

TempBEV: Improving Learned BEV Encoders with Combined Image and BEV Space Temporal Aggregation

Thomas Monninger^{1,2}, Vandana Dokkadi^{3,*}, Md Zafar Anwar^{4,*}, Steffen Staab^{2,5}

Abstract—Autonomous driving requires an accurate representation of the environment. A strategy toward high accuracy is to fuse data from several sensors. Learned Bird’s-Eye View (BEV) encoders can achieve this by mapping data from individual sensors into one joint latent space. For cost-efficient camera-only systems, this provides an effective mechanism to fuse data from multiple cameras with different views. Accuracy can further be improved by aggregating sensor information over time. This is especially important in monocular camera systems to account for the lack of explicit depth and velocity measurements, such that decision-critical information about distance distances and motions of other objects is easily accessible. Thereby, the effectiveness of developed BEV encoders crucially depends on the operators used to aggregate temporal information and on the used latent representation spaces.

We analyze BEV encoders proposed in the literature and compare their effectiveness, quantifying the effects of aggregation operators and latent representations. While most existing approaches aggregate temporal information *either* in image *or* in BEV latent space, our analyses and performance comparisons suggest that these latent representations exhibit complementary strengths. Therefore, we develop a novel temporal BEV encoder, TempBEV, which integrates aggregated temporal information from both latent spaces. We consider subsequent image frames as stereo through time and leverage methods from optical flow estimation for temporal stereo encoding.

Empirical evaluation on the NuScenes dataset shows a significant improvement by TempBEV over the baseline for 3D object detection and BEV segmentation. The ablation uncovers a strong synergy of joint temporal aggregation in the image and BEV latent space. These results indicate the overall effectiveness of our approach and make a strong case for aggregating temporal information in both image and BEV latent spaces.

I. INTRODUCTION

Autonomous driving requires a perception system that can derive an accurate representation of the environment, which includes both dynamic objects and static elements. Perception tasks include 3D object detection [1] (estimating bounding boxes for dynamic objects), Bird’s-Eye View (BEV) segmentation [2] (performing pixel-level semantic classification in 2D top-down view), motion tracking, etc. Autonomous vehicles depend on the outputs of these tasks to understand the structure of their environment and make informed decisions. While monocular multi-camera systems




Fig. 1: Exemplary traffic scene at three time steps represented in high-resolution Image vs. low-resolution BEV space. Same car is colored differently at different time steps only for better visibility.

offer a cost-effective solution, they provide only projected views and lack explicit depth information [3]. This poses challenges for tasks inherently requiring 3D perception.

One way of improving depth estimation is temporal aggregation, a process of combining sensor information over multiple time steps, for example through temporal stereo [4, 5]. Temporal aggregation also allows for mitigating occlusions, reducing measurement uncertainty, and estimating motions of other dynamic objects by considering several views from the same cameras at different times, reaching a comprehensive representation of the traffic scene [6].

Temporal aggregation of camera frames requires determining the corresponding features across subsequent frames. The camera pose changes due to the ego-motion of the vehicle [7], on which the cameras are mounted, and other dynamic objects [8] move in the 3D environment affecting their position in the different frames. Image space represents a 2D projection of the 3D environment, making linear motions of dynamic objects appear non-linear. Moreover, objects in image space might also change appearance over time, for example through changes in perspective, lighting, or occlusion, adding to the complexity of finding correspondences.

Most recent works [7-13] perform temporal aggregation in BEV space. BEV space is an orthogonal projection of the scene into a 2D top-down view, preserving a linear appearance of linear motions. Lifting is performed to transform the camera frames from projected image space into 3D space, followed by an orthogonal projection to BEV space. Thereby, BEV space representation reduces the complexity of finding correspondences. However, the uncertainty of depth estimation [3] as well as the lower resolution of BEV space [14] affect the estimate of object positions in BEV representation. This is illustrated in Fig. 1, which shows a traffic scene in image vs. BEV representation. Aggregating only in BEV space misses features present in high-resolution

¹Mercedes-Benz Research & Development North America, Sunnyvale, CA, USA (email: thomas.monninger@mercedes-benz.com)

²University of Stuttgart, Institute for Artificial Intelligence, Stuttgart, Germany (email: steffen.staab@ki.uni-stuttgart.de)

³University of Massachusetts Amherst, Amherst, MA, USA

⁴Pennsylvania State University, University Park, PA, USA

⁵University of Southampton, Southampton, United Kingdom

*Work was done during an internship at Mercedes-Benz Research & Development North America.

Table I: Survey of learning-based temporal BEV aggregation methods on 3D object detection (Det) and BEV segmentation (Seg).

