# Timely Fusion of Surround Radar/Lidar for Object Detection in Autonomous Driving Systems

Wenjing Xie

Tao Hu

City University of Hong Kong City University of Hong Kong wenjing.xie@my.cityu.edu.hk taohu9-c@my.cityu.edu.hk

Chun Jason Xue Mohamed bin Zayed University of Artificial Intelligence jason.xue@mbzuai.ac.ae

Abstract—Fusing Radar and Lidar sensor data can fully utilize their complementary advantages and provide more accurate reconstruction of the surrounding for autonomous driving systems. Surround Radar/Lidar can provide  $360^{\circ}$  view sampling with the minimal cost, which are promising sensing hardware solutions for autonomous driving systems. However, due to the intrinsic physical constraints, the rotating speed of surround Radar, and thus the frequency to generate Radar data frames, is much lower than surround Lidar. Existing Radar/Lidar fusion methods have to work at the low frequency of surround Radar, which cannot meet the high responsiveness requirement of autonomous driving systems. This paper develops techniques to fuse surround Radar/Lidar with working frequency only limited by the faster surround Lidar instead of the slower surround Radar, based on widely-used object detection model called MVDNet. The basic idea of our approach is simple: we let MVDNet work with temporally unaligned data from Radar/Lidar, so that fusion can take place at any time when a new Lidar data frame arrives, instead of waiting for the slow Radar data frame. However, directly applying MVDNet to temporally unaligned Radar/Lidar data greatly degrades its object detection accuracy. The key information revealed in this paper is that we can achieve high output frequency with little accuracy loss by enhancing the training procedure to explore the temporal redundancy in MVDNet so that it can tolerate the temporal unalignment of input data. We explore several different ways of training enhancement and compare them quantitatively with experiments.

# I. INTRODUCTION

Accurate perception of the surroundings is a fundamental requirement for autonomous driving systems (ADS). Among the various sensors used in ADS, Lidar and Radar have emerged as highly desirable options for perception tasks [1]. While Lidar provides detailed point clouds in favorable weather conditions, it struggles in adverse weather (e.g., fog, snow), whereas Radar is less affected by adverse conditions but offers lower precision [2]. To overcome these limitations, fusing Lidar and Radar data has shown promise in enhancing object detection accuracy.

To obtain a comprehensive view of the surrounding environment, it is necessary for both Lidar and Radar sensors to cover the entire  $360^{\circ}$  angle. One traditional approach is to mount multiple sensors with different orientations, each covering a specific angle. The data from these sensors can then be combined to construct a  $360^{\circ}$  view. However, this approach comes with the drawback of high overall costs due Neiwen Ling Yale University 7 neiwen.ling@yale.edu

Guoliang Xing The Chinese University of Hong Kong glxing@cuhk.edu.hk

Nan Guan City University of Hong Kong nanguan@cityu.edu.hk

to the requirement of multiple sensors and their associated hardware. Alternatively, a more cost-effective solution is to utilize surround Radar/Lidar systems to capture the full 360° view. The temporal resolution of surround Radar/Lidar, i.e., the frequency at which they produce full-view data frames, is limited by the rotating speed. Radar, which operates using millimeter-waves, travels much slower than Lidar, which uses light, resulting in lower rotating speeds for Radar under similar operating conditions. For example, a popular surround Lidar sensor, Velodyne HDL-32E, rotates at a frequency of 20Hz (20 cycles per second) [3], while the state-of-the-art surround Radar sensor, NavTech CTS350-X, rotates at a frequency of 4Hz (4 cycles per second) [4]. Consequently, it is necessary to address the inconsistent input frequencies between surround Radar and Lidar when fusing their data. One straightforward solution is to down-sample the Lidar data to match the Radar frequency. Recent research has also explored combining consecutive Lidar data frames to produce artificial Lidar frames that align with the low frequency of Radar [2], [5]. In both cases, fusion is performed at the frequency dictated by the Radar rotation. For instance, with a surround Lidar rotating at 20Hz and a surround Radar rotating at 4Hz, fusion is performed and thus generates full-view reconstruction 4 times per second. Typically, ADS requires to capture events and take proper reaction in a short time, e.g., 100ms. Therefore, the low frequency of surround Radar/Lidar fusion makes it unsuitable for ADS from the realtime performance perspective.

