

MemFusionMap: Working Memory Fusion for Online Vectorized HD Map Construction

Jingyu Song^{1*} Xudong Chen² Liupei Lu² Jie Li² Katherine A. Skinner¹

¹University of Michigan ²NVIDIA

Abstract

High-definition (HD) maps provide environmental information for autonomous driving systems and are essential for safe planning. While existing methods with single-frame input achieve impressive performance for online vectorized HD map construction, they still struggle with complex scenarios and occlusions. We propose MemFusionMap, a novel temporal fusion model with enhanced temporal reasoning capabilities for online HD map construction. Specifically, we contribute a working memory fusion module that improves the model’s memory capacity to reason across a history of frames. We also design a novel temporal overlap heatmap to explicitly inform the model about the temporal overlap information and vehicle trajectory in the Bird’s Eye View space. By integrating these two designs, MemFusionMap significantly outperforms existing methods while also maintaining a versatile design for scalability. We conduct extensive evaluation on open-source benchmarks and demonstrate a maximum improvement of 5.4% in mAP over state-of-the-art methods. The project page for MemFusionMap is <https://song-jingyu.github.io/MemFusionMap>.

1. Introduction

With recent advances in leveraging Bird’s Eye View (BEV) representations, perception capabilities of Autonomous Vehicles (AVs) have been largely enhanced [18, 23, 27, 31, 41]. Among all the BEV perception tasks, constructing High-Definition (HD) maps in an online manner has received increasing attention from the AV research community [17, 19, 22, 39]. Specifically, online HD map construction aims to use onboard vehicle sensor data to reconstruct road elements such as lanes, pedestrian crossings and road boundaries on-the-fly. The reconstructed HD maps are essential for downstream planning tasks such as prediction and planning [13]. This task is of great significance in advancing large-scale deployment of AVs as it avoids the need

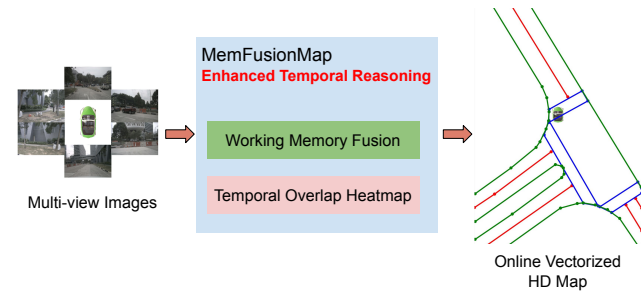


Figure 1. We propose MemFusionMap to improve the temporal reasoning capability for online vectorized HD map construction. MemFusionMap takes in multi-view images and leverages the proposed working memory fusion module and temporal overlap heatmap to output a vectorized HD map online.

for a sophisticated and labor-intensive offline HD map generation process [22].

Existing works have demonstrated the effectiveness of formulating online HD map construction in an end-to-end framework [22, 24, 28, 39]. Though significant advances have been made in this field, many methods fall short of leveraging temporal information, which has been shown to be very beneficial, especially during instances of occlusion [4, 33, 39]. StreamMapNet [39] is the first end-to-end vectorized HD map construction work leveraging temporal cues. By maintaining a recurrent BEV memory feature map and propagating map queries, StreamMapNet conducts temporal fusion effectively and significantly outperforms previous non-temporal baselines [19, 22], demonstrating the importance of incorporating temporal cues for online HD map construction. The success of StreamMapNet [39] has inspired follow-up works to further improve it by incorporating auxiliary learning tasks or by integrating ideas from the tracking literature [4, 33].

Still, we highlight the challenges of directly accumulating all the temporal features into a single memory feature map, as proposed in [39]. The network can struggle to reason about the entire history due to limited memory capacity in complex road environments since the network only has access to the latest memory feature and current BEV fea-

*Work done during an internship at NVIDIA. Corresponding at: jingyuso@umich.edu

ture. Furthermore, this design can struggle with occlusion. A typical example is when a sudden occlusion occurs due to a moving truck near the ego-vehicle. This will trigger a bad update to the memory feature due to challenging 2D-to-BEV projection, and thus will affect all the future predictions. A similar issue is also noted in the video object segmentation field [7,43]. We share the same insight that maintaining working memory features from a subset of historical frames is helpful, which is also pointed out by [4]. Lastly, we argue the importance of maintaining explicit temporal overlap information, as the model can leverage this input to reason about temporal information more effectively.

Given these insights, we propose MemFusionMap (Fig. 1), a novel memory fusion framework designed for online vectorized HD Map construction. We design a working memory buffer to maintain working memory features in a fixed-lag manner. The working memory features are propagated recursively to align with the current field-of-view of the ego-vehicle. The fixed-lag design ensures bounded memory usage to fit real-world deployment needs. In addition, we propose a novel use of a temporal overlap heatmap, represented as a BEV image with a single channel. The heatmap score indicates the number of times each grid cell falls within the field-of-view. We recurrently propagate the temporal overlap heatmap. To facilitate the model’s temporal reasoning capabilities, the temporal overlap heatmap is fused with the working memory features by a convolutional temporal fusion block, whose output is fed into a standard detector head to generate road element predictions.

To summarize, our main contributions are as follows:

- We propose a simple yet effective model to fuse working memory features in BEV space for online vectorized HD map construction. MemFusionMap focuses on improving the network’s temporal reasoning capability while also maintaining a versatile design for scalability and compatibility.
- We propose a novel design of maintaining a temporal overlap heatmap, providing a strong cue for the model to reason across a history of frames and also implicitly encoding valuable insights of the vehicle’s trajectory.
- We conduct extensive evaluation on nuScenes [1] and Argoverse2 [35] to demonstrate the effectiveness of MemFusionMap. The proposed method significantly outperforms the state-of-the-art method [39], achieving a maximum improvement of 5.4% in mAP.

2. Related Works

2.1. Online Vectorized HD Map Construction

Constructing HD maps is one of the key tasks for AV perception. While traditional SLAM-based methods [30, 40]

suffer from high cost during offline construction and maintenance, online HD map construction solutions have obtained increasing attention as they can effectively simplify the map construction process and handle map changes. HDMaNet [17] marks a significant step of moving towards vectorized HD map representations, which are preferred by the downstream tasks over the rasterized map representations that are generated from previous segmentation methods [18, 27, 42] or lane detection methods [5, 12, 32]. VectorMapNet [22] proposes the first end-to-end vectorized HD map construction framework where a DETR [2] based map decoder is used to generate the map elements’ prediction with an auto-regressive transformer. Its performance is further improved by MapTR [19] and MapTRV2 [20] by treating the map construction task as a point set prediction problem and adding auxiliary tasks and decoder design improvement. Other works focus on different map element representations: BeMapNet [28] and PivotNet [10] explore using Bézier curves and pivot-based points, respectively. Additionally, other methods investigate network design improvement and training pipeline enhancements [3, 21, 24, 38]. More recently, StreamMapNet [39] and its follow-up works [4, 33] demonstrate notable improvement by leveraging temporal fusion for this task, which will be discussed in Sec. 2.2.