Methods Publication	Operator	Aggregation			Feature Space		Motions Object-based	Learning Tasks		
		Recurrent	Parallel	# Steps	Image	BEV		Det	Seg	Other
BEVDepth [9]	Convolution		✓	2		✓		✓		
BEVDet4D [7]	Convolution		✓	2		✓		✓		
BEVFormer [10]	Deformable Attention	✓		4		✓		✓	✓	
BEVFormer v2 [11]	Convolution		✓	4		✓		✓		
BEVStereo [15]	Convolution		✓	2	✓	✓		✓		
BEVStitch [16]	Max pooling		✓	4	✓				✓	
BEVerse [12]	Convolution		✓	3		✓		✓	✓	✓
DORT [13]	Convolution	✓		16		✓	✓	✓		✓
DfM [5]	Convolution		✓	2	✓			✓		
DynamicBEV [17]	Attention	✓		8		✓	✓	✓		
FIERY [18]	Attention+Convolution		✓	3		✓			✓	✓
Fast-BEV [19, 20]	Convolution		✓	4		✓				
HoP [21]	Deformable Attention	✓	✓	4 / 2		✓		✓		
Img2Maps [22]	Axial Attention		✓	4		✓			✓	
MaGNet [23]	Convolution		✓	5	✓					✓
MVSNet [24]	Convolution		✓	3	✓					✓
OCBEV [25]	Deformable Attention	✓		4		✓		✓		
PETRv2 [26]	Attention		✓	2	✓			✓	✓	✓
PointBEV [27]	Sub-manifold Attention			8		✓			✓	
PolarDETR [28]	Attention		✓	2	✓			✓		✓
PolarFormer [29]	Attention		✓	2		✓		✓		
ST-P3 [30]	Attention+Convolution		✓	3		✓	✓		✓	✓
STS [31]	Group-wise correlation		✓	2	✓			✓		
SoloFusion [14]	Convolution		✓	17	✓			✓		
SparseBEV [32]	Attention		✓	8		✓	✓	✓		
StreamPETR [8]	Attention	✓		4		✓	✓	✓		✓
TBP-Former [33]	Deformable Attention		✓	3		✓			✓	✓
UVTR [34]	Convolution		✓	5		✓		✓		
UniFusion [35]	Deformable Attention		✓	7		✓			✓	
VideoBEV [36]	Convolution	✓		8		✓		✓	✓	✓

image space, indicating the complementary nature of the image and BEV space.

We argue that it is beneficial to consider both image and BEV space for temporal aggregation to combine the different representational strengths. To design a suitable model, we survey existing literature and perform a comparative study of different aggregation mechanisms including attention, convolution, and max pooling. With the insights gained from this study, we propose a novel **Temporal BEV** encoder, **TempBEV**, that effectively combines temporal aggregation in image and BEV space. In summary, the main contributions of this paper are:

- We provide a survey on temporal aggregation mechanisms for learned BEV encoders.
- We compare learned temporal aggregation operators through systematic experiments and investigate their effect on object detection and BEV segmentation tasks.
- We propose TempBEV to combine temporal aggregation in image and BEV space and investigate its effectiveness through experimental evaluation.
- We apply methods from optical flow for image space aggregation and show the benefit of transfer learning.

II. SURVEY ON TEMPORAL AGGREGATION

We survey temporal aggregation mechanisms in recent BEV models (cf. Table I). Models are distinguished based on the following five key aspects: aggregation operator, recurrent or parallel aggregation and number of steps, feature space, motions, and learning tasks.

A. Aggregation Operator

The aggregation operator defines the mathematical operation that is used to combine information from multiple time steps. Commonly used temporal aggregation operators are attention, convolution, and max pooling. Most approaches use the attention mechanism, owing to its effectiveness and expressiveness. More specifically, deformable self-attention [37] is frequently employed to address computational complexity for real-time applications. A simple alternative is max pooling, which requires no additional parameters (e.g., BEVStitch [16]). A few other approaches exist, such as STS [31], which employs group-wise correlations processed by Multi-Layer Perceptrons.

B. Recurrent or Parallel Aggregation

For a given time step t , let $U_t = \{u_t^1, \dots, u_t^n\}$ be the set of n inputs (e.g., corresponding to the images from n different cameras). A key choice is whether $U_{t-k:t}$ is aggregated in a recurrent or parallel manner to combine features into a latent state X . Recurrent aggregation conditions the current state X_t on X_{t-1} . This enables information flow over a longer time horizon, primarily limited by the capacity of the latent state. All works in this survey that uses recurrent aggregation (e.g., [10]) perform it in BEV space with X being a BEV grid. Mathematically, at time t , the latent state X_t can be expressed as a function $f_{\text{recurrent}}$ of the input U_t and the previous latent state X_{t-1} :

$$X_t = f_{\text{recurrent}}(U_t, X_{t-1}) \quad (1)$$

Alternatively, parallel aggregation directly combines input from a fixed number of time steps (*e.g.*, [11]). This approach can learn individual aggregations specific for each time step. However, computation grows linearly, limiting the number of time steps. This effectively constrains the available temporal information to a limited time horizon. Mathematically, X_t can be expressed as a function $f_{parallel}$ of the inputs from the last k time steps.

$$X_t = f_{parallel}(U_{t-k:t}) \quad (2)$$

The survey indicates that both parallel and recurrent aggregation techniques are prevalent, with no clear advantage of one over another. To the best of our knowledge, HoP [21] is the only model using both parallel and recurrent aggregation. Their setup performs recurrent aggregation of the BEV grids and also adds a parallel aggregation module to process two BEV grids directly. Furthermore, it is important to specify the number of time steps considered in each approach. For parallel aggregation, this refers to the number of aggregated time steps. For recurrent aggregation, it denotes the extent to which the recurrent computation graph unfolds during training, while the inference setting is not limited to a specific number of time steps. Several studies have investigated the impact of the number of time steps, yielding diverse conclusions. BEVFormer [10] employs 4 time steps (equivalent to 2s), VideoBEV [36] and StreamPETR [8] utilize 8 time steps (4s), and SOLOFusion [14] extends this further to 17 time steps (8.5s). It is worth noting that these studies do not adequately analyze the potential of long time horizons for detecting static elements in the scene. They primarily focus on object detection, where a quick performance saturation is expected due to the dynamic nature of objects.

C. Feature Space

Related works mainly use representation in two feature spaces: image and BEV (see Fig. 1). Image feature space is derived from camera images that contain information in projected camera views. BEV feature space is a joint latent space where features from multiple sensor views are mapped into and from which BEV and 3D representations are decoded, imposing spatial constraints [38]. This concept was pioneered by Phillion and Fidler [39], where they lift each image to a 3D volume and then map all volumes into a joint BEV grid. Typically learned BEV spaces are of low resolution due to computational constraints, as illustrated in Fig. 1.

