

TMA: Temporal Motion Aggregation for Event-based Optical Flow

Haotian Liu¹ Guang Chen^{1*} Sanqing Qu¹ Yanping Zhang¹
 Zhijun Li¹ Alois Knoll² Changjun Jiang¹
¹Tongji University ²Technical University of Munich

Abstract

Event cameras have the ability to record continuous and detailed trajectories of objects with high temporal resolution, thereby providing intuitive motion cues for optical flow estimation. Nevertheless, most existing learning-based approaches for event optical flow estimation directly remould the paradigm of conventional images by representing the consecutive event stream as static frames, ignoring the inherent temporal continuity of event data. In this paper, we argue that temporal continuity is a vital element of event-based optical flow and propose a novel Temporal Motion Aggregation (TMA) approach to unlock its potential. Technically, TMA comprises three components: an event splitting strategy to incorporate intermediate motion information underlying the temporal context, a linear lookup strategy to align temporally fine-grained motion features and a novel motion pattern aggregation module to emphasize consistent patterns for motion feature enhancement. By incorporating temporally fine-grained motion information, TMA can derive better flow estimates than existing methods at early stages, which not only enables TMA to obtain more accurate final predictions, but also greatly reduces the demand for a number of refinements. Extensive experiments on DSEC-Flow and MVSEC datasets verify the effectiveness and superiority of our TMA. Remarkably, compared to E-RAFT, TMA achieves a 6% improvement in accuracy and a 40% reduction in inference time on DSEC-Flow. Code will be available at <https://github.com/ispc-lab/TMA>.

1. Introduction

Optical flow aims to compute the velocity of objects on the image plane without geometry prior and is a fundamental topic in event-based vision [8, 42]. It plays an important role in many applications, such as ego-motion estimation [38, 45], image reconstruction [24], and video frame interpolation [32, 31]. Recently, learning-based methods [29, 18, 30] have dominated frame-based optical flow




Figure 1: **Accuracy illustrations.** (a) Reference image from DSEC-Flow. (b) Corresponding event frame. (c) The final flow prediction and visualization of the motion feature from E-RAFT [11]. (d) The final flow prediction and visualization of the motion feature from our proposed method. We visualize the motion feature at the initial stage by taking the average across channels. E-RAFT fails to generate informative motion features and results in blurred boundaries. In contrast, our method utilizes temporally fine-grained motion information to address the information scarcity issue in motion features, generating high-quality predictions.

by employing correlation volumes to derive motion features for flow regression. Inspired by this, several event-based approaches [11, 34] adopt a similar paradigm by converting the consecutive event stream into grid-like representations. Despite encouraging progress, these methods still suffer from non-negligible limitations, due to the great differences between event data and conventional images. Distinct from dense and colorful conventional (RGB) images, event data features spatial sparsity and a lack of intensity information. It is prone to result in close matching scores and invalid regions in the correlation volume. Consequently, less informative motion features are derived and regress inaccurate

*corresponding author: guangchen@tongji.edu.cn

predictions. Figure 1 illustrates this issue in E-RAFT [11].

As the saying goes, each coin has two sides. Though event cameras encode visual information in a sparse manner, they are capable of capturing continuous and detailed object trajectories with high temporal resolution, thereby providing rich motion cues for optical flow estimation. Several model-based methods [9, 27] benefit from the temporal continuity by warping events along point trajectories, achieving decent performance. Drawing on this observation, we contend that the temporal continuity is a vital element of event-based optical flow and propose to unleash its potential within a learning-based framework.

To materialize our idea, we propose a novel Temporal Motion Aggregation (TMA) approach to explore the inherent temporal continuity of event data. Technically, we revamp the dominant learning-based architectures with three distinct components. We first introduce an event splitting strategy for temporally-dense correlation volumes computation. By splitting the event stream into multiple segments and extracting their features, feature similarities are compared between the first feature and all others, which record rich intermediate motion information. Aware of that correlation volumes record motions between different time spans, we then design a linear lookup strategy to sample each correlation volume based on the corresponding flow estimate to encode motion features. As a result, fine-grained motions are warped to the same pixels across motion features. Furthermore, we consider the incorrect motion patterns in intermediate motion features due to the manual lookup. Therefore, we propose a novel pattern aggregation module for motion feature enhancement by aggregating consistent patterns between each intermediate motion feature and the last one (without manual lookup). Thanks to the incorporation of temporally fine-grained motion information, TMA can effectively address the information scarcity in motion features and thus generate good flow estimates at early stages, which not only enable TMA to obtain accurate flow predictions, but also greatly reduces the demand for a number of time-consuming refinements.

We evaluate our TMA on DSEC-Flow [10] and MVSEC [44] datasets. Extensive experiments demonstrate the superior advantages of TMA.

In summary, our main contributions are as follows:

- We argue that the temporal continuity of event data is a vital element of event-based optical flow and propose to unlock its potential in a learning-based framework.
- We propose a novel Temporal Motion Aggregation (TMA) approach, which comprises three components: an event splitting strategy, a linear lookup strategy and a motion pattern aggregation module.
- The incorporation of temporally fine-grained motion information enables TMA to achieve high accuracy

while maintaining high efficiency. Compared with E-RAFT, TMA achieves a 6% improvement in accuracy and a 40% reduction in inference time on DSEC-Flow.