In this paper, we propose techniques to address the aforementioned problem of low fusion frequency in Radar/Lidar fusion. Our approach aims to increase the fusion frequency beyond the limitations imposed by the low rotation speed of surround Radar. Our work is based on MVDNet, a widely-used Radar/Lidar data fusion deep neural network model in object detection. The core idea of our approach is to increase the fusion frequency by fusion the latest available Radar data frame with a new Lidar frame. However, this introduces temporal inconsistency between the Radar and Lidar sensor data, which can significantly degrade detection accuracy. To address this, we explore different methods to enhance the training procedure and evaluate their effectiveness through quantitative experiments.

979-8-3503-8795-7/24/\$31.00 ©2024 IEEE

DOI 10.1109/RTCSA62462.2024.00014 Authorized licensed use limited to: Yale University Library. Downloaded on November 15,2024 at 00:10:14 UTC from IEEE Xplore. Restrictions apply.

## II. RELATED WORK

Many studies have focused on fusing different sensor modalities, such as cameras, Lidars, and Radars, to improve 3D perception in autonomous driving. MVDNet [2] employs an attention mechanism to fuse Radar and Lidar data for vehicle detection in foggy weather and achieves state-of-the-art results on the Oxford Radar RobotCar (ORR) dataset [6]. Other works, such as ST-MVDNet [5], propose methods to handle missing sensor data in multi-modal vehicle detection. Realtime Radar/Lidar fusion methods for road-object detection and tracking in ADS have been proposed by Farag et al. [7]. LiRaNet [8] and EZFfusion [9] focus on end-to-end trajectory prediction and multi-modal 3D object detection and tracking, respectively. DEF [10] and FUTR3D [11] introduce fusion frameworks applicable to various sensor configurations. Bi-LRFusion [12] enhances 3D detection for dynamic objects through bi-directional LiDAR-Radar fusion. LiRaFusion [13] proposes early and middle fusion modules for joint voxel feature encoding and adaptive feature map fusion. However, existing Radar and Lidar fusion models assume perfect data synchronization. Down-sampling the faster Lidar data to match the Radar frequency is a common approach [2], [5], [7]–[11]. Some researchers have explored real-time infrastructure-vehicle cooperative approaches for timely fusion and perception [14]-[16]. In this paper, we propose techniques to fuse temporally unaligned Radar/Lidar data and achieve a high output frequency with minimal accuracy loss.

# **III. PRELIMINARY**

We build our work based on MVDNet [2], first object detection network that fuses Lidar and high-resolution 360° radar signals. MVDNet fuses Radar intensity maps with Lidar point clouds, which harnesses their complementary capabilities. The architecture of MVDNet is shown in Fig. 1. MVDNet consists of two stages. The region proposal network (MVD-RPN) extracts feature maps from Lidar and Radar inputs and generates proposals. The region fusion network (MVD-RFN) pools and fuses region-wise features of the two sensors' data frames and outputs oriented bounding boxes of detected vehicles. MVDNet assumes that the input Radar and Lidar



Fig. 1. The architecture of MVDNet [2].

frames have same timestamps. However, in the training data of MVDNet, the frequency of raw Radar frames, denoted by  $F_r$ , is different from the frequency of raw Lidar frames, denoted by  $F_l$ . MVDNet solves this problem by combining several consecutive raw Lidar frames into an artificial concatenated Lidar frame, which is generated with the same frequency as



Fig. 2. The Radar/Lidar frames alignment policy in MVDNet. Radar data frames, as shown in Fig. 2. In this way, the fusion frequency of MVDNet equals  $F_r$ . We define

$$ratio = \left\lfloor \frac{F_l}{F_r} \right\rfloor$$

and ratio+1+1 raw Lidar frames are aggregated to produce one concatenated Lidar data frame. More precisely, a subset of cloud points that covers 120° view in each Lidar frame is selected, and these subsets are combined into a concatenated Lidar frame as shown in Fig. 2. Each concatenated Lidar frame and corresponding Radar frame are paired and sent to MVDNet. To improve accuracy, when fusing the latest paired Radar/Lidar frames, MVDNet uses historical frames. The number of historical frames is denoted by num\_history. In MVDNet, num\_history=4 by default. Decreasing num history can reduce the computation workload of MVDNet inference, but at the cost of accuracy loss, as discussed in IV-E and IV-B.