Most related works aggregate information in BEV space. Representation in BEV space is affected by the uncertainty in depth estimation that is propagated through the lifting step [3], but also preserves a linear appearance of linear motions, simplifying the aggregation process. Only a few models, such as PETRv2 [26], STS [31], DfM [5], BEVStitch [16], and PolarDETR [28] conduct aggregation purely in image feature space. These models typically leverage methods for learning-based temporal stereo. To the best of our knowledge, SoloFusion [14] and BEVStereo [15] are the only approaches that perform temporal aggregation in both feature spaces.

BEVStereo uses depth estimation from two consecutive frames to improve the quality of the lifting and performs parallel temporal aggregation in BEV space. SoloFusion employs parallel aggregation in image space and BEV space. A high-resolution cost volume is created from image space to perform stereo matching between current and previous camera frame. This is complemented by a low-resolution aggregation in BEV space over a longer time horizon. Both approaches have a fixed time horizon by performing parallel temporal aggregation.

D. Motions

Sensors equipped on an autonomous vehicle observe a superposition of two kinds of motions: ego-motion [7], which alters the reference frame for all sensors on the vehicle, and the dynamic motion of objects [8, 13] in the environment.

Ego-motion refers to the movement of the ego vehicle over time relative to static elements in the scene. This concept entails that all sensors mounted on the vehicle undergo the same transformation. In BEV space, ego-motion is compensated by applying the inverse of the rotation and translation to transform previous measurements into the current frame of reference. This process ensures that static elements are consistently represented at the same spatial locations across different time steps, thereby facilitating the aggregation. BEVDet4D [7] offers deeper insights into the benefits of ego-motion compensation to the learning process. All works in this survey that perform temporal aggregation in BEV space apply ego-motion compensation to X_{t-1} . Ego-motion compensation in image space is a more challenging problem that only some works address [31].

The second kind of motion in the scene is the motion of other objects. For the motion of other objects, few approaches [8, 13, 30, 32] use an explicit representation of dynamic objects in the latent state X . The latent state of the model is constrained to capture dynamic objects and their motions effectively. In Table I this is referred to as object-based motion. A notable example is StreamPETR [8], which limits the latent space X_t to cover only sparse object queries from previous time steps, guiding the model to focus on dynamic objects in the scene.

E. Learning Tasks

The learning tasks encompass object detection [1], BEV segmentation [2], and other miscellaneous objectives. Object detection involves estimating the 3D bounding boxes of objects within the scene and is a fundamental task addressed by almost all approaches. BEV segmentation, on the other hand, focuses on representing static elements of the environment, and only some approaches include this task in their scope. Integrating other tasks into an end-to-end learning setup is covered by only a few works so far. The "other" category may involve predicting trajectories [12, 18, 33], object tracking [8, 13, 36], occupancy maps [30], or other downstream tasks.

III. APPROACH

A. Problem Statement

For a given time step t , let $U_t = \{u_t^1, \dots, u_t^n\}$ be the set of image frames from the n monocular cameras mounted on the ego vehicle, e_t be the ego-motion vectors representing translation and rotation from $t-1$ to t , B_t be the set of 6D vectors representing 3D location and dimension of bounding boxes of dynamic objects visible from the ego vehicle, and S_t be the segmentation in BEV space as a 2D grid representing the ground plane with the origin at the ego-vehicle, where each grid pixel represents the class of the static element at that location. The goal is to find a function h that returns bounding boxes B_t and BEV segmentation S_t for a given sequence of sets of image frames $U_{t-k:t}$:

$$(B_t, S_t) = h(U_{t-k:t}, e_{t-k:t}). \quad (3)$$

Function h is typically realized in an encoder-decoder fashion where the encoder f performs temporal aggregation and g is the task-specific decoder: $h(x) = g(f(x))$. The objective of our study is to propose a temporal encoder f that effectively extracts temporal features that can be used by g to estimate B_t and S_t .

B. Approach for Temporal Aggregation

Temporal aggregation is the process of combining inputs $U_{t-k:t}$ to improve the prediction of B_t and S_t at time t . Our survey shows that previous research has predominantly focused on temporal aggregation either in BEV space or image space (see Sec. II-C). Eq. 4 formalizes aggregation in BEV space, Eq. 5 in image space, where lift is the operation that encodes and lifts projected camera features into 3D space, followed by projection into BEV space.

$$(B_t, S_t) = g(f_{\text{BEV}}(\text{lift}(U_{t-k:t}))) \quad (4)$$

$$(B_t, S_t) = g(\text{lift}(f_{\text{img}}(U_{t-k:t}))) \quad (5)$$

We analyze the pros and cons of temporal aggregation in image vs. BEV space using Fig. 1. Consecutive frames in image space offer precise cues of motion due to high resolution and low uncertainty, as visible in Fig. 1 for time t and $t+1$. Those visual cues are desirable for detecting the motion of dynamic objects over short time horizons. Image space aggregation over long time horizons is more challenging since ego-motion compensation cannot be directly applied to features in image space [31]. This can make even static elements appear increasingly different over long time horizons, increasing the difficulty of finding correspondences for aggregation, as visible in Fig. 1 for time t and $t+k$.

Conversely, changes across short time horizons are less apparent in BEV space due to lower resolution and higher uncertainty induced by lifting [3], as visible in BEV space in Fig. 1. However, representing features in BEV space allows for ego-motion compensation, which keeps static elements at the same location in the grid. This enables aggregation over longer time horizons, where a broader range of observations must be aggregated from diverse perspectives and varying occlusions. Our understanding of the time horizon

is consistent with the work in SOLOFusion [14], where they demonstrate that temporal aggregation over short and long time horizons is complementary.