2. Related Work

Learning methods for optical flow estimation. Recently, correlation volume based approaches dominate the learning-based optical flow estimation [25, 17, 29, 18, 37, 14]. As a milestone, PWCNet [29] summarizes a simple principle in architecture construction: pyramidal processing, warping and correlation volumes computation. Following the well-established pipeline, LiteFlowNet series [16, 17, 15] propose cascaded feature pyramids with a lightweight convolutional burden. VCN [37] proposes 4D volumetric U-Net architectures to enlarge the receptive fields for correlation volumes. IRR [18] notes the reusability of pyramidal refinement for flow update and proposes a shared-weights refiner across all scales, which can be combined with mainstream learning methods.

However, the correlation volumes used in previous works are computed locally, which cannot fully deal with long-term motions. RAFT [30] firstly introduces the full correlation volume to capture global visual similarities. Separable Flow [41] proposes a separable correlation volume module as a drop-in replacement to correlation volumes and uses non-local aggregation layers [39, 40] to refine the correlation volumes efficiently. Besides, transformer components show promising performance on optical flow [36, 19, 35]. GMA [19] notices the underlying motion clues in the context and introduces an attention layer to aggregate hidden motions, improving the accuracy based on RAFT architecture. GMFlow [36] reformulates optical flow as a global matching problem by attention blocks, suggesting a new paradigm for flow estimation.

We draw inspiration from GMA [19] to design our motion pattern aggregation module. In contrast to GMA which aggregates hidden motion patterns in the spatial context for motion feature enhancement, we extend the pattern aggregation to the temporal dimension and leverage feature cross-similarities to enhance the motion feature.

Event-based optical flow estimation. A number of traditional works remould frame-based approaches into event-based optical flow estimation. Benosman et al. [4] extends the classic Lucas-Kanade method to event-based optical flow. Liu et al. [21, 22] implement event-based corner keypoint detection and novel block-matching optical flow. Gallego et al. [9] propose a unifying contrast maximization (CM) framework and MultiCM [27] extends it to complex scenes. Besides, plane fitting [1, 3], filter banks [6] and time surface matching [2] are proposed to boost performance.

More recently, learning-based methods have dominated frame-based optical flow, inspiring attempts at event-based optical flow. Early learning methods [43, 45, 38] represent

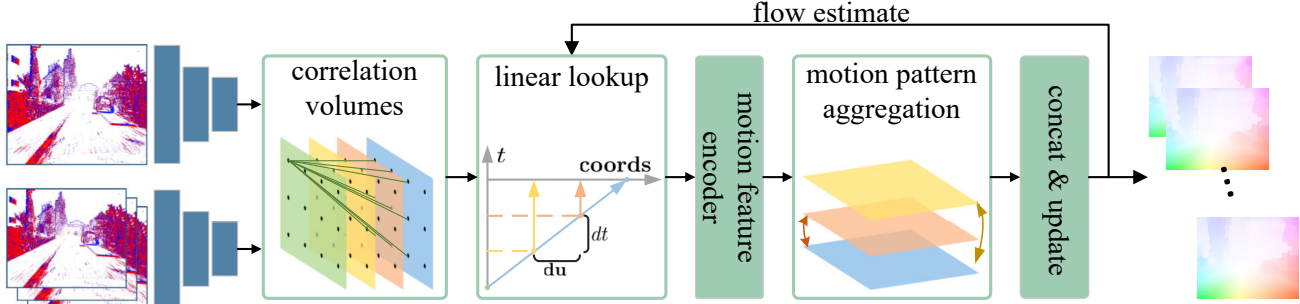


Figure 2: **The overall architecture of TMA.** First, we split the event stream into segments and extract features with a share-weights feature extractor. Then, regarding the first feature as reference, we correlate all other features to compute temporally-dense correlation volumes, which contain rich embedded motion information. Next, aware of the different time spans of correlation volumes, we look up each correlation volume based on a corresponding flow estimate. All sampled correlation maps are delivered into a share-weights motion feature encoder to derive motion features. We further propose a novel motion pattern aggregation module to enrich the spatial information of motion features. By concatenating the enhanced motion features and utilizing several refinements, high-quality flow predictions are generated.

the event stream as a static frame and utilize the architecture of U-Net [26] to predict optical flow, which can only address small motions. Mainstream works [7, 34, 11] introduce two consecutive event frames and follow the correlation volume based pipeline (PWCNet and RAFT), achieving higher accuracy on large-scales datasets [44, 10]. ERAFT [11, 12] extends RAFT [30] and introduces full correlation volume into event-based vision. DCEIFlow [34] combines conventional images and event data to boost higher accuracy. Note that our work and STE-FlowNet [7] share a similar motivation to introduce temporal continuity with fine-grained discretization. However, STE-FlowNet takes a frame-by-frame manner [18] to update flow, leading to an insufficient inference process.

Contrary to previous learning-based methods, we do not merely adopt the classic correlation volume based pipeline. Instead, we regard the temporal continuity of event data as a vital component of optical flow and incorporate the intermediate motion information into a learning-based framework, which boosts higher accuracy and efficiency.