#### IV. THE PROPOSED METHOD

# A. Dataset Preparation

The original ORR data are collected by a vehicle equipped with a NavTech CTS350-X Radar [4] at the roof center, colocated with two Velodyne HDL-32E [3] Lidars whose point clouds are combined. MVDNet [2] manually generates the ground-truth labels based on the ORR Lidar point clouds. They create 3D bounding boxes for vehicles in one out of every 20 frames (i.e., 1s) using the Scalable annotation tool. Labels of the remaining 19 frames are interpolated using the visual odometry data provided in ORR and manually adjusted to align with the corresponding vehicles. To demonstrate the effectiveness of our proposed timely fusion methods, we use the same processed dataset as MVDNet.

# B. Increasing the Radar/Lidar Fusion Frequency

Our method enables fusion to be executed immediately upon the arrival of a raw Lidar frame, instead of waiting for the slow Radar frames, so fusion can be triggered with a higher frequency. The fusion frequency of our method, denoted as  $F_f$ , is different from existing methods [2], [5], [7]-[11], [14]-[16], where the  $F_f$  is equal to the Radar frequency  $F_r$ . In our method, the fusion frequency  $F_f$  can be increased within the range:

$$F_r \le F_f \le F_l \tag{1}$$



Fig. 3. Our proposed fusion method.

Each arrival of a new raw Lidar frame provides an opportunity for fusion. In practice, fusion is feasible whenever the fusion period  $(1/F_f)$  is an integral number multiple of the Lidar frame period  $(1/F_l)$ . This establishes a direct relationship between the fusion frequency and the Lidar frequency, represented by the parameter  $\alpha$ :

$$\alpha = \frac{F_l}{F_f} \tag{2}$$

Where  $\alpha$  should satisfy the condition:

$$\alpha \in \{1, 2, 3, ..., ratio\}$$
(3)

If  $\alpha > ratio$ , the fusion of Radar and Lidar frames becomes meaningless due to the absence of temporal overlap. The fusion frequency in our method is limited only by the faster Lidar instead of the slower Radar.

Fig. 3 shows two examples to demonstrate our fusion method. When a new raw Lidar frame arrives and the matching raw Radar frame for corresponding interval has not yet arrived, our method utilizes the latest available raw Radar frame for fusion. In Fig. 3(a), we aim to achieve a high fusion frequency  $F_f = F_l$ . According to Equation (2), we need to set  $\alpha = 1$ , indicating performing timely fusion when each raw Lidar frame arrives. Consider that the first fusion event occurs at time  $t_7^l$ , when the second and third new raw Lidar frames arrive at time  $t_8^l$  and  $t_9^l$  respectively, the Radar frame which is perfectly synchronized in time has not yet arrived and is expected to arrive at time  $t_2^r$ . Under these situations, fusion is conducted between the latest concatenated Lidar frames (the orange box) and the latest available raw Radar frame (the orange box) generated at  $t_1^r$  respectively. Fig. 3(b) shows another example with fusion frequency  $F_f = F_l/3$ . The fusion should be performed every three raw Lidar arrives. Our method not only increases the fusion frequency but also allows us to adjust the frequency according to the actual situation. In practice, the fusion results do not always have to be generated at the highest frequency. Based on the current vehicle speed and overall workload of the vehicle, the ADS is expected to provide a fusion frequency that can be adjusted.