2.2. Temporal Fusion for HD Map Construction

Fusing temporal information has proven effective in various works in domains like 3D object detection and scene completion [14, 18, 26, 34, 36]. HDMaNet [17] is the first to demonstrate the benefits of temporal fusion in online HD map construction. It directly applies max pooling to compress temporal information into a BEV feature map, which is fed into the decoder. StreamMapNet [39] develops the first temporal fusion pipeline in an end-to-end framework for this task. In addition to propagating and reusing queries in the decoder, it proposes a BEV temporal fusion paradigm leveraging the streaming strategy following [14, 34]. StreamMapNet [39] compresses the history of information into a latent memory feature and propagates the latent memory feature recurrently. It applies a Gated Recurrent Unit (GRU) [8] to fuse the latent memory feature with the current BEV feature from the BEV feature encoder.

The success of StreamMapNet has motivated subsequent works such as SQD-MapNet [33] and MapTracker [4]. Inspired by DN-DETR [16], SQD-MapNet [33] designs a stream query denoising approach to further improve over StreamMapNet. However, SQD-MapNet is only evaluated on the original splits of nuScenes [1] and Argoverse2 [35] that are prone to overfitting, contradicting the main motivation of online HD map construction. MapTracker [4] explores leveraging additional ground truth of tracked road elements to enhance the query propagation paradigm. Mem-

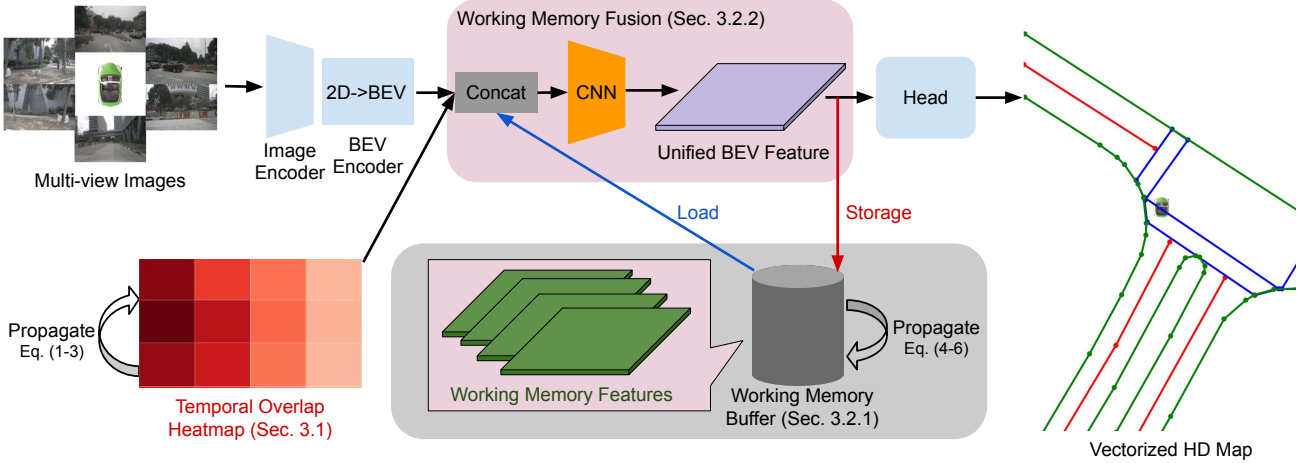


Figure 2. Overall architecture of MemFusionMap. MemFusionMap takes in multi-view images and extracts a BEV feature via an image encoder and a BEV feature encoder. The encoded BEV feature is fused with working memory features and the temporal overlap heatmap to obtain a unified BEV feature. The unified feature is fed into a decoder head to generate HD map prediction.

FusionMap shares the same intuition as MapTracker [4] that the redundancy of maintaining a subset of past frames can improve the model’s robustness. However, we design MemFusionMap with a strong emphasis on scalability and versatility for practical application. Notably, while MapTracker fuses the last 20 frames of memory latents selected according to the distance strides, our proposed method, MemFusionMap, designs a simple yet effective memory fusion module, which only uses 4 past frames to avoid computational overhead and to achieve better training efficiency and runtime. MemFusionMap also covers the perception range of $100 \times 50 m$, which has more practical value for real-world AV deployment, while MapTracker [4] omits this extended perception range. Furthermore, SQD-MapNet [33] and MapTracker [4] require implementing additional task and loss modules, which can cause extra overhead when integrating into existing multi-task AV perception pipelines. MapTracker also requires additional post-processing to generate the tracks of map elements, which can be time-consuming and potentially incompatible with large-scale databases that use an auto labeling process that cannot guarantee quality assured labels for all frames. In contrast, MemFusionMap maintains a compact design with better versatility that can be directly applied to existing pipelines by switching the BEV temporal fusion module. Lastly, MemFusionMap also proposes a novel method for maintaining and integrating a temporal overlap heatmap, which is demonstrated to significantly improve the model’s temporal reasoning capability.

3. MemFusionMap

The overall framework of MemFusionMap is shown in Fig. 2. The pipeline of MemFusionMap is inspired by

StreamMapNet [39]. However, MemFusionMap proposes a novel structure for the temporal BEV fusion module. MemFusionMap firstly uses a shared ResNet-50 [15] image encoder to extract image features. A BEV feature encoder further projects the image features in 2D to BEV space. The projected BEV feature is denoted as $\mathcal{F}_{\text{BEV}} \in \mathbb{R}^{C \times H \times W}$, where C is the BEV feature dimension, and H and W represent the spatial dimensions of the BEV feature. The BEV feature is fed into the memory fusion block, which will be introduced in the following sections. The output of the memory fusion block is a unified BEV feature map containing temporal cues, denoted as $\tilde{\mathcal{F}}_{\text{BEV}} \in \mathbb{R}^{C \times H \times W}$. Following [39], we use a DETR [2] based transformer decoder with a multi-point attention mechanism for extended perception range. We also maintain a memory buffer to properly propagate and reuse history map queries.