3. Methodology

Given an input event stream $\{\mathcal{E}\}$ from t_0 to t_1 , the goal of optical flow is to estimate a dense displacement field $\mathbf{u} : \mathbb{R}^2 \rightarrow \mathbb{R}^2$ which maps each per-pixel (x, y) at t_0 to its corresponding coordinate (x', y') at t_1 . To address this, mainstream learning-based approaches compute correlation volumes from two event frames as a matching prior. Such principle inspires several methods [11, 34] by squeezing event data into consecutive frames. Unfortunately, the characteristics of event data, spatial sparsity and lack of intensity records, are prone to result in close matching scores and invalid regions in the correlation volume. Consequently, less informative motion features are derived and regress inaccurate predictions. Here we address the information scarcity

of motion features from the attractive temporal continuity of event data. Temporally fine-grained motion information in event data is sufficient to supplement the spatial richness of motion features, yielding high-quality flow predictions.

In the following, we first describe the preliminary and event representation, and then present the detailed architecture of the proposed TMA.

3.1. Preliminary and Event Representation

Event cameras record light intensity changes in an asynchronous way. Each pixel serves as an independent trigger and generates an event instantly whenever the log intensity changes exceeding a threshold C . An event $e_k = (x_k, p_k, t_k)$ is a triplet containing triggered pixel location $\mathbf{x}_k = (x_k, y_k)^T$, timestamp t_k and signed intensity change ($p_k = \pm 1$).

To ease the feature extraction and preserve the temporal continuity of the event data, we firstly transform the event stream into a 3D volume $\mathbf{V}(x, y, t)$ along the temporal dimension following existing works [45, 11, 24]. Given an event stream $\{\mathcal{E}\}$, $\mathbf{V}(x, y, t)$ is generated as:

$$t_i^* = (B - 1)(t_i - t_1)/(t_{N_e} - t_1) \quad (1)$$

$$\mathbf{V}(x, y, t) = \sum_i p_i k_b(x - x_i) k_b(y - y_i) k_b(t - t_i^*) \quad (2)$$

$$k_b(a) = \max(0, 1 - |a|), \quad (3)$$

where t_i^* represents the i -th time bin. B and N_e represent the length of time bins and the number of events, respectively. $k_b(a)$ is the bilinear interpolation function.

3.2. TMA Architecture

We develop our framework based on the successful RAFT architecture [30]. Figure 2 summarizes the architecture of the proposed TMA, which consists of three components, *an event splitting strategy, a linear lookup strategy*

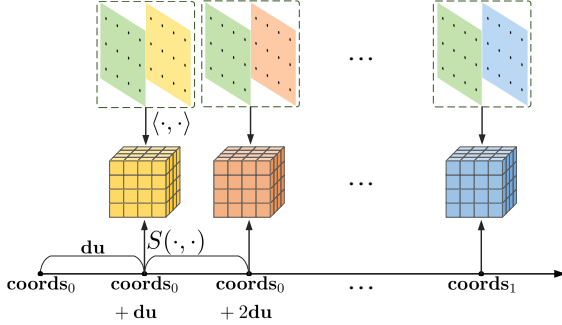


Figure 3: **Linear lookup strategy.** We compute correlation volumes by comparing the similarity of first feature with all others, which are misaligned due to different time spans. To rectify this, we narrow down the lookup coordinates linearly to warp correlation maps to spatially aligned pixels.

and a *motion pattern aggregation module*. With the proposed method, we successfully unlock the potential of temporal continuity to predict event-based optical flow within a learning-based framework.

Event splitting. Instead of viewing the continuous event stream as a static frame which tends to be a matter of common knowledge in recent works [45, 11, 34], we prefer to identify it as a high frame rate video clip, which records continuous and detailed trajectories of objects, namely, provides intuitive motion cues for optical flow estimation. To introduce rich intermediate motion information into our method, we first propose an event splitting strategy, which splits the event stream into multiple segments. Therefore, more intermediate event frames are retained for correlation volume computation.

In particular, given an event stream $\{\mathcal{E}\}$ from t_0 to t_1 , we split it into g segments equally and the time span between adjacent segments is dt . Besides, we introduce an auxiliary event segment from $t_0 - dt$ to t_0 as reference. Then, $g + 1$ event representations are generated following Sec. 3.1. All $g + 1$ event representations are then delivered into a share-weights encoder to achieve features $F_i \in \mathbb{R}^{H \times W \times C}$, $i = 0, 1, 2, \dots, g$, where H , W and D represent the height, width and channel of features. We use all features to compute temporally-dense correlation volumes C_i , $i = 1, 2, \dots, g$, which are given by:

$$C_i = \frac{F_0 F_i^T}{\sqrt{D}} \in \mathbb{R}^{H \times W \times H \times W}, i = 1, 2, \dots, g, \quad (4)$$

where F_0 is the first feature and F_i is the i -th feature derived from the i -th event frame.

Temporally-dense correlation volumes compare feature similarities across multiple time spans, enabling refined records of relative motions. This construction is vastly different from the classic full correlation volume that only contains feature similarities between a fixed time span.