# C. Training Enhancement to Maintain Accuracy

However, when directly applying the mentioned timely fusion method, there can be a temporal misalignment between the input Radar/Lidar frames to the MVDNet. This misalignment often leads to a significant decrease in object detection accuracy. In fact, each pair of unaligned Radar/Lidar frames exhibits a time offset, as illustrated in Fig. 4. We use the term *offset* to represent the number of complete raw Lidar frames that arrive ahead of the Radar in each paired Radar/Lidar frame set. The *offset* must satisfy the condition:

$$offset \in \{0, 1, 2, ..., ratio\}$$
 (4)

A larger *offset* value indicates a longer unaligned time interval between the paired Radar/Lidar frames. Consequently, if the *offset* exceeds the *ratio*, the fusion between the paired frames becomes meaningless because they do not overlap in time.

Following the MVDNet paper, we evaluate the original MVDNet model on unaligned Radar/Lidar frames using the ORR dataset ( $F_l = 20$ Hz,  $F_r = 4$ Hz, *ratio*=5). To ensure a fair comparison, we utilize average precision (AP) in COCO evaluation [17] with Intersection-over-Union (IoU) thresholds of 0.5, 0.65, and 0.8. The accuracy results for various offset settings are shown in Table I. Our observations are as follows:

- The accuracy of the MVDnet model significantly decreases when applied to unaligned frames (*offset* =  $\{1, 2, 3, 4, 5\}$ ) compared to aligned frames (*offset*=0) across different IoU thresholds. This indicates that using the original MVDNet model on unaligned Radar/Lidar frames significantly degrades object detection accuracy.
- Specifically, when considering a specific IoU value, such as IoU=0.8, we observe a further decrease in object detection accuracy as the *offset* increases. This highlights the critical impact of the *offset* on the accuracy of MVDNet.

TABLE I Overall Performance of Directly Applying MVDNet on Radar/Lidar Frames: AP of Oriented Bounding Boxes in Bird's Eye View (BEV).

Model	offset	0				1			2			3			4		5		
	IoU	0.5	0.65	0.8	0.5	0.65	0.8	0.5	0.65	0.8	0.5	0.65	0.8	0.5	0.65	0.8	0.5	0.65	0.8
MVI	DNet	0.897	0.877	0.735	0.652	0.619	0.486	0.616	0.587	0.462	0.583	0.559	0.441	0.51	0.498	0.411	0.521	0.508	0.407

TABLE II

OVERALL PERFORMANCE OF OUR FUSION METHOD ON RADAR/LIDAR FRAMES: AP OF ORIENTED BOUNDING BOXES IN BEV.

Model	offset	1				2			3			4		5		
	IoU	0.5	0.65	0.8	0.5	0.65	0.8	0.5	0.65	0.8	0.5	0.65	0.8	0.5	0.65	0.8
MVD	Net_1	0.898	0.869	0.722	0.867	0.841	0.668	0.847	0.827	0.652	0.828	0.809	0.643	0.826	0.801	0.644
MVD	Net_2	0.87	0.85	0.67	0.899	0.877	0.724	0.889	0.863	0.684	0.876	0.848	0.705	0.868	0.847	0.702
MVD	Net_3	0.867	0.846	0.699	0.868	0.847	0.701	0.89	0.867	0.722	0.879	0.858	0.709	0.878	0.849	0.7
MVD	Net_4	0.877	0.855	0.686	0.881	0.853	0.683	0.874	0.853	0.688	0.888	0.86	0.709	0.875	0.854	0.688
MVD	Net_5	0.855	0.833	0.679	0.862	0.834	0.687	0.865	0.843	0.689	0.884	0.864	0.704	0.885	0.865	0.707