The following sub-sections discuss the key contributions of the proposed architecture in more detail. The core contribution of MemFusionMap is its temporal BEV fusion design, which consists of a temporal overlap heatmap (Sec. 3.1) and the working memory fusion module (Sec. 3.2).

3.1. Temporal Overlap Heatmap

We propose a novel paradigm to maintain a temporal overlap heatmap to explicitly inform the model of the temporal overlap information. Figure 3 illustrates the process of maintaining the heatmap. Starting from the first frame of each clip, we initialize the temporal overlap heatmap \mathbf{H}_{t_0} as follows:

$$\mathbf{H}_{t_0} = \mathbf{1}^{1 \times H \times W}. \quad (1)$$

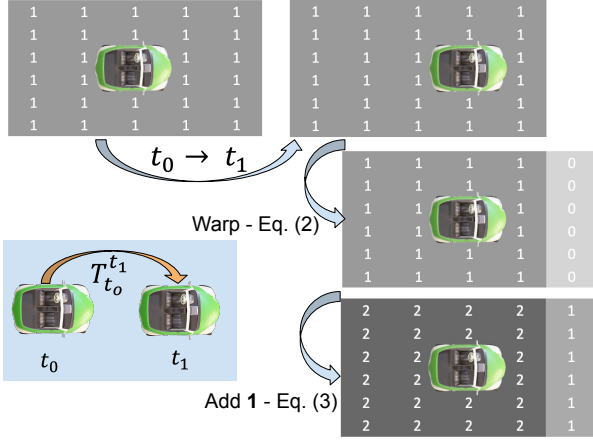


Figure 3. Example propagation process of the temporal overlap heatmap from t_0 to t_1 .

Specifically, \mathbf{H}_{t_0} has the same spatial size as \mathcal{F}_{BEV} and is single-channel. Each pixel of \mathbf{H} indicates the temporal overlap score of this cell in the corresponding BEV space. \mathbf{H}_{t_0} is initialized with all 1s, meaning the whole BEV has not been seen before. Then, as the second frame (t_1) comes in, we first propagate the temporal overlap heatmap to the new vehicle position as follows:

$$\hat{\mathbf{H}}_{t_1} = \text{Warp}(\mathbf{H}_{t_0}, T_{t_0}^{t_1}), \quad (2)$$

where $\hat{\mathbf{H}}_{t_1}$ is the propagated and warped overlap heatmap from t_0 to t_1 according to the egomotion $T_{t_0}^{t_1}$. The warping is implemented using the grid sample function in PyTorch with zero padding mode for the out-of-bound grids. We then obtain the final heatmap for t_1 as:

$$\mathbf{H}_{t_1} = \hat{\mathbf{H}}_{t_1} + \mathbf{1}^{1 \times H \times W}. \quad (3)$$

Following this strategy, we keep updating the temporal overlap heatmap recurrently. At each timestamp, the overlap heatmap is fed into the memory fusion block, which will be discussed in Sec. 3.2.

One main benefit of maintaining this temporal overlap heatmap is that it provides informative insights into the regions with heavy overlaps and those without. With this information, the model can more adaptively reason across the current BEV feature and the historical features. Specifically, in the overlap heatmap \mathbf{H}_t , a smaller value indicates that a region is less temporally overlapped, meaning more information from the current BEV feature should be trusted. Conversely, a larger value suggests a region with greater overlap, implying that information from the memory latents could be more valuable. Furthermore, the overlap heatmap naturally encodes the vehicle trajectory, providing the model with additional knowledge about vehicle motion to enable more adaptive temporal reasoning.

We demonstrate this benefit with two visual examples. As shown in Fig. 4, we present two example overlap




Figure 4. Example temporal overlap heatmap resulted from different speeds. The comparison demonstrates that the overlap heatmap encodes vehicle moving speed.

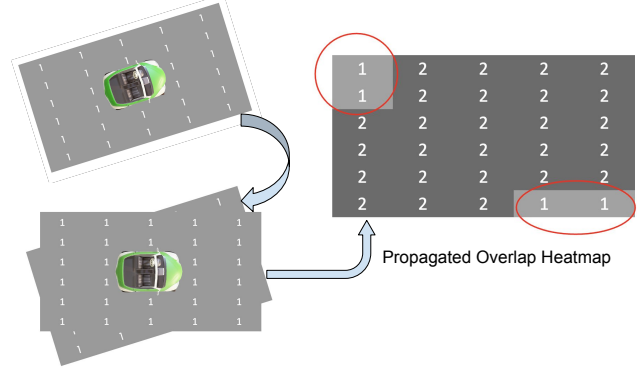


Figure 5. Example temporal overlap heatmap resulted from a vehicle turning. We highlight the two noticeable edges at the corner. This pattern encodes insights into the vehicle’s trajectory.

heatmaps resulting from different speeds. The heatmap on the right exhibits a larger area with a value of 1, indicating that a larger area is newly seen, which is a result of higher-speed egomotion. We present another example of a vehicle turning in Fig. 5. The propagated overlap heatmap reveals two noticeable edges at the corner. These low-level feature patterns can be easily detected by a convolutional neural network, providing critical insights into the vehicle’s trajectory.

3.2. Memory Fusion Block

We design a memory fusion block to synthesize memory features with the temporal overlap heatmap. As discussed in Sec. 1, we highlight the importance of increasing the memory capacity for better temporal reasoning capability and robustness against occlusion. We aggregate a fixed number of BEV features from historical frames to fuse with the current BEV feature. We follow [7] to use the terminology used in cognitive science [9, 11, 37] and refer to the subset of historical frames as working memory features. As shown in Fig. 2, we develop a working memory buffer to properly handle storage and loading of working memory features. We also design a memory fusion block to temporally reason across working memory features, aided by the temporal overlap heatmap.