Linear lookup. Recall that each correlation volume C_i

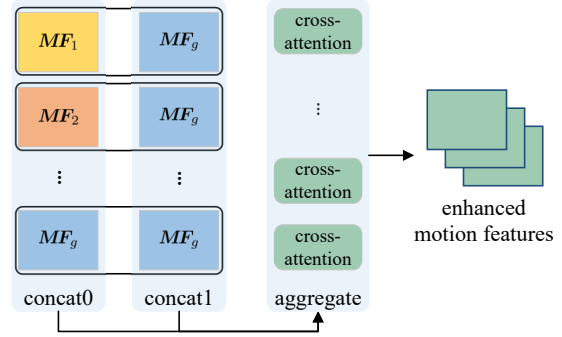


Figure 4: **Motion pattern aggregation.** We pair each intermediate motion feature to the last motion feature. Then each pair is sent to cross-attention blocks to aggregate consistent patterns. Hence, each intermediate motion feature carries information from the last one and consistent patterns receive emphasis.

is achieved by correlating F_0 and F_i , which means the recorded relation motions are misaligned in time. While our goal is to predict optical flow from t_0 to t_1 , we hope that all sampled correlation maps could present the relative motions within the same time span. Therefore, based on that correlation volumes are evenly constructed in time and magnitudes of optical flow change linearly over a short period, we narrow down lookup coordinates in a linear fashion to match the timestamp of each correlation volume.

Figure 3 displays the linear lookup strategy. In detail, given current flow estimate \mathbf{u} , we split it into g segments to sample each correlation volume. The linear lookup operation can be expressed as:

$$\mathbf{du} = (\mathbf{coords}_1 - \mathbf{coords}_0)/g = \mathbf{u}/g \quad (5)$$

$$\mathbf{corr}_i = S(C_i, \mathbf{coords}_0 + i\mathbf{du}), i = 1, 2, \dots, g, \quad (6)$$

where \mathbf{coords}_0 is the initialized coordinates at t_0 (reference time), $S(\cdot, \cdot)$ represents the lookup operation, \mathbf{coords}_i are lookup coordinates for C_i and \mathbf{corr}_i is the correlation map sampled from C_i . After delivering all correlation maps into a share-weights motion feature encoder, aligned motion features are achieved.

From the perspective of image contrast [9], the standard lookup is responsible for warping possible matching points between two images to the same pixel. Our linear lookup extends it by warping matching points of a temporally fine-grained motion to the same pixel. As a result, similar motions are presented at same pixels across correlation maps, showing high contrast.

Motion pattern aggregation. A standard operation for next step is to concatenate the motion features as a whole to update flow and the improvements made so far already alleviate the information scarcity in motion features. However, we notice that the concatenation operation treats all motion features equally, which means the majority of motion features, intermediate ones, dominate the flow regression. Due




Figure 5: **Qualitative examples of DSEC-Flow.** For each two columns, we show the reference image and event frame and compare our method with a state-of-the-art baseline E-RAFT [11]. Significant improvements are highlighted by red boxes.

to the linear optical flow assumption and manual lookup operation, these motion features are not perfectly correct. A natural thought is to put more emphasis on the last motion feature that is without manual lookup and retain the patterns in intermediate motion features that are consistent with the last motion feature. To this end, we balance the weights between the last motion feature and others through cross-similarities computed by transformers [33, 19].

Figure 4 shows the overview of motion pattern aggregation module. Given a set of motion features $\mathbf{MF}_i \in \mathbb{R}^{H \times W \times D'}$, $i = 1, 2, \dots, g$, we pair each intermediate motion feature to the last one. Then each pair $(\mathbf{MF}_i, \mathbf{MF}_g)$ is sent to cross-attention blocks to compare cross-similarities and aggregate consistent patterns. The enhanced motion feature $\widehat{\mathbf{MF}}_i$ is obtained as:

$$\mathbf{A} = \text{Attention}(\mathbf{MF}_i W^Q, \mathbf{MF}_g W^K, \mathbf{MF}_g W^V) \quad (7)$$

$$\widehat{\mathbf{MF}}_i = \mathbf{MF}_i + \text{MLP}([\mathbf{MF}_i, \mathbf{A} W^O]), \quad (8)$$

where $\text{Attention}(Q, K, V) = \text{softmax}\left(\frac{QK^T}{\sqrt{d_k}}\right)V$, d_k is the dimension of K . W^Q, W^K, W^V are projection matrices. $\text{MLP}(\cdot)$ is multi-layer perceptron and W^O is another projection matrix. We omit the normalization functions in the manuscript for simplicity. Motion pattern aggregation module does not change the dimension of motion features, so multiple cross-attention layers can be stacked easily.

Once motion features are enhanced by the motion pattern aggregation module, each $\widehat{\mathbf{MF}}_i$ carries motion patterns that are consistent with \mathbf{MF}_g . Namely, the last motion feature receives high attention and consistent patterns in motion features are fully utilized.

Above three components are seamlessly integrated into a unified pipeline: incorporating intermediate motion information, warping matching motions and emphasizing consistent motion patterns. Attributed to the innovative application of temporal continuity to a learning-based frame-

work, our proposed method is capable of overcoming the information scarcity issue in motion features and generating high-quality flow estimates even at an early refinement stage. This enables our method to achieve accurate final predictions, while reducing the demand for a number of refinements.

3.3. Supervision

Following RAFT [30], we choose L_1 distance between the predictions and ground-truth as the supervision signal for our model. The loss function \mathcal{L} is formally defined as:

$$\mathcal{L} = \sum_{j=1}^N \gamma^{N-j} \|\mathbf{u}^{gt} - \mathbf{u}^j\|_1, \quad (9)$$

where \mathbf{u}^{gt} denotes the flow ground-truth and \mathbf{u}^j the flow prediction at j -th stage. N is the total refinement stages. The weight γ balances different stages of predictions.