To address the issue of accuracy loss issue mentioned above, we leverage an important characteristic of MVDNet that incorporates both the current pair and a specific number of historical pairs (num history=4) during fusion. This results in a total of 1+num\_history paired Radar/Lidar frames being used as input for each fusion. It is important to note that there exists redundant information among consecutive Radar/Lidar frames. We exploit the redundant information to make MVDNet still work well with temporally unaligned Radar/Lidar frames. Furthermore, we discover a straightforward yet highly effective training enhancement method that enables MVDNet to maintain accuracy during the fusion of temporally unaligned Radar/Lidar frames. This method involves synthesizing paired Radar/Lidar frames with corresponding offset values and utilizing them to train the MVDNet from scratch. By incorporating data with different offset values during training, we aim to equip the MVDNet model with the ability to detect objects from

temporally unaligned Radar/Lidar frames. Our methods' results, obtained under the same experimental setting, are summarized in Table II. MVDNet\_offset refers to the MVDNet model trained on Radar/Lidar frames with a specific offset, where offset  $\{1, 2, 3, 4, 5\}$ . Our observations are as follows:

- In table II, for unaligned Radar/Lidar frames with a specific offset=i, the MVDNet\_i model achieves better accuracy performance than other MVDNet\_j, where i, j ∈ {1, 2, 3, 4, 5} and i ≠ j.
- Comparing Table II and Table I, the accuracy of our enhancement training method on unaligned Radar/Lidar frames with each specific *offset* (*offset*  $\in \{1, 2, 3, 4, 5\}$ ) is almost equal to the original MVDNet model on aligned Radar/Lidar frames (*offset*=0).

The results demonstrate that our proposed fusion method maintains accuracy even when the Radar/Lidar frames are temporally unaligned, while also increasing the fusion frequency.

### D. Exploiting Different Enhancement Strategies

In order to enable MVDNet model to deal with Radar/Lidar frames with various *offset*, we exploit two different enhancement strategies.

Separate training strategy: For temporally unaligned Radar/Lidar frames with different offset, the first strategy involves training a dedicated MVDNet model for each specific offset frame. During training, we first train a shared MVDNet model using paired Radar/Lidar frames with offset=[ratio/2]. The parameter weights of its MVD-RPN are then shared among other MVDNet models that handle different offset frames. Next, the MVD-RFN of each MVDNet model is fine-tuned using the corresponding offset frames. This process results in multiple MVDNet models, with each model designed to handle Radar/Lidar frames with a specific offset. The evaluation results on the ORR dataset are summarized in Table III. MVDNet\_separated\_offset refers to an MVDNet model finetuned on Radar/Lidar frames with a specific offset, where offset  $\in \{0, 1, 2, \dots, 5\}$ . We observe that for Radar/Lidar frames with a specific offset, each MVDNet separated offset model, fine-tuned on the corresponding offset frames, achieves the best accuracy performance compared to other models.

**Mixed training strategy.** Another strategy is to mix an equal amount of paired Radar/Lidar frames with various *offset* together. During training, we use the newly mixed training set to train a MVDNet\_*mixed* model from scratch. The goal is to obtain a MVDNet\_*mixed* model capable of handling Radar/Lidar frames with different *offset*. The evaluation results are summarized in Table IV. We observe that the MVD-Net\_*mixed* model achieves similar and good accuracy performance with different IoU settings, indicating its ability to handle Radar/Lidar frames with different *offset*.

Comparing the results of the two strategies, we observe two phenomena. Firstly, for aligned Radar/Lidar frames (*off-set=0*), the separately fine-tuned MVDNet\_separate\_0 model achieves slightly better accuracy than the MVDNet\_*mixed* model. Secondly, for other unaligned Radar/Lidar frames (*off-set* $\in$ {1,2,3,4,5}), the MVDNet\_*mixed* model achieves better accuracy than separately fine-tuned MVDNet\_separated\_1/2/3/4/5 models. Therefore, designing a unified MVDNet model capable of handling diverse Radar/Lidar frames involves a trade-off.