3.2.1 Working Memory Buffer

We design a working memory buffer to manage the memory features and update them recurrently. In the streaming training schedule, at time t , before memory fusion, the working memory feature $\mathcal{F}_{\text{WM}}^t$ is defined as follows:

$$\mathcal{F}_{\text{WM}}^t = \{\tilde{\mathcal{F}}_{\text{BEV}}^{t-T_{\text{WM}}}, \tilde{\mathcal{F}}_{\text{BEV}}^{t-T_{\text{WM}}+1}, \dots, \tilde{\mathcal{F}}_{\text{BEV}}^{t-2}, \tilde{\mathcal{F}}_{\text{BEV}}^{t-1}\}, \quad (4)$$

where $\tilde{\mathcal{F}}_{\text{BEV}}^t$ is the unified BEV feature after memory fusion at time t , and T_{WM} is the working memory capacity. Progressing towards time $t+1$, we first add the new feature and drop the oldest feature:

$$\mathcal{F}_{\text{WM}}^{t+1} = \{\tilde{\mathcal{F}}_{\text{BEV}}^{t-T_{\text{WM}}+1}, \tilde{\mathcal{F}}_{\text{BEV}}^{t-T_{\text{WM}}+2}, \dots, \tilde{\mathcal{F}}_{\text{BEV}}^{t-1}, \tilde{\mathcal{F}}_{\text{BEV}}^t\}. \quad (5)$$

We also warp all the working memory features to align with $\mathcal{F}_{\text{BEV}}^{t+1}$:

$$\mathcal{F}_{\text{WM}}^{t+1} = \text{Warp}(\mathcal{F}_{\text{WM}}^{t+1}, T_t^{t+1}), \quad (6)$$

where T_t^{t+1} is the transformation from t to $t+1$. Additionally, at t_0 , the beginning of a sequence, we initialize $\mathcal{F}_{\text{WM}}^{t_0}$ by repeating $\mathcal{F}_{\text{BEV}}^{t_0}$ for T_{WM} times. After the memory fusion at t_0 , we replace $\mathcal{F}_{\text{BEV}}^{t_0}$ with $\tilde{\mathcal{F}}_{\text{BEV}}^{t_0}$ and warp them to form $\mathcal{F}_{\text{BEV}}^{t_1}$. The following process of updating working memory follows the strategy mentioned above. Meanwhile, the working memory buffer also serves to store the temporal overlap heatmap. The propagation process for the heatmap is discussed in Sec. 3.1.

3.2.2 Working Memory Fusion

We propose a working memory fusion module to obtain a unified BEV feature $\tilde{\mathcal{F}}_{\text{BEV}}$ fused with temporal cues. The goal of the memory fusion block is to effectively utilize the temporal information in working memory features, aided by the temporal overlap heatmap, and to fuse them with the current BEV feature to obtain a unified BEV feature with rich temporal information.

Specifically, at time t , the input to this module consists of the working memory features $\mathcal{F}_{\text{WM}}^t \in \mathbb{R}^{C_{\text{WM}} \times H \times W}$, the temporal overlap heatmap $\mathbf{H}_t \in \mathbb{R}^{1 \times H \times W}$, and the BEV feature $\mathcal{F}_{\text{BEV}}^t \in \mathbb{R}^{C \times H \times W}$ obtained from the BEV feature encoder. Here C is the feature channel dimension of the BEV feature encoder, $C_{\text{WM}} = T_{\text{WM}} \times C$, and H and W represent the spatial size of the BEV feature. The output of the working memory fusion module is the unified feature $\tilde{\mathcal{F}}_{\text{BEV}}^t \in \mathbb{R}^{C \times H \times W}$. In the working memory fusion module, the first step is to extract low-level features from the temporal overlap heatmap:

$$\tilde{\mathbf{H}}_t = \text{sigmoid}(\text{Conv}_{\mathbf{H}}(\mathbf{H}_t)) \in \mathbb{R}^{C_{\mathbf{H}} \times H \times W}, \quad (7)$$

where $\text{Conv}_{\mathbf{H}}$ is a convolution block composed of 3 convolution layers with ReLU, and $C_{\mathbf{H}}$ is the feature dimension of the learned temporal overlap heatmap feature. We

apply a sigmoid function on the output temporal overlap feature to bound its value. Subsequently, we use a convolutional memory fusion block to obtain the unified BEV feature $\tilde{\mathcal{F}}_{\text{BEV}}^t$ following:

$$\tilde{\mathcal{F}}_{\text{BEV}}^t = \text{LayerNorm}(\text{Conv}_{\text{Mem}}(\text{Concat}(\mathcal{F}_{\text{WM}}^t, \tilde{\mathbf{H}}_t, \mathcal{F}_{\text{BEV}}^t))), \quad (8)$$

where Conv_{Mem} is composed of 3 convolution layers with ReLU. We follow [39] to apply layer normalization to improve training stability.

When designing the working memory fusion module, we keep the HD map construction task in mind. We observe that a unique feature for this task is that map elements, such as lane markings, can have elongated shapes, which inspires us to enhance the receptive field of the memory fusion module to let temporal features interact more effectively both temporally and spatially. Therefore, we apply dilated (atrous) convolution [6] in Conv_{Mem} to further improve the working memory fusion module.

3.3. Training Loss

During training, the overall map loss \mathcal{L}_{map} is defined as in [39]:

$$\mathcal{L}_{\text{map}} = \lambda_1 \mathcal{L}_{\text{Focal}} + \lambda_2 \mathcal{L}_{\text{line}} + \lambda_3 \mathcal{L}_{\text{trans}}, \quad (9)$$

where $\mathcal{L}_{\text{Focal}}$ is the classification matching cost, $\mathcal{L}_{\text{line}}$ is the polyline-wise matching cost, $\mathcal{L}_{\text{trans}}$ is the auxiliary transformation loss for temporal query propagation, and λ_1 - λ_3 are pre-defined loss weights. We refer to [39] for more details about the loss function.

4. Experiments

4.1. Datasets and Splits

We evaluate our methods on the open-source nuScenes [1] and Argoverse2 [35] benchmarks. The nuScenes benchmark has an annotation keyframe rate of 2Hz, with 6 cameras synchronized. Argoverse2 is annotated with 10Hz with 7 ring cameras and 2 stereo cameras. We follow [39] to unify the frame rate of the Argoverse2 dataset to 2Hz, and use all 6 cameras of nuScenes and the 7 ring cameras of Argoverse2 as the input to the model.

As noted by [29, 39], the default splits in both nuScenes and Argoverse2 can be prone to overfitting as there is a large portion of overlap between training and validation sets. We highlight that the ultimate goal of the online HD map construction task is to develop a generalizable model that can adapt to unseen environments with robustness to potential map changes, instead of simply memorizing the training set. Therefore, we focus on evaluating MemFusionMap using the split proposed in [29] for nuScenes and the split proposed in [39] for Argoverse2. We refer to these non-overlapped splits as new splits and the default splits as original splits in the following sub-sections.