4. Experiments

Datasets and evaluation setup. Following previous works [11, 45, 7], we provide experimental details and extensive comparison results on two event-based datasets DSEC-Flow [10] and MVSEC [44]. For DSEC-Flow, we train models on the official training set and evaluate DSEC-Flow test set on the public benchmark. For MVSEC, we conduct two kinds of experiments. Basically, following previous methods [43, 27], we train methods on outdoor_day2 sequence with ground-truth corresponding to a time interval of $dt = 1$ and $dt = 4$ gray images and evaluate 800 frames of outdoor_day1 and three indoor_flying sequences. Equally importantly, considering the supervised learning setting of our method and the large gap between outdoor_day and indoor_flying data, we extend the range of training data by including one indoor_flying sequence each time to further boost the accuracy of supervised learning methods.

Methods		EPE	1PE	3PE	Time(ms)
MB	MultiCM [27]	3.47	76.6	30.9	-
SL	EV-FlowNet [†] [43]	2.32	55.4	18.6	5
	E-RAFT [11]	0.79	12.5	2.7	52
	Ours	0.74	10.9	2.3	30

Table 1: **Results on DSEC-Flow.** Best accuracy in bold. [†] denotes the result is taken from E-RAFT [11]. Model-based (MB) methods need no training data; and supervised learning (SL) methods need ground-truth.

Metrics. The end-point-error (EPE) is used as core metric for prediction accuracy for both DSEC-Flow and MVSEC. To evaluate the robustness against large displacements, we compute the percentage of pixels with EPE greater than M pixels (MPE , $M = 1, 3$) for DSEC-Flow. Percentage of pixels with EPE above 3 and 5% of the magnitude of the flow ground-truth (% Outlier) are presented for MVSEC. On DSEC-Flow, metrics are measured over pixels with valid ground-truth. While on MVSEC, metrics are measured over pixels with valid ground-truth and at least one triggered event.

Implementation details. The proposed TMA is implemented using PyTorch. We set the event segment $g = 5$. For experiments on DSEC-Flow, the channel of event representation B for each segment is set to 3. For MVSEC, B is set to 1 when $dt = 1$ and 3 when $dt = 4$, respectively. The sum of channels keeps the same with the state-of-the-art E-RAFT [11]. We only stack 1 layer of motion pattern aggregation module. Other components are basically identical to RAFT’s model, including feature extractor, motion feature encoder and flow updater. Note that we use a small feature extractor whose final channel is 128. We predict 6 levels of optical flow and the weight γ is set to 0.8 for supervision. In training, we exploit AdamW [23] optimizer and One-cycle learning rate scheduler [28]. We train models on DSEC-Flow for 200K steps and MVSEC for 100K steps, both with a learning rate of 0.0002 and a batch size of 6.

4.1. DSEC-Flow

Accuracy comparison. Evaluation results on DSEC-Flow benchmark¹ are provided in Table 1. Compared with E-RAFT, our method improves the EPE from 0.79 to 0.74. Our method is also leading in 1PE and 3PE with a margin of 1.6 and 0.4, showing better robustness.

Qualitative results are exhibited in Figure 5. In challenging regions such as textureless areas (the wall in second two columns, the shrubs in last two columns), existing correlation volume contains similar matching scores and fails to generate informative motion features, leading to low-quality predictions. On the contrary, our method aggregates

¹The ground-truth of DSEC-Flow dataset is unavailable, so we use last three checkpoints to evaluate the online benchmark. <https://dsec.ifi.uzh.ch/uzh/dsec-flow-optical-flow-benchmark/>.




Figure 6: **EPE vs. number of iterations at inference.** The figure exhibits the prediction results on DSEC-Flow at different stages during inference. Our method achieves an accuracy (EPE = 0.81) comparable to E-RAFT with 2 iterations while E-RAFT needs 6 iterations.

temporally consistent motion patterns to enhance motion features, distinguishing different moving objects and generating high-quality predictions.

Efficiency comparison. The last column of Table 1 exhibits the inference time. Our method enjoys faster inference speed compared with E-RAFT, which is mainly attributed to fewer iterative updates. The reasonable explanation is that our method utilizes temporal motion information to enrich motion features, thus being able to generate good predictions even at an initial stage. To prove this, we compare the accuracy with E-RAFT at early inference stages in Figure 6. Our method achieves better flow initialization at early stages, supporting our argument. The ablation experiment on different iterations (Table 4e) also demonstrates that our method requires fewer iterations (ours: 6, E-RAFT: 12) to achieve surpassing accuracy.

4.2. MVSEC

Inter-domain evaluation. Table 2 reports the results on MVSEC training with outdoor_day2 sequence. The top of Table 2 reports the results corresponding to $dt = 1$ grayscale frame and the bottom corresponds to $dt = 4$ grayscale frames. Our method shows comparable accuracy on outdoor_day1 for $dt = 1$ and highest accuracy on outdoor_day1 for $dt = 4$ among all compared methods. Unfortunately, like two other supervised learning methods, DCEI-Flow [34] and E-RAFT [11], our method fails to achieve superior results on indoor_flying sequences. One possible reason is that model based methods and unsupervised learning methods usually have better generalization ability, while supervised learning methods are limited if training data has a significantly large gap with test data [36, 34]. Since outdoor_day sequences and indoor_flying sequences have a great domain gap, it is challenging for supervised learning methods to generalize well on indoor_flying sequences.