TABLE III
OVERALL PERFORMANCE OF SEPARATELY TRAINED MVDNET MODELS : AP OF ORIENTED BOUNDING BOXES IN BEV

Model	offset		0			1			2			3			4			5		
	IoU	0.5	0.65	0.8	0.5	0.65	0.8	0.5	0.65	0.8	0.5	0.65	0.8	0.5	0.65	0.8	0.5	0.65	0.8	
MVDNet_separated_0		0.899	0.878	0.736	0.848	0.826	0.679	0.838	0.817	0.675	0.83	0.809	0.68	0.802	0.79	0.671	0.803	0.792	0.663	
MVDNet_separated_1		0.887	0.866	0.719	0.887	0.866	0.722	0.884	0.864	0.722	0.884	0.856	0.722	0.876	0.855	0.711	0.874	0.847	0.711	
MVDNet_separated_2		0.867	0.846	0.708	0.876	0.856	0.656	0.886	0.865	0.722	0.883	0.856	0.722	0.874	0.855	0.718	0.867	0.854	0.712	
MVDNet_separated_3		0.861	0.84	0.689	0.868	0.848	0.699	0.877	0.856	0.71	0.888	0.868	0.723	0.88	0.859	0.714	0.879	0.859	0.713	
MVDNet_separated_4		0.841	0.82	0.682	0.848	0.828	0.694	0.858	0.838	0.703	0.878	0.859	0.725	0.887	0.868	0.725	0.879	0.86	0.723	
MVDNet_separated_5 0.869 0.848 0.		0.702	0.867	0.847	0.699	0.875	0.855	0.701	0.885	0.865	0.711	0.887	0.867	0.711	0.887	0.867	0.712			

 TABLE IV

 Overall Performance of The Mixed Training MVDNet Model: AP of Oriented Bounding Boxes in BEV.

Model	offset		0		1			2				3			4		5		
WIGGET	IoU	0.5	0.65	0.8	0.5	0.65	0.8	0.5	0.65	0.8	0.5	0.65	0.8	0.5	0.65	0.8	0.5	0.65	0.8
MVDNe	et_mixed	0.897	0.877	0.735	0.896	0.876	0.734	0.896	0.877	0.734	0.895	0.875	0.733	0.895	0.876	0.727	0.895	0.875	0.726

OVERALL PERFORMANCE OF FRAME SKIPPING METHOD: AP OF ORIENTED BOUNDING BOXES IN BEV.

Strategy	offset		0			1			2			3			4		5		
	IoU	0.5	0.65	0.8	0.5	0.65	0.8	0.5	0.65	0.8	0.5	0.65	0.8	0.5	0.65	0.8	0.5	0.65	0.8
Without frame skipping, num_history=4		0.897	0.877	0.735	0.898	0.869	0.722	0.899	0.877	0.724	0.89	0.867	0.722	0.888	0.86	0.709	0.885	0.865	0.707
Without frame skipping, num_history=2		0.871	0.85	0.725	0.883	0.863	0.72	0.885	0.866	0.721	0.885	0.864	0.718	0.869	0.856	0.712	0.878	0.866	0.723
With frame skipping, num_history=2		0.895	0.877	0.724	0.885	0.865	0.715	0.873	0.852	0.696	0.887	0.856	0.721	0.887	0.857	0.7	0.885	0.864	0.706

TABLE VI
OVERALL PERFORMANCE OF HISTORICAL FRAME ALIGNMENT STRATEGY: AP OF ORIENTED BOUNDING BOXES IN BEV.

Strategy	offset	1				2	_		3	_		4		5			
	IoU	0.5	0.65	0.8	0.5	0.65	0.8	0.5	0.65	0.8	0.5	0.65	0.8	0.5	0.65	0.8	
unaligned		0.898	0.869	0.722	0.899	0.877	0.724	0.89	0.867	0.722	0.888	0.86	0.709	0.885	0.865	0.707	
aligned		0.889	0.859	0.694	0.895	0.874	0.732	0.88	0.859	0.718	0.879	0.859	0.705	0.856	0.836	0.687	



Fig. 5. Structure of the multi-branch unified MVDNet.

Based on the separated training strategy, we propose a multibranch unified MVDNet architecture, as shown in Fig. 5. The architecture consists of three main modules. The first module consists of feature extractors designed for Radar/Lidar frames. As neural networks tend to extract low-level features in the lower layers [18], we leverage this property and utilize the low-level feature extractors as shared feature extractors for Radar/Lidar frames with different *offset*. The second module is the fusion module, which incorporates multiple branches. Each branch is dedicated to fusing Radar/Lidar frames with a specific *offset*. The third module introduces a new input *offset* to determine which branch should be chosen for inference.