Based on our analysis, we hypothesize that aggregation in image and BEV space is complementary and that combining them leverages the strengths of both representations. To this end, we propose a model that effectively combines temporal aggregation in image and BEV space by extracting temporal features from both representations as formalized in Eq. 6.

$$(B_t, S_t) = g(f_{\text{BEV}}(\text{lift}(U_{t-k:t}, f_{\text{img}}(U_{t-k:t})))) \quad (6)$$

C. TempBEV Model

We design a novel model TempBEV. As per common practice [21, 40, 41], we use BEVFormer [10] as the starting point of our implementation. Eq. 7 formalizes the temporal aggregation of TempBEV, refining Eq. 6. For image space aggregation f_{img} , we add a temporal stereo encoder [42] that performs parallel aggregation of $U_{t-1:t}$. For BEV space aggregation, we keep the recurrent mechanism of BEVFormer to cover a long time horizon, so f_{BEV} uses X_{t-1} . In addition, f_{BEV} also integrates the lifted encodings from U_t and f_{img} .

$$(B_t, S_t) = g(f_{\text{BEV}}(\text{lift}(U_t, f_{\text{img}}(U_{t-1:t})), X_{t-1})) \quad (7)$$

By leveraging both image and BEV feature space for temporal aggregation, TempBEV can learn from data which information to aggregate in BEV space and what temporal features to extract directly from the image space.

Figure 2 is inspired by BEVFormer [10] and provides a visualization of the architecture, illustrating the temporal aggregation mechanisms employed by TempBEV. In image space, a parallel aggregation module combines camera features encoded by an image encoder and temporal stereo features encoded by a temporal stereo encoder. In BEV space, a transformer encoder layer performs temporal self-attention and spatial cross-attention. To quantify the effect of temporal stereo encoding on the detection of dynamic objects and static elements separately, we perform 3D object detection and map BEV segmentation using task-specific heads. Details on the aggregation mechanisms are explained below:

1) *TempBEV Aggregation in Image Space*: Camera frames from multiple views U_t are individually encoded with a shared image encoder to generate camera features. In addition, a temporal stereo encoder f_{img} extracts temporal features from individual image pairs u_{t-1}^k and u_t^k with $k = \{1, \dots, n\}$. The temporal stereo encoder draws inspiration from optical flow approaches. The intuition is that motions in the scene induce optical flow, so extracting features for optical flow from image space helps capture these motions. To confirm our intuition, we select the simplest model for learned optical flow, FlowNet [43], and use its encoder as the temporal stereo encoder in our TempBEV architecture. The FlowNet encoder consists of 10 convolution layers, each followed by Batch Normalization and Leaky ReLU activation. The temporal stereo encoder is shared between all cameras to minimize model size.

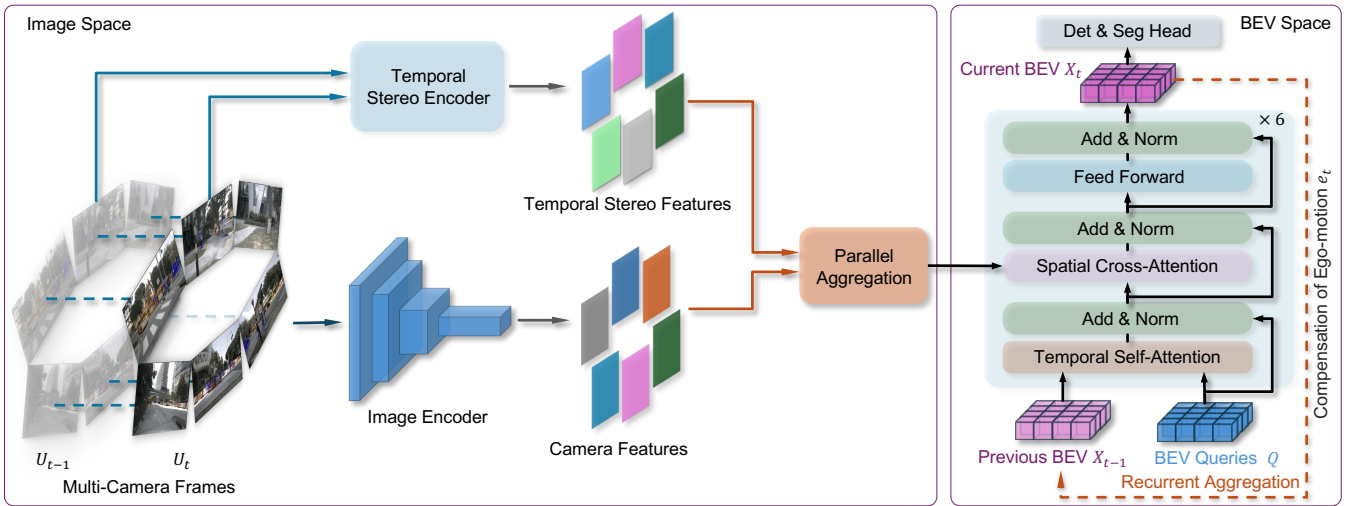


Fig. 2: Proposed TempBEV model architecture with image space and BEV space temporal aggregation mechanisms colored orange.

Image encoder and temporal stereo encoder both provide 4 outputs of intermediate layers. A parallel aggregation module combines temporal stereo features and camera features on the 4 different intermediate resolutions using a 2-layer CNN each. This two-step parallel aggregation captures temporal features over a short time horizon in image feature space.

2) *TempBEV Aggregation in BEV Space*: In BEV space, a transformer encoder layer performs temporal self-attention with the previous BEV grid X_{t-1} in a recurrent fashion. Furthermore, spatial cross-attention is used to lift the aggregation of camera features extracted from U_t and temporal stereo features extracted by f_{img} into BEV space. Learned BEV queries Q are iteratively refined by applying 6 instances of this transformer encoder layer before passing it to the task-specific heads. This setup creates a recurrent chain that facilitates the propagation of information across multiple time steps, enabling the capture of temporal features over long time horizons.