Intra-domain evaluation. To verify the possible cause and better analyze the accuracy of our method under su-

		Input	indoor_flying1		indoor_flying2		indoor_flying3		outdoor_day1	
			EPE	% Outlier	EPE	% Outlier	EPE	% Outlier	EPE	% Outlier
<i>dt = 1</i>										
MB	Brebion et al. [5]	E	0.52	0.10	0.98	5.50	0.71	2.10	0.53	0.20
	MultiCM (Burgers') [27]	E	0.42	0.10	0.60	0.59	0.50	0.28	0.30	0.10
SSL	EV-FlowNet [43]	E	1.03	2.20	1.72	15.10	1.53	11.90	0.49	0.20
	Spike-FlowNet [20]	E	0.84	-	1.28	-	1.11	-	0.49	-
	STE-FlowNet [7]	E	0.57	0.10	0.79	1.60	0.72	1.30	0.42	0.00
USL	Hagenaars et al. [13]	E	0.60	0.51	1.17	8.06	0.93	5.64	0.47	0.25
	Zhu et al. [45]	E	0.58	0.00	1.02	4.00	0.87	3.00	0.32	0.00
	ECN [38]	E	-	-	-	-	-	-	0.30	0.02
SL	E-RAFT [11]	E	1.10	5.72	1.94	30.79	1.66	25.20	0.24	0.00
	DCEiFlow [34]	E + I	0.75	1.55	0.90	2.10	0.80	1.77	0.22	0.00
	Ours	E	1.06	3.63	1.81	27.29	1.58	23.26	0.25	0.07
<i>dt = 4</i>										
MB	MultiCM (Burgers') [27]	E	1.69	12.95	2.49	26.35	2.06	19.03	1.25	9.21
SSL	EV-FlowNet [43]	E	2.25	24.70	4.05	45.30	3.45	39.70	1.23	7.30
	Spike-FlowNet [20]	E	2.24	-	3.83	-	3.18	-	1.09	-
	STE-FlowNet [7]	E	1.77	14.70	2.52	26.10	2.23	22.10	0.99	3.90
USL	Zhu et al. [45]	E	2.18	24.20	3.85	46.80	3.18	47.80	1.30	9.70
	Hagenaars et al. [13]	E	2.16	21.51	3.90	40.72	3.00	29.60	1.69	12.50
SL	DCEiFlow [34]	E + I	2.08	21.47	3.48	42.05	2.51	29.73	0.89	3.19
	E-RAFT [11]	E	2.81	40.25	5.09	64.19	4.46	57.11	0.72	1.12
	Ours	E	2.43	29.91	4.32	52.74	3.60	42.02	0.70	1.08

Table 2: **Evaluation results on MVSEC [44] training with outdoor_day2 sequence.** Model-based (MB) methods need no training data; semi-supervised learning (SSL) methods use grayscale images for supervision; unsupervised learning (USL) methods only require events; and supervised learning (SL) methods need ground-truth. E means the method only uses events as input; and I means the method requires images as input.

<i>dt = 4</i>	Train Set +	indoor_flying1		indoor_flying2		indoor_flying3		outdoor_day1	
		EPE	% Outlier	EPE	% Outlier	EPE	% Outlier	EPE	% Outlier
E-RAFT [11]	indoor_flying1	-	-	2.35	26.24	1.84	18.07	0.72	1.38
		-	-	2.15	20.41	1.68	14.30	0.64	0.98
E-RAFT [11]	indoor_flying2	1.41	8.14	-	-	1.50	9.37	0.71	1.13
		1.32	5.94	-	-	1.50	8.88	0.64	0.96
E-RAFT [11]	indoor_flying3	1.39	8.38	1.77	17.52	-	-	0.72	1.46
		1.33	5.82	1.67	12.80	-	-	0.67	1.10

Table 3: **Evaluation results on MVSEC [44] training with outdoor_day2 and one indoor_flying sequence.** + means the indoor_flying sequence to be included into training set.

perervised learning setting, we incorporate one indoor_flying sequence into training set each time and make comparisons with the supervised learning state-of-the-art E-RAFT. Table 3 shows the evaluation results corresponding to $dt = 4$ grayscale frames (For $dt = 1$, refer to appendix). Compared with the results in Table 2, E-RAFT and our method both achieve great improvement on indoor_flying sequences (EPE of indoor_flying2, ours: 4.32 \rightarrow 1.67, E-RAFT: 5.09 \rightarrow 1.77), confirming the presence of domain gap between

outdoor_day and indoor_flying data. Compared with the accuracy of E-RAFT with the same training set, on the one hand, our method shows better generalization ability, which generates lower EPE and % Outlier without seeing indoor_flying sequences. On the other hand, our method achieves higher accuracy on four sequences by incorporating indoor_flying sequences into training set. Overall, our method achieves a new state-of-the-art result, demonstrating the effectiveness of our proposed method.

splits	EPE	1PE	3PE	Time (ms)
1	0.79	11.91	2.75	16
3	0.76	10.98	2.45	18
<u>5</u>	0.74	10.86	2.30	30
7	0.74	11.13	2.26	50

(a) **Event splitting.** Intermediate motion information contributes.