## E. Impact of Historical Information

In order to navigate the trade-off between latency and accuracy performance, we exploit two strategies for selecting and matching historical frames, aiming to reduce inference latency.

Historical frame skipping strategy. The historical frame skipping strategy aims to reduce inference latency while maintaining acceptable accuracy. In the basic fusion method (Fig. 3), the current and historical paired Radar/Lidar frames are selected consecutively in time. For instance, with a num\_history=2 configuration (Fig. 6(a)), the current frame and two consecutive paired frames are chosen as input for one fusion. Recognizing the presence of redundant information across consecutive frames, we have explored a frame skipping method that allows us to represent the same amount of information using fewer frames compared to the basic fusion method. Fig. 6(b) illustrates our frame skipping method with the same num\_history=2 configuration. In this approach, historical paired Radar/Lidar frames are selected every two intervals, effectively skipping several frames. Even though the fusion process only utilize two historical frames, the overall amount of information represented remains equivalent to that of the *num\_history=4* configuration.

Table V shows the evaluation results. When considering *num\_history*=2 configuration and a specific *offset* (*off-set*  $\in$  {0,1,2,...,5}), our frame skipping method achieves comparable or better accuracy compared to the method without frame skipping, indicating the effectiveness of redundant data. Additionally, compared to the *num\_history*=4 without skipping configuration, our method *num\_history*=2 with skipping configuration only experiences a slight decrease in accuracy. The inference latency increases as the number of input frames used



Fig. 6. Historical frame selection strategy.

for fusion increases. Therefore, our method effectively reduces latency with minimal sacrifice in accuracy, offering a novel approach for achieving an accuracy-latency trade-off.



Historical frame alignment strategy. In the basic fusion method (Fig. 3), misalignment between the current paired Radar/Lidar frames results in misalignment for all other historical paired frames with the same offset. However, we observe that apart from the current frames, other historical concatenated Lidar frames can find their aligned Radar frames. As shown in Fig. 7, while the current concatenated Lidar frame at  $t_i^l$ can only be paired with the latest available raw Radar frame arriving at  $t_i^r$ , each historical concatenated Lidar frame can be aligned perfectly with a reconstructed Radar frame. This new reconstructed Radar frame is obtained by concatenating two adjacent original raw Radar frames. Based on this observation, we propose a historical frame alignment strategy to align all historical paired Radar/Lidar frames in time. The basic idea is, when selecting each historical paired frame, reconstructing a new Radar frame that aligns with the concatenated Lidar frame according to the Lidar timestamp.

Table VI compares our strategy with the basic method. We observe that our alignment strategy does not yield any accuracy improvement for temporally unaligned Radar/Lidar frames with a specific *offset*, where *offset*  $\in$  {1,2,...,5}.

#### V. CONCLUSION

In this paper, we propose techniques to fuse surround Radar/Lidar data by leveraging the faster sampling rate of Lidar instead of the slower Radar. Our approach is based on the widely used MVDNet model for Radar/Lidar fusionbased object detection. We aim to maximize the benefits of Lidar's fast sampling and achieve real-time fusion of temporally unaligned Radar/Lidar frames to provide timely road condition information. Experimental results demonstrate that our fusion method significantly increases fusion frequency of unaligned Radar/Lidar data with minimal impact on detection accuracy.

#### ACKNOWLEDGMENT

This work is partially supported by Hong Kong GRF under grant no. 15206221 and 11208522.