IV. EXPERIMENTS

A. Dataset and Evaluation Metrics

We use the NuScenes dataset [44] for our experiments, which provides data points at 2 Hz. Those include images from 6 monocular cameras U_t and corresponding ground truth bounding boxes B_t and BEV segmentation S_t for static elements including road surface (“Road”), dividers (“Lane”), and pedestrian crossings (“Cross”). The performance on 3D object detection is evaluated using NuScenes Detection Score (NDS) [44] and mean Average Precision (mAP). Intersection over Union (IoU) is evaluated per class for the BEV segmentation task. We provide the results for both tasks separately to quantify the effect of temporal stereo encoding on detecting dynamic objects and static elements.

B. Baseline and Implementation

For uniform comparison, we set BEVFormer [10] as the baseline model and our implementation follows the published code from both BEVFormer [10] and UniAD [41]. The default BEVFormer implementation comprises 6

transformer encoder layers, each composed of self-attention, cross-attention, and feed-forward operations, with layer normalization applied after each steps. Temporal aggregation is facilitated by the self-attention mechanism, which combines information from the previous BEV grid and queries in a recurrent refinement. We adopt the training configuration of BEVFormer [10] with 24 epochs and batch size 1. The model training is performed in parallel on 8/16 NVIDIA A100 GPUs. AdamW is used for optimization with a learning rate of 2×10^{-4} . The size of the BEV grid is 200×200 with perception range $[-51.2 \text{ m}, 51.2 \text{ m}]$ in each BEV direction, resulting in a grid cell size of 0.512 m.

C. Comparative Study

To address the lack of direct comparison of temporal aggregation operators in literature, we conduct comparative experiments. For the lifting mechanism of BEV models, this has been done by Harley *et al.*, showing with SimpleBEV [38] that most performance differences come not from the mechanism itself but from other hyperparameter changes such as batch size. In this study, we evaluate various temporal aggregation operators: temporal self-attention, max pooling, and convolution. We replace the self-attention mechanism in our baseline model with the alternative temporal aggregation operators while keeping the remaining architecture constant. We also evaluate the effect of the time horizon on temporal self-attention. All experiments are trained from scratch on the full NuScenes dataset to be most representative.

Table II shows the quantitative results of different temporal aggregation methods, the first part covers the comparative study. Similar to Table I, we list aggregation operator, recurrent or parallel aggregation, number of time steps, and feature space. Additionally, the model size is shown by the number of learned parameters. All models are trained on 3D object detection and BEV segmentation in a multi-task learning setup. Results for both tasks are reported, going beyond most works that only report results on object detection.

Attention (BEVFormer) is the baseline implementation of the BEVFormer model. The size of the model is 97.69 Mio.

Table II: Quantitative results of comparative study (first section) and TempBEV model (second section) on NuScenes *val* dataset. Results reported for 3D object detection and BEV segmentation learning task. Aggregation is parallel (P) or recurrent (R). Results improving on the BEVFormer baseline are marked in bold, the best results are underlined.

Methods Aggregation Operator	Aggregation			Feature Space		# Params [Mio.]	3D Detection [%]		BEV Seg. IoU [%]		
	R	P	# Steps	Image	BEV		NDS	mAP	Road	Lane	Cross
No temporal aggregation			1		✓	96.31	43.78	36.35	75.14	37.64	22.05
Attention (BEVFormer) - Baseline	✓		4		✓	97.69	50.25	39.82	75.65	38.22	23.89
Max pooling	✓		4		✓	96.31	41.74	34.12	73.84	36.09	19.98
Convolution 1x1	✓		4		✓	97.10	47.42	37.95	74.72	37.05	22.57
Convolution 3x3	✓		4		✓	103.39	48.36	38.48	75.37	37.72	23.74
Convolution 5x5	✓		4		✓	115.98	48.95	39.61	76.51	39.40	25.64
Temporal Stereo Encoder (256C, pret.)		✓	2	✓		101.49	44.23	36.47	75.74	38.21	23.17
TempBEV (64C, single resolution)	✓	✓	4 / 2	✓	✓	98.67	50.44	40.56	76.04	38.74	24.49
TempBEV (64C)	✓	✓	4 / 2	✓	✓	99.55	50.76	39.80	76.30	38.37	23.57
TempBEV (64C, camera-specific)	✓	✓	4 / 2	✓	✓	101.48	50.62	40.21	75.79	38.27	23.56
TempBEV (256C)	✓	✓	4 / 2	✓	✓	102.87	50.84	40.84	76.76	39.53	25.82
TempBEV (256C, pretrained)	✓	✓	4 / 2	✓	✓	102.87	51.31	41.26	76.85	39.34	25.74
TempBEV (1024C, pretrained)	✓	✓	4 / 2	✓	✓	121.75	51.28	41.52	76.35	39.43	25.30

parameters, consisting of the encoder and object detection head as well as the BEV segmentation head (taken from UniAD [41]). Our NDS and mAP scores are roughly on par with the reported values in the BEVFormer paper [10] (NDS 50.3% vs. 52.0%, mAP 39.8% vs. 41.2%). Also, our IoU scores are in the same range (Road IoU 75.7% vs. 77.5%, Lane IoU 38.2% vs. 23.9%). Results slightly differ because our implementation of the BEV segmentation task excludes classes that represent dynamic objects. To quantify the benefit of temporal aggregation, we train the baseline with no temporal aggregation by removing the temporal self-attention mechanism. This shrinks the model size from 97.69 Mio. to 96.31 Mio. parameters. The performance drop on object detection tasks is very strong and roughly matches what is reported in the BEVFormer paper (−6.5 percentage points (pts) vs. −7.2 pts in NDS). We additionally report the BEV segmentation values, which reveal that the temporal aggregation of BEVFormer only has a small benefit for segmenting static elements (e.g., Road IoU only drops from 75.65% to 75.14%).