setup	EPE	1PE	3PE	Time (ms)
<u>basic</u>	0.74	10.86	2.30	30
+v_proj	0.75	10.90	2.35	30
circular	0.76	11.25	2.44	30

(c) **Pattern aggregation module design.**

iterations	EPE	1PE	3PE	Time (ms)
2	0.83	11.04	2.37	10
4	0.76	11.45	2.43	16
<u>6</u>	0.74	10.86	2.30	30
8	0.74	10.81	2.32	54

(e) **Iterations.** TMA achieves good results with a few iterations.

style	EPE	1PE	3PE
w/o	0.78	12.00	2.66
<u>linear</u>	0.74	10.86	2.30
same	0.76	11.45	2.48

(b) **Lookup style.** Sampling correlation volumes with a set of linear coordinates improves the accuracy.

layers	EPE	1PE	3PE	Time (ms)
0	0.75	11.36	2.33	20
<u>1</u>	0.74	10.86	2.30	30
2	0.75	10.90	2.32	38

(d) **Layers of pattern aggregation module.** Stacking 1 pattern aggregation module showcases high accuracy.

radius	EPE	1PE	3PE
1	0.75	11.14	2.37
2	0.76	11.16	2.39
<u>3</u>	0.74	10.86	2.30
4	0.75	10.93	2.30

(f) **Searching radius.** A radius of 3 leads to high accuracy.

Table 4: **TMA ablations.** Models are trained on DSEC-Flow. Settings used in our final model are underlined.

4.3. Ablations

We conduct a series of ablation experiments to validate the proposed improvements. All ablation models are trained on DSEC-Flow and evaluated on the public benchmark.

Event splitting. We compare different numbers of event segments in Table 4a. We choose different splits in 1, 3, 5 and 7. With the split set to 1, our model degrades into a classic correlation volume based framework and exhibits similar accuracy in comparison with E-RAFT. More splits lead to higher accuracy but lower inference efficiency, so we use 5 splits to achieve accuracy-efficiency trade-off.

Lookup style. Based on the assumption that the optical flow changes linearly over a short period, we propose the linear lookup strategy. To validate it, we try to cancel the lookup operation or just use same sampling coordinates on intermediate correlation volumes. As Table 4b shows, sampling the correlation maps with a set of linear coordinates leads to high accuracy.

Pattern aggregation module design. We ablate different components of pattern aggregation module in Table 4c. By default, we set the projection matrix of value W^V as an identity matrix because the motion features are encoded by a motion feature encoder beforehand. In addition, we replace the pairing mode of motion features with circular pairing. Circular pairing gives equal treatment to all motion features without highlighting the importance of the last motion feature, which fails to improve accuracy.

Layers of pattern aggregation module. We compare the accuracy with different layers of pattern aggregation modules in Table 4d. We note that stacking 1 layer of pat-

tern aggregation module showcases high accuracy. More layers may lead to difficulties in optimization, which in turn causes a decrease in accuracy.

Iterations. Iterative learning plays a role in numerous works [18, 30]. Results in Table 4e show that 6 iterative updates reach the sweet spot of accuracy and efficiency.

Searching radius. We compare different radii for looking up correlation volumes in Table 4f. Searching radius of 3 leads to high accuracy. Different from RAFT [30] that achieves better predictions with an increased radius, our model does not benefit a lot from a greater radius.

5. Conclusion

In this paper, we delve into event-based optical flow estimation problem. Different from existing methods that represent the event stream as static frames to adopt the classic correlation volume based pipeline, we contend that the temporal continuity is a vital element of event-based optical flow. We propose a novel Temporal Motion Aggregation (TMA) approach to unlock its potential within a learning-based framework. By incorporating temporally fine-grained motion information, TMA generates high-quality flow estimates at early stages, which not only enables TMA to obtain accurate final predictions, but also greatly reduces the demand for a number of refinements. Extensive experiments on DSEC-Flow and MVSEC datasets verify the effectiveness and superiority of TMA. Meanwhile, we notice the limitations of supervised learning in generalization ability and expect further improvement to extend our idea into unsupervised learning frameworks.

Acknowledgments: This work was supported in part by Shanghai Municipal Science and Technology Major Project under Grant 2018SHZDZX01, ZJ Lab, the Shanghai Center for Brain Science and Brain-Inspired Technology, in part by Shanghai Rising Star Program under Grant 21QC1400900, in part by Tongji-Qomolo Autonomous Driving Commercial Vehicle Joint Lab Project, and in part by Xiaomi Young Talents Program.