#### REFERENCES

- Q. Zhang, H. Sun, Z. Wei, and Z. Feng, "Sensing and communication integrated system for autonomous driving vehicles," in *IEEE INFOCOM* 2020-IEEE Conference on Computer Communications Workshops (IN-FOCOM WKSHPS). IEEE, 2020, pp. 1278–1279.
- [2] K. Qian, S. Zhu, X. Zhang, and L. E. Li, "Robust multimodal vehicle detection in foggy weather using complementary lidar and radar signals," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2021, pp. 444–453.
- [3] Velodyne hdl-32e lidar. [Online]. Available: https://velodynelidar.com/ products/hdl-32e/
- [4] Navtech cts350-x radar. [Online]. Available: https://navtechradar.com/?s= CTS350-X
- [5] Y.-J. Li, J. Park, M. O'Toole, and K. Kitani, "Modality-agnostic learning for radar-lidar fusion in vehicle detection," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2022, pp. 918–927.
- [6] D. Barnes, M. Gadd, P. Murcutt, P. Newman, and I. Posner, "The oxford radar robotcar dataset: A radar extension to the oxford robotcar dataset," in 2020 IEEE International Conference on Robotics and Automation (ICRA). IEEE, 2020, pp. 6433–6438.
- [7] W. Farag, "Real-time lidar and radar fusion for road-objects detection and tracking," *International Journal of Computational Science and Engineering*, vol. 24, no. 5, pp. 517–529, 2021.
- [8] M. Shah, Z. Huang, A. Laddha, M. Langford, B. Barber, S. Zhang, C. Vallespi-Gonzalez, and R. Urtasun, "Liranet: End-to-end trajectory prediction using spatio-temporal radar fusion," *arXiv preprint arXiv:2010.00731*, 2020.
- [9] Y. Li, J. Deng, Y. Zhang, J. Ji, H. Li, and Y. Zhang, "Ezfusion: A close look at the integration of lidar, millimeter-wave radar, and camera for accurate 3d object detection and tracking," *IEEE Robotics and Automation Letters*, vol. 7, no. 4, pp. 11182–11189, 2022.
- [10] M. Bijelic, T. Gruber, F. Mannan, F. Kraus, W. Ritter, K. Dietmayer, and F. Heide, "Seeing through fog without seeing fog: Deep multimodal sensor fusion in unseen adverse weather," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2020, pp. 11 682–11 692.
- [11] X. Chen, T. Zhang, Y. Wang, Y. Wang, and H. Zhao, "Futr3d: A unified sensor fusion framework for 3d detection," *arXiv preprint* arXiv:2203.10642, 2022.
- [12] Y. Wang, J. Deng, Y. Li, J. Hu, C. Liu, Y. Zhang, J. Ji, W. Ouyang, and Y. Zhang, "Bi-Irfusion: Bi-directional lidar-radar fusion for 3d dynamic object detection," in *Proceedings of the IEEE/CVF Conference* on Computer Vision and Pattern Recognition, 2023, pp. 13 394–13 403.
- [13] J. Song, L. Zhao, and K. A. Skinner, "Lirafusion: Deep adaptive lidarradar fusion for 3d object detection," arXiv preprint arXiv:2402.11735, 2024.
- [14] A. Prakash, K. Chitta, and A. Geiger, "Multi-modal fusion transformer for end-to-end autonomous driving," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2021, pp. 7077– 7087.
- [15] J. Cui, H. Qiu, D. Chen, P. Stone, and Y. Zhu, "Coopernaut: Endto-end driving with cooperative perception for networked vehicles," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2022, pp. 17252–17262.
- [16] Y. He, L. Ma, Z. Jiang, Y. Tang, and G. Xing, "Vi-eye: semantic-based 3d point cloud registration for infrastructure-assisted autonomous driving," in *Proceedings of the 27th Annual International Conference on Mobile Computing and Networking*, 2021, pp. 573–586.
- [17] T.-Y. Lin, M. Maire, S. Belongie, J. Hays, P. Perona, D. Ramanan, P. Dollár, and C. L. Zitnick, "Microsoft coco: Common objects in context," in *European conference on computer vision*. Springer, 2014, pp. 740–755.
- [18] A. Søgaard and Y. Goldberg, "Deep multi-task learning with low level tasks supervised at lower layers," in *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers)*, 2016, pp. 231–235.