Range	Method	Backbone	Image Size	Epoch	AP _{ped}	AP _{div}	AP _{bound}	mAP
60 × 30 m	VectorMapNet [†] [22]	R50	256 × 480	120	15.8	17.0	21.2	18.0
	MapTR [†] [19]	R50	480 × 800	24	6.4	20.7	35.5	20.9
	StreamMapNet [39]	R50	480 × 800	24	30.7	29.9	41.8	34.1
	MemFusionMap (ours)	R50	480 × 800	24	38.3	32.1	43.6	38.0
100 × 50 m	VectorMapNet [†] [22]	R50	256 × 480	120	12.0	8.1	6.3	8.8
	MapTR [†] [19]	R50	480 × 800	24	8.3	16.0	20.0	14.8
	StreamMapNet [39]	R50	480 × 800	24	25.4	19	22.2	22.2
	MemFusionMap (ours)	R50	480 × 800	24	30.2	22.6	29.9	27.6

Table 1. Performance comparison with baseline methods on the new nuScenes split at both 30 m and 50 m perception ranges. [†]: results reported in StreamMapNet [39] paper. Abbreviations: pedestrian crossing (ped), lane divider (div), road boundary (bound). For StreamMapNet, we report our reproduced results, which match the original paper.

4.2. Evaluation Metrics

We follow [33, 39] to consider both a small perceptual range (60 × 30 m) and a large perceptual range (100 × 50 m). We mainly use the geofenced new splits proposed by [29, 39]. We use Average Precision (AP) as the evaluation metric. We calculate AP using the distance threshold of {0.5m, 1.0m, 1.5m} for the small perceptual range and {1.0m, 1.5m, 2.0m} for the large perceptual range. Following common practices in existing works, we evaluate three types of map elements: pedestrian crossings, lane dividers, and road boundaries. We calculate the mean AP (mAP) score of the three types of map elements.

4.3. Implementation Details

We build MemFusionMap upon the codebase of StreamMapNet [39]. We use ResNet-50 [15] as the image feature encoder. We use BEVFormer [18] as the BEV feature encoder to project features from 2D to BEV space. C is set as 256 across all experiments. We set H as 50 and W as 100. In the memory fusion module, we set T_{WM} as 4 and C_H as 32. The model is trained on 8 NVIDIA V100 GPUs with a total batch size of 32. The learning rate is set to 5×10^{-4} and an AdamW [25] optimizer is used. To train MemFusionMap, we apply the streaming training strategy [14, 34] to detach the memory tensor so gradients do not propagate back to previous frames. Following [26, 39], we apply a two-stage training procedure, where the input of the first stage is single-frame. It is followed by a temporal training stage, where the input is consecutive sequences obtained from dividing the raw training sequence into 2 random sequences. We include more implementation details in the supplementary material.

4.4. Quantitative Comparison

We compare MemFusionMap against state-of-the-art (SOTA) baselines across the nuScenes [1] and Argoverse2 [35] datasets with two perceptual ranges using the new splits [29, 39]. The baselines in the comprehensive comparison include StreamMapNet [39] and single-frame

baselines, including VectorMapNet [22] and MapTR [19]. As to the follow-up works of StreamMapNet (i.e., SQD-MapNet [33] and MapTracker [4]), we conduct separate comparisons. This is due to the lack of support for either the extended 100 × 50 m perception range in [4], which holds greater practical value for AV deployment, or the absence of open-source implementation and results evaluated using the new splits in [33]. We report the number of epochs to demonstrate the training efficiency and measure the runtime on a desktop equipped with an NVIDIA A6000 GPU.

We first show an overall comparison on the new split [29] of the nuScenes [1] benchmark. At both ranges, MemFusionMap achieves significant improvement over baselines across all metrics, demonstrating the effectiveness of the proposed memory fusion architecture. Similarly, as shown in Tab. 2, MemFusionMap demonstrates superior performance over other methods across all categories on the Argoverse2 benchmark, validating the consistent improvement brought by MemFusionMap. We also highlight the greater improvement in the long perception range setting compared to the short range setting on both benchmarks, which is more applicable to real-world deployment. As MemFusionMap is built upon StreamMapNet and shares a similar training objective, we emphasize the advancement brought by MemFusionMap over StreamMapNet, which directly shows the improvement brought by the proposed working memory fusion framework over the previous GRU-based temporal fusion paradigm.

Furthermore, we show a comparison with open-source temporal baselines on mAP and runtime in Tab. 3. Though MemFusionMap significantly outperforms StreamMapNet [39] while maintaining a similarly fast runtime, it falls behind the recent SOTA work MapTracker [4] in mAP. We recognize the contribution of MapTracker of formulating the HD map construction task as a tracking problem, and think merging MemFusionMap and MapTracker is an interesting future direction. However, it is worth mentioning that MapTracker is benefited from additional ground truth of tracked map elements and requires a time-consuming post-processing process. This may have scalability and compat-

Range	Method	Backbone	Image Size	Epoch	AP_{ped}	AP_{div}	AP_{bound}	mAP
$60 \times 30 m$	VectorMapNet [†] [22]	R50	384×384	120	35.6	34.9	37.8	36.1
	MapTR [†] [19]	R50	608×608	30	48.1	50.4	55	51.1
	StreamMapNet [39]	R50	608×608	30	57.5	55.9	61.5	58.3
	MemFusionMap (ours)	R50	608×608	30	59.3	57.2	65.1	60.6
$100 \times 50 m$	VectorMapNet [†] [22]	R50	384×384	120	32.4	20.6	24.3	25.7
	MapTR [†] [19]	R50	608×608	30	46.3	36.3	38	40.2
	StreamMapNet [39]	R50	608×608	30	60.2	45.5	49	51.5
	MemFusionMap (ours)	R50	608×608	30	64.1	47.3	51.5	54.3

Table 2. Performance comparison with baseline methods on the new Argoverse2 split at both $30 m$ and $50 m$ perception ranges. [†]: results reported in StreamMapNet [39] paper. Abbreviations: pedestrian crossing (ped), lane divider (div), road boundary (bound). For StreamMapNet, we report our reproduced results, which match the original paper.

Method	Epoch	mAP	FPS
StreamMapNet [39]	24	34.1	15.3
MapTracker [4]	72	40.3	11.6
MemFusionMap (ours)	24	38.0	15.1

Table 3. Quantitative comparison of training efficiency and runtime on the new nuScenes split at $30 m$ range. MapTracker [4] lacks support for the extended $50 m$ range setting. SQD-MapNet [33] is not included as it is not open-source.