Max pooling is used as a simple, parameter-free approach (96.3 Mio. parameters, equal to BEVFormer with no temporal aggregation) to assess the relevance of learnable weights for temporal aggregation. In this study, we utilize top-k max pooling [45] with $k = 256$ to map the 512-channel concatenation of X_{t-1} and queries Q to 256 channels. The results of max pooling are worse than the model with no temporal aggregation. We assume the reason is that max pooling cannot separate the inputs from times $t - 1$ and t , and mixes features, resulting in reduced performance.

As indicated in the survey, convolution is another widely employed mechanism for temporal aggregation. In our comparative experiments, we utilize one simple convolution layer with 512 input channels and 256 output channels to aggregate information from X_{t-1} and Q . We use kernel sizes 1x1, 3x3, and 5x5 to evaluate the impact of local context on aggregation performance. As expected, the object detection performance of the simplest convolution operation, i.e., 1x1 (NDS 47.42%), is between no temporal aggregation (NDS

43.78%) and attention (NDS 50.25%). When adding spatial context with a 3x3 or 5x5 kernel, the metrics further increase (NDS 48.36% and 48.95%) However, bigger kernel sizes also increase the number of parameters, with 97.10 Mio. for 1x1, 103.39 Mio. for 3x3, and 115.99 Mio. for 5x5. Overall, compared to the more sophisticated temporal self-attention operator, a basic 1-layer convolution captures a substantial fraction of the benefits of aggregating temporal information.

Flow (256D, pretrained) is the ablated TempBEV model and reports the performance of flow-based image space aggregation without the recurrent BEV space aggregation. The results show slight improvements on all metrics over no temporal aggregation, indicating limited effectiveness of flow-based parallel aggregation in image space by itself.

D. Quantitative Results of TempBEV Model

The second part of Table II shows the quantitative results of different variants of the TempBEV model, serving as an ablation study of the architecture. TempBEV is characterized by the hidden size of its temporal stereo encoder, which is the number of channels of the latent space created by f_{img} . Already in its simplest form, the TempBEV model with a hidden size of 64 (64C) gives an improvement over the BEVFormer baseline with a similar NDS and improved mAP (40.56% vs. 39.82%). When combining the temporal stereo features with the intermediate camera features on all resolutions in the parallel aggregation module, the NDS increases to 50.76%. This however comes with reductions in mAP, Lane IoU, and Pedestrian Crossing IoU. We do not choose the single resolution approach since we assume that combining features in multiple resolutions could be beneficial for bigger hidden sizes. The temporal stereo encoder operates on raw images, which have very different flow patterns based on the camera mounting position. One idea is that camera-specific instances of the temporal stereo encoder could better learn the particular flow patterns of each camera mounting position (e.g., flow for the front camera moves away from the center, but towards the center for the rear camera). The results are mixed with a small benefit only on mAP (increasing to 40.21% from 39.80%).




Fig. 3: Qualitative result of TempBEV model on 3D object detection shown on the left (green: prediction, orange: GT) and BEV segmentation shown on the right (green: road, orange: lane, purple: crossing, black: ego vehicle).

Since the number of parameters increases from 99.55 Mio. to 101.48 Mio., we stick with one shared temporal stereo encoder. We assume that the NuScenes dataset is not large enough to provide sufficient guidance to train camera-specific temporal stereo encoders individually. For this reason, we evaluate transfer learning by using a temporal stereo encoder pretrained on optical flow datasets. We use the encoder of a FlowNetS, the simple variant of FlowNet [43], that is pretrained on Flying Chair Dataset [43]. On a FlowNet with a hidden size of 256, we see a strong improvement in object detection (NDS 51.31% vs. 50.84%). The BEV segmentation results remain similar, indicating that transfer learning from optical flow estimation is primarily helpful in detecting dynamic objects in the scene. The hidden size of the FlowNet encoder is ablated. Increasing the hidden size from 64 to 256 without pretraining shows a strong improvement on all metrics with the number of parameters slightly increasing from 99.55 Mio. to 102.87 Mio. Further increasing hidden size to 1024 for FlowNet with pretraining increases the number of parameters drastically to 121.75 Mio., but shows no additional improvement (NDS 51.31% vs. 51.28%). Hence, for our final model we define TempBEV using a pretrained optical flow encoder with hidden size 256. This adds 6.18 Mio. parameters (+5.3%) to the BEVFormer baseline and changes inference speed on an Nvidia A6000 GPU from 2.18 FPS to 1.94 FPS (-11.0%).

With TempBEV, we combine temporal aggregation in image and BEV space to test our hypothesis from Section III-B. Compared to the BEVFormer baseline model, even with the simplest optical flow encoder from FlowNet, TempBEV performs better by +1.06 pts NDS, +1.44 pts mAP, +1.20 pts Road IoU, +1.12 pts Lane IoU, and +1.85 pts Pedestrian Crossing IoU. This confirms our hypothesis of the complementary nature of image and BEV space temporal aggregation. Starting from the baseline with no temporal aggregation and adding the image space aggregation improves NDS by just +0.45 pts. Starting from the baseline with BEV space aggregation and adding image space aggregation improves NDS even more by +1.06 pts. This highlights that image and BEV space aggregation are not just complementary but also show synergy effects. The results show the effectiveness of our adjustments to the temporal aggregation mechanism. We leave it for future work to use these insights to improve the latest state of the art models.

