, vol. abs/2212.10156, 2022.
- [48] X. Zhou, V. Koltun, and P. Krährenbühl, “Tracking objects as points,” in *Eur. Conf. Comput. Vis. (ECCV)*, 2020.
- [49] M. Chaabane, P. Zhang, J. R. Beveridge, and S. O’Hara, “DEFT: detection embeddings for tracking,” *CoRR*, vol. abs/2102.02267, 2021.
- [50] P. Li and J. Jin, “Time3d: End-to-end joint monocular 3d object detection and tracking for autonomous driving,” in *IEEE/CVF Conf. Comput. Vis. Pattern Recog. (CVPR)*, 2022.
- [51] H. Hu, Y. Yang, T. Fischer, T. Darrell, F. Yu, and M. Sun, “Monocular quasi-dense 3d object tracking,” *IEEE Trans. Pattern Anal. Mach. Intell. (TPAMI)*, vol. 45, no. 2, pp. 1992–2008, 2023.
- [52] N. Marinello, M. Proesmans, and L. V. Gool, “Tripletrack: 3d object tracking using triplet embeddings and LSTM,” in *IEEE Conf. Comput. Vis. Pattern Recog. Worksh. (CVPR Workshop)*, 2022.
- [53] T. Zhang, X. Chen, Y. Wang, Y. Wang, and H. Zhao, “MUTR3D: A multi-camera tracking framework via 3d-to-2d queries,” in *IEEE Conf. Comput. Vis. Pattern Recog. Worksh. (CVPR Workshop)*, 2022.
- [54] S. Chen, X. Wang, T. Cheng, Q. Zhang, C. Huang, and W. Liu, “Polar parametrization for vision-based surround-view 3d detection,” *CoRR*, vol. abs/2206.10965, 2022.
- [55] Y. Shi, J. Shen, Y. Sun, Y. Wang, J. Li, S. Sun, K. Jiang, and D. Yang, “SRCN3D: sparse R-CNN 3d surround-view camera object detection and tracking for autonomous driving,” *CoRR*, vol. abs/2206.14451, 2022.
- [56] T. Fischer, Y. Yang, S. Kumar, M. Sun, and F. Yu, “CC-3DT: panoramic 3d object tracking via cross-camera fusion,” *CoRR*, vol. abs/2212.01247, 2022.
- [57] Z. Pang, J. Li, P. Tokmakov, D. Chen, S. Zagoruyko, and Y. Wang, “Standing between past and future: Spatio-temporal modeling for multi-camera 3d multi-object tracking,” *CoRR*, vol. abs/2302.03802, 2023.
- [58] J. Yang, E. Yu, Z. Li, X. Li, and W. Tao, “Quality matters: Embracing quality clues for robust 3d multi-object tracking,” *CoRR*, vol. abs/2208.10976, 2022.
- [59] X. Lin, T. Lin, Z. Pei, L. Huang, and Z. Su, “Sparse4d: Multi-view 3d object detection with sparse spatial-temporal fusion,” *CoRR*, vol. abs/2211.10581, 2022.
- [60] M. Liang, B. Yang, W. Zeng, Y. Chen, R. Hu, S. Casas, and R. Urtasun, “Pnpnet: End-to-end perception and prediction with tracking in the loop,” in *IEEE/CVF Conf. Comput. Vis. Pattern Recog. (CVPR)*, 2020.
- [61] H. Song, W. Ding, Y. Chen, S. Shen, M. Y. Wang, and Q. Chen, “Pip: Planning-informed trajectory prediction for autonomous driving,” in *Eur. Conf. Comput. Vis. (ECCV)*, 2020.
- [62] O. Ronneberger, P. Fischer, and T. Brox, “U-net: Convolutional networks for biomedical image segmentation,” in *Int. Conf. Medical Image Comput. Assist. Interv. (MICCAI)*, 2015.
- [63] N. Peri, J. Luiten, M. Li, A. Osep, L. Leal-Taixé, and D. Ramanan, “Forecasting from lidar via future object detection,” in *IEEE/CVF Conf. Comput. Vis. Pattern Recog. (CVPR)*, 2022.