Dataset	Method	Epoch	mAP
nuScenes	StreamMapNet [39]	30*	63.4
	SQD-MapNet [33]	24	65
	MemFusionMap (ours)	30*	65.2
Argoverse2	StreamMapNet [†] [39]	30	61.5
	SQD-MapNet [33]	30	63.3
	MemFusionMap (ours)	30	63.8

Table 4. Performance comparison on the original split of nuScenes at $30 m$ range. *: we follow the practice of the StreamMapNet code repository to train 30 epochs on the original split. [†]: results reported in SQD-MapNet [33] paper.

ibility issues with large-scale databases that use an auto labeling process that cannot guarantee all frames are quality assured, which could render a direct comparison on mAP inappropriate. Additionally, as shown in Tab. 3, MemFusionMap demonstrates strong performance without requiring a long training schedule, which is very valuable for integrating to large-scale databases. Compared with MapTracker [4], MemFusionMap also improves runtime significantly. Considering these factors, we still think MemFusionMap has unique benefits for the community.

We also show additional comparison on the original split (Tab. 4) for completeness since SQD-MapNet [33] does not report results on the new split and is not open-source. The evaluation is based on the $60 \times 30 m$ range since SQD-MapNet uses a larger BEV spatial dimension for the extended range. Results in Tab. 4 demonstrate that MemFusionMap still outperforms the baselines, though we note that this split is prone to overfitting.

Index	Method	mAP
(a)	StreamMapNet	22.2
(b)	– GRU BEV fusion	19.5
(c)	+ Working Memory Fusion	25.5
(d)	+ Dilated convolution	25.9
(e)	+ Temporal overlap heatmap	27.6
	MemFusionMap	27.6

Table 5. Ablation study of each proposed component in MemFusionMap evaluated on nuScenes with the new split at $50 m$ range.

Module	Memory Capacity T_{WM}	mAP	FPS
Memory Fusion	1	24.3	15.4
	2	25.9	15.2
	4	27.6	15.1
	6	26.2	14.8
	8	27.0	14.7

Table 6. Ablation study of the working memory capacity.

4.5. Ablation Studies

We conduct extensive ablation studies to further break down the performance improvement brought by MemFusionMap and justify our design choice. We conduct all of the ablation studies using the new split of nuScenes at $100 \times 50 m$ range. We include some ablation studies in the supplementary material.

We show an ablation study in Tab. 5. We first remove the GRU-based temporal BEV fusion from StreamMapNet. We notice a notable performance drop, which shows the importance of temporal fusion for this task. We then test adding a convolutional working memory fusion module. It is observed that a simple convolution block can effectively improve temporal reasoning capability with the help of increased memory capacity by maintaining working memory features. The comparison between (a) and (c) also demonstrates the challenge of accumulating the entire history into a single memory latent. The improvement from (c) to (d) validates the importance of using dilated convolution for larger receptive field. Progressing from (d) to (e), the 1.7% improvement on an already strong model justifies the signif-

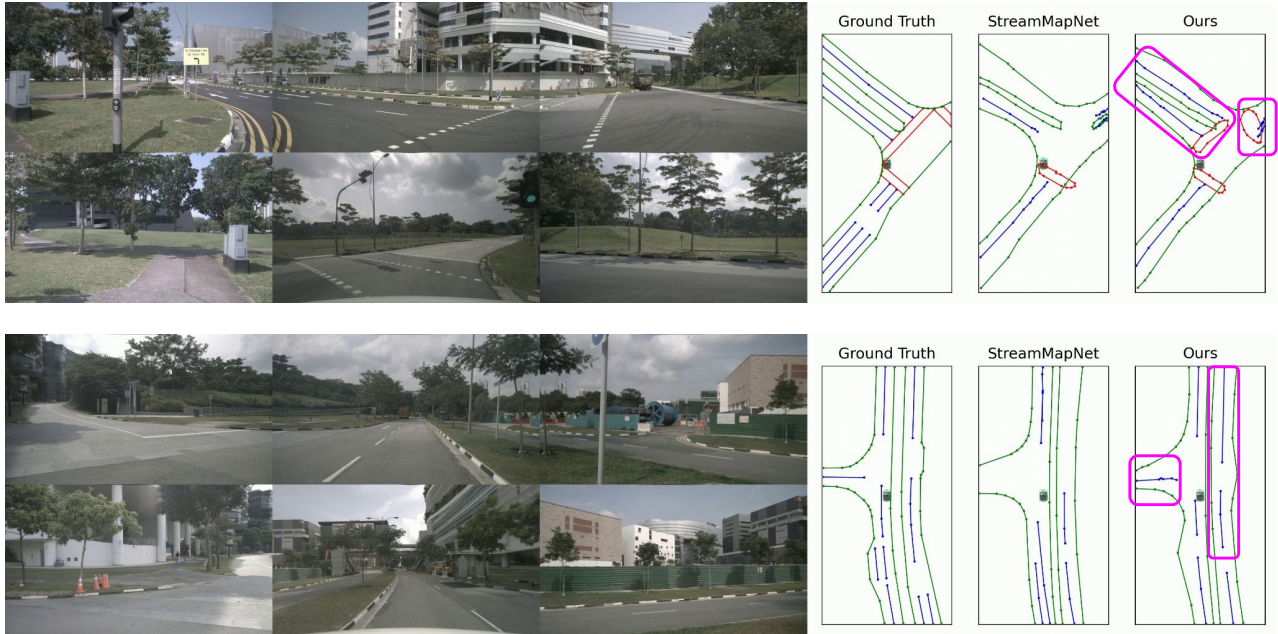


Figure 6. We show qualitative results on nuScenes [1] at the extended 50 m range. We show the input multi-view images on the left. In the HD maps, green denotes road boundaries, red denotes pedestrian crossings, and blue denotes lane lines. Best viewed in color.

Module	Conv Dilation	mAP
Memory	(1,1)	26.4
	(2,2)	27.6
Fusion	(2,3)	26.4
	(3,3)	27.0

Table 7. Ablation study of the dilation choice of convolution layers in the working memory fusion module. The conv dilation parameter corresponds to the BEV spatial dimension (H, W).

ificance of our novel design of the temporal overlap heatmap.

Furthermore, we study the effect of varying the working memory capacity T_{WM} . As shown in Tab. 6, we find that $T_{WM} = 4$ is a sweet spot that delivers the best performance while also ensuring good runtime performance. We conjecture that the performance drop from further increasing T_{WM} may be due to the unnecessary dilution of the current BEV feature when fusing with additional temporal information.