E. Qualitative Results

Fig. 3 shows qualitative results of TempBEV on 3D Object Detection and BEV Segmentation tasks. Frames from 6 cameras are overlaid with ground truth (GT) and prediction bounding boxes projected into the camera view. On the right, GT and predicted BEV segmentation are visualized. TempBEV accurately predicts the bounding boxes of dynamic objects in 3D and the BEV segmentation of the road class. It predicts dividers (class lane) that separate the bus lane behind the ego vehicle, that are missing in GT.

V. CONCLUSION

In this paper, we presented a survey and a comparative study on temporal aggregation mechanisms for BEV encoders. Based on the gained insights, we proposed TempBEV, a novel BEV model that combines recurrent aggregation in BEV space with parallel aggregation in image space. An optical flow encoder is used to extract temporal stereo features directly from pairs of subsequent camera frames in image space. Experiments show that TempBEV provides a significant increase in performance by using a simple optical flow encoder, with further potential to be expected from more sophisticated approaches. The results suggest a synergy effect between temporal aggregation in different representations, making a strong case for combined aggregation in both image space and BEV space. We validated our hypothesis based on the BEVFormer model, which is commonly used for BEV encoding [40, 41], and expect similar synergy effects when applied to other state of the art BEV encoders.

REFERENCES

- [1] X. Chen, K. Kundu, Z. Zhang, H. Ma, S. Fidler, and R. Urtasun, "Monocular 3d object detection for autonomous driving," in *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2016, pp. 2147–2156.
- [2] B. Pan, J. Sun, H. Y. T. Leung, A. Andonian, and B. Zhou, "Cross-view semantic segmentation for sensing surroundings," *IEEE Robotics and Automation Letters*, vol. 5, no. 3, pp. 4867–4873, 2020.
- [3] Y. Lu, X. Ma, L. Yang, T. Zhang, Y. Liu, Q. Chu, J. Yan, and W. Ouyang, "Geometry uncertainty projection network for monocular 3d object detection," in *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 2021, pp. 3111–3121.
- [4] J. Davis, R. Ramamoorthi, and S. Rusinkiewicz, "Spacetime stereo: A unifying framework for depth from triangulation," in *2003 IEEE Computer Society Conference on Computer Vision and Pattern Recognition, 2003. Proceedings.*, vol. 2. IEEE, 2003, pp. II–359.