Additionally, we validate the design choice of dilation parameters of the convolution layers in the proposed working memory fusion module. Empirical results in Tab. 7 demonstrate that maintaining a symmetric dilation of (2, 2) brings the most improvement.

4.6. Qualitative Results

We present qualitative comparison of MemFusionMap and StreamMapNet on nuScenes [1] at the extended 100 × 50 m range, which is the most challenging configuration. We present several examples of MemFusionMap outperforming baseline model StreamMapNet [39] in Fig. 6. Thanks to the proposed working memory fusion module

with the temporal overlap heatmap, MemFusionMap predicts more accurate map elements with better geometry and fewer missed elements. We show more qualitative sequential examples in the supplementary material.

5. Conclusion

In this paper, we have proposed MemFusionMap, a novel approach for effective online vectorized HD map construction with enhanced temporal reasoning capability. We propose a working memory fusion module to improve the memory capacity of the model. We also design a temporal overlap heatmap to aid the temporal reasoning process across history frames. Our proposed MemFusionMap demonstrates state-of-the-art performance on open-source benchmarks with the geofenced splits and maintains a fast runtime and versatile design. We hope MemFusionMap will inspire future research works to leverage our proposed working memory fusion module in other BEV perception tasks. We think integrating MemFusionMap in the MapTracker [4] framework to leverage the tracking paradigm to further enhance the performance and temporal consistency is also worth exploring.

6. Acknowledgement

We deeply appreciate the fruitful discussion with Ziqi Pang. We also thank Xinyu Xie for the insights on leveraging working memory for enhancing temporal reasoning capabilities.

References

- [1] Holger Caesar, Varun Bankiti, Alex H Lang, Sourabh Vora, Venice Erin Liong, Qiang Xu, Anush Krishnan, Yu Pan, Giancarlo Baldan, and Oscar Beijbom. nuscenes: A multi-modal dataset for autonomous driving. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 11621–11631, 2020. [2](#), [5](#), [6](#), [8](#)
- [2] Nicolas Carion, Francisco Massa, Gabriel Synnaeve, Nicolas Usunier, Alexander Kirillov, and Sergey Zagoruyko. End-to-end object detection with transformers. In *European Conference on Computer Vision (ECCV)*, pages 213–229, 2020. [2](#), [3](#)
- [3] Jiacheng Chen, Ruizhi Deng, and Yasutaka Furukawa. Polydiffuse: Polygonal shape reconstruction via guided set diffusion models. *Advances in Neural Information Processing Systems (NIPS)*, 36, 2024. [2](#)
- [4] Jiacheng Chen, Yuefan Wu, Jiaqi Tan, Hang Ma, and Yasutaka Furukawa. Maptracker: Tracking with strided memory fusion for consistent vector hd mapping. In *European Conference on Computer Vision (ECCV)*, pages 90–107, 2024. [1](#), [2](#), [3](#), [6](#), [7](#), [8](#)
- [5] Li Chen, Chonghao Sima, Yang Li, Zehan Zheng, Jiajie Xu, Xiangwei Geng, Hongyang Li, Conghui He, Jianping Shi, Yu Qiao, et al. Persformer: 3d lane detection via perspective transformer and the openlane benchmark. In *European Conference on Computer Vision (ECCV)*, pages 550–567, 2022. [2](#)
- [6] Liang-Chieh Chen, George Papandreou, Iasonas Kokkinos, Kevin Murphy, and Alan L Yuille. Deeplab: Semantic image segmentation with deep convolutional nets, atrous convolution, and fully connected crfs. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 40(4):834–848, 2017. [5](#)
- [7] Ho Kei Cheng and Alexander G Schwing. Xmem: Long-term video object segmentation with an atkinson-shiffrin memory model. In *European Conference on Computer Vision (ECCV)*, pages 640–658, 2022. [2](#), [4](#)
- [8] Junyoung Chung, Caglar Gulcehre, KyungHyun Cho, and Yoshua Bengio. Empirical evaluation of gated recurrent neural networks on sequence modeling. *arXiv preprint arXiv:1412.3555*, 2014. [2](#)
- [9] Nelson Cowan. The magical mystery four: How is working memory capacity limited, and why? *Current Directions in Psychological Science*, 19(1):51–57, 2010. [4](#)
- [10] Wenjie Ding, Limeng Qiao, Xi Qiu, and Chi Zhang. Pivotnet: Vectorized pivot learning for end-to-end hd map construction. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 3672–3682, 2023. [2](#)
- [11] Jason M Doherty, Clement Belletier, Stephen Rhodes, Agnieszka Jaroslawska, Pierre Barrouillet, Valerie Camos, Nelson Cowan, Moshe Naveh-Benjamin, and Robert H Logie. Dual-task costs in working memory: An adversarial collaboration. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 45(9):1529, 2019. [4](#)
- [12] Zhengyang Feng, Shaohua Guo, Xin Tan, Ke Xu, Min Wang, and Lizhuang Ma. Rethinking efficient lane detection via curve modeling. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 17062–17070, 2022. [2](#)
- [13] Xunjiang Gu, Guanyu Song, Igor Gilitschenski, Marco Pavone, and Boris Ivanovic. Producing and leveraging online map uncertainty in trajectory prediction. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 14521–14530, 2024. [1](#)
- [14] Chunrui Han, Jinrong Yang, Jianjian Sun, Zheng Ge, Runpei Dong, Hongyu Zhou, Weixin Mao, Yuang Peng, and Xiangyu Zhang. Exploring recurrent long-term temporal fusion for multi-view 3d perception. *IEEE Robotics and Automation Letters*, 2024. [2](#), [6](#)
- [15] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2016. [3](#), [6](#)
- [16] Feng Li, Hao Zhang, Shilong Liu, Jian Guo, Lionel M. Ni, and Lei Zhang. Dn-detr: Accelerate detr training by introducing query denoising. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 13619–13627, 2022. [2](#)
- [17] Qi Li, Yue Wang, Yilun Wang, and Hang Zhao. Hdmapnet: An online hd map construction and evaluation framework. In *2022 International Conference on Robotics and Automation (ICRA)*, pages 4628–4634, 2022. [1](#), [2](#)
- [18] Zhiqi Li, Wenhai Wang, Hongyang Li, Enze Xie, Chonghao Sima, Tong Lu, Yu Qiao, and Jifeng Dai. Bevformer: Learning bird’s-eye-view representation from multi-camera images via spatiotemporal transformers. In *European Conference on Computer Vision (ECCV)*, pages 1–18, 2022. [1](#), [2](#), [6](#)
- [19] Bencheng Liao, Shaoyu Chen, Xinggang Wang, Tianheng Cheng, Qian Zhang, Wenyu Liu, and Chang Huang. MapTR: Structured modeling and learning for online vectorized HD map construction. In *The Eleventh International Conference on Learning Representations (ICLR)*, 2023. [1](#), [2](#), [6](#), [7](#)
- [20] Bencheng Liao, Shaoyu Chen, Yunchi Zhang, Bo Jiang, Qian Zhang, Wenyu Liu, Chang Huang, and Xinggang Wang. Maptrv2: An end-to-end framework for online vectorized hd map construction. *International Journal of Computer Vision (IJCV)*, pages 1–23, 2024. [2](#)
- [21] Xiaolu Liu, Song Wang, Wentong Li, Ruizi Yang, Junbo Chen, and Jianke Zhu. Mgmap: Mask-guided learning for online vectorized hd map construction. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 14812–14821, 2024. [2](#)
- [22] Yicheng Liu, Tianyuan Yuan, Yue Wang, Yilun Wang, and Hang Zhao. Vectormapnet: End-to-end vectorized hd map learning. In *International Conference on Machine Learning (ICML)*, pages 22352–22369, 2023. [1](#), [2](#), [6](#), [7](#)
- [23] Zhijian Liu, Haotian Tang, Alexander Amini, Xinyu Yang, Huizi Mao, Daniela L Rus, and Song Han. Bevfusion: Multi-task multi-sensor fusion with unified bird’s-eye view representation. In *2023 IEEE International Conference on Robotics and Automation (ICRA)*, pages 2774–2781, 2023. [1](#)

- [24] Zihao Liu, Xiaoyu Zhang, Guangwei Liu, Ji Zhao, and Ningyi Xu. Leveraging enhanced queries of point sets for vectorized map construction. *arXiv preprint arXiv:2402.17430*, 2024. 1, 2
- [25] Ilya Loshchilov, Frank Hutter, et al. Fixing weight decay regularization in adam. *arXiv preprint arXiv:1711.05101*, 5, 2017. 6
- [26] Jinhyung Park, Chenfeng Xu, Shijia Yang, Kurt Keutzer, Kris M Kitani, Masayoshi Tomizuka, and Wei Zhan. Time will tell: New outlooks and a baseline for temporal multi-view 3d object detection. In *The Eleventh International Conference on Learning Representations (ICLR)*, 2022. 2, 6
- [27] Jonah Philion and Sanja Fidler. Lift, splat, shoot: Encoding images from arbitrary camera rigs by implicitly unprojecting to 3d. In *European Conference on Computer Vision (ECCV)*, pages 194–210, 2020. 1, 2
- [28] Limeng Qiao, Wenjie Ding, Xi Qiu, and Chi Zhang. End-to-end vectorized hd-map construction with piecewise bezier curve. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 13218–13228, 2023. 1, 2
- [29] Thomas Roddick and Roberto Cipolla. Predicting semantic map representations from images using pyramid occupancy networks. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2020. 5, 6
- [30] Tixiao Shan and Brendan Englot. Lego-loam: Lightweight and ground-optimized lidar odometry and mapping on variable terrain. In *2018 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 4758–4765, 2018. 2
- [31] Jingyu Song, Lingjun Zhao, and Katherine A Skinner. Lirafusion: Deep adaptive lidar-radar fusion for 3d object detection. In *2024 IEEE International Conference on Robotics and Automation (ICRA)*, pages 18250–18257, 2024. 1
- [32] Ruihao Wang, Jian Qin, Kaiying Li, Yaochen Li, Dong Cao, and Jintao Xu. Bev-lanedet: An efficient 3d lane detection based on virtual camera via key-points. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 1002–1011, 2023. 2
- [33] Shuo Wang, Fan Jia, Yingfei Liu, Yucheng Zhao, Zehui Chen, Tiancai Wang, Chi Zhang, Xiangyu Zhang, and Feng Zhao. Stream query denoising for vectorized hd map construction. *arXiv preprint arXiv:2401.09112*, 2024. 1, 2, 3, 6, 7
- [34] Shihao Wang, Yingfei Liu, Tiancai Wang, Ying Li, and Xiangyu Zhang. Exploring object-centric temporal modeling for efficient multi-view 3d object detection. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 3621–3631, 2023. 2, 6
- [35] Benjamin Wilson, William Qi, Tanmay Agarwal, John Lambert, Jagjeet Singh, Siddhesh Khandelwal, Bowen Pan, Ratnesh Kumar, Andrew Hartnett, Jhony Kaesemodel Pontes, et al. Argoverse 2: Next generation datasets for self-driving perception and forecasting. *arXiv preprint arXiv:2301.00493*, 2023. 2, 5, 6
- [36] Joey Wilson, Jingyu Song, Yuewei Fu, Arthur Zhang, Andrew Capodiceci, Paramsothy Jayakumar, Kira Barton, and Maani Ghaffari. Motionsc: Data set and network for real-time semantic mapping in dynamic environments. *IEEE Robotics and Automation Letters*, 7(3):8439–8446, 2022. 2
- [37] Xinyu Xie and Jarrod Moss. Working memory capacity predicts more effective problem space exploration. In *Proceedings of the Annual Meeting of the Cognitive Science Society*, volume 44, 2022. 4
- [38] Zhenhua Xu, Kenneth KY Wong, and Hengshuang Zhao. Insightmapper: A closer look at inner-instance information for vectorized high-definition mapping. *arXiv preprint arXiv:2308.08543*, 2023. 2
- [39] Tianyuan Yuan, Yicheng Liu, Yue Wang, Yilun Wang, and Hang Zhao. Streammapnet: Streaming mapping network for vectorized online hd map construction. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision (WACV)*, pages 7356–7365, 2024. 1, 2, 3, 5, 6, 7, 8
- [40] Ji Zhang, Sanjiv Singh, et al. Loam: Lidar odometry and mapping in real-time. In *Robotics: Science and systems*, volume 2, pages 1–9, 2014. 2
- [41] Lingjun Zhao, Jingyu Song, and Katherine A Skinner. Crkd: Enhanced camera-radar object detection with cross-modality knowledge distillation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 15470–15480, 2024. 1
- [42] Brady Zhou and Philipp Krähenbühl. Cross-view transformers for real-time map-view semantic segmentation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 13760–13769, 2022. 2
- [43] Junbao Zhou, Ziqi Pang, and Yu-Xiong Wang. Rmem: Restricted memory banks improve video object segmentation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 18602–18611, 2024. 2