- [5] T. Wang, J. Pang, and D. Lin, "Monocular 3d object detection with depth from motion," in *European Conference on Computer Vision*. Springer, 2022, pp. 386–403.
- [6] T. Moninger, J. Schmidt, J. Rupperecht, D. Raba, J. Jordan, D. Frank, S. Staab, and K. Dietmayer, "Scene: Reasoning about traffic scenes using heterogeneous graph neural networks," *IEEE Robot. Autom. Lett.*, 2023.
- [7] J. Huang and G. Huang, "Bevdet4d: Exploit temporal cues in multi-camera 3d object detection," *arXiv preprint arXiv:2203.17054*, 2022.
- [8] S. Wang, Y.-H. Liu, T. Wang, Y. Li, and X. Zhang, "Exploring object-centric temporal modeling for efficient multi-view 3d object detection," *2023 IEEE/CVF International Conference on Computer Vision (ICCV)*, pp. 3598–3608, 2023.
- [9] 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 *Proceedings of the AAAI Conference on Artificial Intelligence*, vol. 37, no. 2, 2023, pp. 1477–1485.
- [10] Z. Li, W. Wang, H. Li, E. Xie, C. Sima, T. Lu, Y. Qiao, and J. Dai, "Befvformer: Learning bird's-eye-view representation from multi-camera images via spatiotemporal transformers," in *European conference on computer vision*. Springer, 2022, pp. 1–18.
- [11] C. Yang, Y. Chen, H. Tian, C. Tao, X. Zhu, Z. Zhang, G. Huang, H. Li, Y. Qiao, L. Lu *et al.*, "Befvformer v2: Adapting modern image backbones to bird's-eye-view recognition via perspective supervision," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2023, pp. 17 830–17 839.
- [12] Y. Zhang, Z. Zhu, W. Zheng, J. Huang, G. Huang, J. Zhou, and J. Lu, "Beverse: Unified perception and prediction in birds-eye-view for vision-centric autonomous driving," *arXiv preprint arXiv:2205.09743*, 2022.
- [13] L. Qing, T. Wang, D. Lin, and J. Pang, "Dort: Modeling dynamic objects in recurrent for multi-camera 3d object detection and tracking," in *Conference on Robot Learning*. PMLR, 2023, pp. 3749–3765.
- [14] 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," 2023.
- [15] Y. Li, H. Bao, Z. Ge, J. Yang, J. Sun, and Z. Li, "Bevstereo: Enhancing depth estimation in multi-view 3d object detection with temporal stereo," in *Proceedings of the AAAI Conference on Artificial Intelligence*, vol. 37, no. 2, 2023, pp. 1486–1494.
- [16] Y. B. Can, A. Liniger, O. Unal, D. Paudel, and L. Van Gool, "Understanding bird's-eye view of road semantics using an onboard camera," *IEEE Robotics and Automation Letters*, vol. 7, no. 2, pp. 3302–3309, 2022.
- [17] J. Yao and Y. Lai, "Dynamicbev: Leveraging dynamic queries and temporal context for 3d object detection," *arXiv preprint arXiv:2310.05989*, 2023.
- [18] 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 *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 2021, pp. 15 273–15 282.
- [19] 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," *arXiv preprint arXiv:2301.07870*, 2023.
- [20] Y. Li, B. Huang, Z. Chen, Y. Cui, F. Liang, M. Shen, F. Liu, E. Xie, L. Sheng, W. Ouyang *et al.*, "Fast-bev: A fast and strong bird's-eye view perception baseline," *arXiv preprint arXiv:2301.12511*, 2023.
- [21] 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," 2023.
- [22] A. Saha, O. Mendez, C. Russell, and R. Bowden, "Translating images into maps," in *2022 International conference on robotics and automation (ICRA)*. IEEE, 2022, pp. 9200–9206.
- [23] G. Bae, I. Budvytis, and R. Cipolla, "Multi-view depth estimation by fusing single-view depth probability with multi-view geometry," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2022, pp. 2842–2851.
- [24] Y. Yao, Z. Luo, S. Li, T. Fang, and L. Quan, "Mvsnet: Depth inference for unstructured multi-view stereo," in *Proceedings of the European conference on computer vision (ECCV)*, 2018, pp. 767–783.
- [25] Z. Qi, J. Wang, X. Wu, and H. Zhao, "Ocbbev: Object-centric bev transformer for multi-view 3d object detection," *arXiv preprint arXiv:2306.01738*, 2023.
- [26] Y. Liu, J. Yan, F. Jia, S. Li, A. Gao, T. Wang, and X. Zhang, "PetrV2: A unified framework for 3d perception from multi-camera images," in *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, October 2023, pp. 3262–3272.
- [27] L. Chambon, E. Zablocki, M. Chen, F. Bartoccioni, P. Perez, and M. Cord, "Pointbev: A sparse approach to bev predictions," *arXiv preprint arXiv:2312.00703*, 2023.
- [28] S. Chen, X. Wang, T. Cheng, Q. Zhang, C. Huang, and W. Liu, "Polar parametrization for vision-based surround-view 3d detection," *arXiv preprint arXiv:2206.10965*, 2022.
- [29] Y. Jiang, L. Zhang, Z. Miao, X. Zhu, J. Gao, W. Hu, and Y.-G. Jiang, "Polarformer: Multi-camera 3d object detection with polar transformer," in *Proceedings of the AAAI Conference on Artificial Intelligence*, vol. 37, no. 1, 2023, pp. 1042–1050.
- [30] S. Hu, L. Chen, P. Wu, H. Li, J. Yan, and D. Tao, "St-p3: End-to-end vision-based autonomous driving via spatial-temporal feature learning," in *European Conference on Computer Vision*. Springer, 2022, pp. 533–549.
- [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," 2022.
- [32] H. Liu, Y. Teng, T. Lu, H. Wang, and L. Wang, "Sparsebev: High-performance sparse 3d object detection from multi-camera videos," in *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 2023, pp. 18 580–18 590.
- [33] S. Fang, Z. Wang, Y. Zhong, J. Ge, and S. Chen, "Tbp-former: Learning temporal bird's-eye-view pyramid for joint perception and prediction in vision-centric autonomous driving," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2023, pp. 1368–1378.
- [34] Y. Li, Y. Chen, X. Qi, Z. Li, J. Sun, and J. Jia, "Unifying voxel-based representation with transformer for 3d object detection," *Advances in Neural Information Processing Systems*, vol. 35, pp. 18 442–18 455, 2022.
- [35] Z. Qin, J. Chen, C. Chen, X. Chen, and X. Li, "Unifusion: Unified multi-view fusion transformer for spatial-temporal representation in bird's-eye-view," in *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 2023, pp. 8690–8699.
- [36] C. Han, J. Sun, Z. Ge, J. Yang, R. Dong, H. Zhou, W. Mao, Y. Peng, and X. Zhang, "Exploring recurrent long-term temporal fusion for multi-view 3d perception," *arXiv preprint arXiv:2303.05970*, 2023.
- [37] X. Zhu, W. Su, L. Lu, B. Li, X. Wang, and J. Dai, "Deformable {detr}: Deformable transformers for end-to-end object detection," in *International Conference on Learning Representations*, 2021.
- [38] A. W. Harley, Z. Fang, J. Li, R. Ambrus, and K. Fragkiadaki, "Simplebev: What really matters for multi-sensor bev perception?" in *2023 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2023, pp. 2759–2765.
- [39] J. Philion and S. Fidler, "Lift, splat, shoot: Encoding images from arbitrary camera rigs by implicitly unprojecting to 3d," in *Proceedings of the European Conference on Computer Vision*, 2020.
- [40] T. Yuan, Y. Liu, Y. Wang, Y. Wang, and H. Zhao, "Streammapnet: Streaming mapping network for vectorized online hd map construction," in *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision*, 2024, pp. 7356–7365.
- [41] 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, "Planning-oriented autonomous driving," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2023.
- [42] H. Laga, L. V. Jospin, F. Boussaid, and M. Bennamoun, "A survey on deep learning techniques for stereo-based depth estimation," *IEEE transactions on pattern analysis and machine intelligence*, vol. 44, no. 4, pp. 1738–1764, 2020.
- [43] A. Dosovitskiy, P. Fischer, E. Ilg, P. Hausser, C. Hazirbas, V. Golkov, P. Van Der Smagt, D. Cremers, and T. Brox, "Flownet: Learning optical flow with convolutional networks," in *Proceedings of the IEEE international conference on computer vision*, 2015, pp. 2758–2766.
- [44] H. Caesar, V. Bankiti, A. H. Lang, S. Vora, V. E. Liong, Q. Xu, A. Krishnan, Y. Pan, G. Baldan, and O. Beijbom, "nusenes: A multimodal dataset for autonomous driving," in *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 2020, pp. 11 621–11 631.
- [45] Q. Liu, F. Yu, S. Wu, and L. Wang, "A convolutional click prediction model," in *Proceedings of the 24th ACM international on conference on information and knowledge management*, 2015, pp. 1743–1746.