References

- [1] Himanshu Akolkar, Sio-Hoi Ieng, and Ryad Benosman. Real-time high speed motion prediction using fast aperture-robust event-driven visual flow. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 44(1):361–372, 2020. [2](#)
- [2] Mohammed Almatrafi, Raymond Baldwin, Kiyoharu Aizawa, and Keigo Hirakawa. Distance surface for event-based optical flow. *IEEE transactions on pattern analysis and machine intelligence*, 42(7):1547–1556, 2020. [2](#)
- [3] Ryad Benosman, Charles Clercq, Xavier Lagorce, Sio-Hoi Ieng, and Chiara Bartolozzi. Event-based visual flow. *IEEE transactions on neural networks and learning systems*, 25(2):407–417, 2013. [2](#)
- [4] Ryad Benosman, Sio-Hoi Ieng, Charles Clercq, Chiara Bartolozzi, and Mandyam Srinivasan. Asynchronous frameless event-based optical flow. *Neural Networks*, 27:32–37, 2012. [2](#)
- [5] Vincent Brebion, Julien Moreau, and Franck Davoine. Real-time optical flow for vehicular perception with low-and high-resolution event cameras. *IEEE Transactions on Intelligent Transportation Systems*, 2021. [7](#)
- [6] Tobias Brosch, Stephan Tschechne, and Heiko Neumann. On event-based optical flow detection. *Frontiers in neuroscience*, 9:137, 2015. [2](#)
- [7] Ziluo Ding, Rui Zhao, Jiyuan Zhang, Tianxiao Gao, Ruiqin Xiong, Zhaofei Yu, and Tiejun Huang. Spatio-temporal recurrent networks for event-based optical flow estimation. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 36, pages 525–533, 2022. [3](#), [5](#), [7](#)
- [8] Guillermo Gallego, Tobi Delbrück, Garrick Orchard, Chiara Bartolozzi, Brian Taba, Andrea Censi, Stefan Leutenegger, Andrew J Davison, Jörg Conradt, Kostas Daniilidis, et al. Event-based vision: A survey. *IEEE transactions on pattern analysis and machine intelligence*, 44(1):154–180, 2020. [1](#)
- [9] Guillermo Gallego, Henri Rebecq, and Davide Scaramuzza. A unifying contrast maximization framework for event cameras, with applications to motion, depth, and optical flow estimation. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 3867–3876, 2018. [2](#), [4](#)
- [10] Mathias Gehrig, Willem Aarents, Daniel Gehrig, and Davide Scaramuzza. Dsec: A stereo event camera dataset for driving scenarios. *IEEE Robotics and Automation Letters*, 6(3):4947–4954, 2021. [2](#), [3](#), [5](#)
- [11] Mathias Gehrig, Mario Millhäusler, Daniel Gehrig, and Davide Scaramuzza. E-raft: Dense optical flow from event cameras. In *2021 International Conference on 3D Vision (3DV)*, pages 197–206. IEEE, 2021. [1](#), [2](#), [3](#), [4](#), [5](#), [6](#), [7](#)
- [12] Mathias Gehrig, Manasi Muglikar, and Davide Scaramuzza. Dense continuous-time optical flow from events and frames. *arXiv preprint arXiv:2203.13674*, 2022. [3](#)
- [13] Jesse Hagenaars, Federico Paredes-Vallés, and Guido De Croon. Self-supervised learning of event-based optical flow with spiking neural networks. *Advances in Neural Information Processing Systems*, 34:7167–7179, 2021. [7](#)
- [14] Markus Hofinger, Samuel Rota Bulò, Lorenzo Porzi, Arno Knapitsch, Thomas Pock, and Peter Kotschieder. Improving optical flow on a pyramid level. In *European Conference on Computer Vision*, pages 770–786. Springer, 2020. [2](#)
- [15] Tak-Wai Hui and Chen Change Loy. Liteflownet3: Resolving correspondence ambiguity for more accurate optical flow estimation. In *European Conference on Computer Vision*, pages 169–184. Springer, 2020. [2](#)
- [16] Tak-Wai Hui, Xiaoou Tang, and Chen Change Loy. Liteflownet: A lightweight convolutional neural network for optical flow estimation. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 8981–8989, 2018. [2](#)
- [17] Tak-Wai Hui, Xiaoou Tang, and Chen Change Loy. A lightweight optical flow cnn—revisiting data fidelity and regularization. *IEEE transactions on pattern analysis and machine intelligence*, 43(8):2555–2569, 2020. [2](#)
- [18] Junhwa Hur and Stefan Roth. Iterative residual refinement for joint optical flow and occlusion estimation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 5754–5763, 2019. [1](#), [2](#), [3](#), [8](#)
- [19] Shihao Jiang, Dylan Campbell, Yao Lu, Hongdong Li, and Richard Hartley. Learning to estimate hidden motions with global motion aggregation. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 9772–9781, 2021. [2](#), [5](#)
- [20] Chankyu Lee, Adarsh Kumar Kosta, Alex Zihao Zhu, Kenneth Chaney, Kostas Daniilidis, and Kaushik Roy. Spikeflownet: event-based optical flow estimation with energy-efficient hybrid neural networks. In *European Conference on Computer Vision*, pages 366–382. Springer, 2020. [7](#)
- [21] Min Liu and Tobi Delbrück. Adaptive time-slice block-matching optical flow algorithm for dynamic vision sensors. *BMVC*, 2018. [2](#)
- [22] Min Liu and Tobi Delbrück. Edflow: Event driven optical flow camera with keypoint detection and adaptive block matching. *IEEE Transactions on Circuits and Systems for Video Technology*, 32(9):5776–5789, 2022. [2](#)
- [23] Ilya Loshchilov and Frank Hutter. Fixing weight decay regularization in adam. 2017. [6](#)
- [24] Federico Paredes-Vallés and Guido CHE de Croon. Back to event basics: Self-supervised learning of image reconstruction for event cameras via photometric constancy. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 3446–3455, 2021. [1](#), [3](#)
- [25] Anurag Ranjan and Michael J Black. Optical flow estimation using a spatial pyramid network. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 4161–4170, 2017. [2](#)

- [26] Olaf Ronneberger, Philipp Fischer, and Thomas Brox. U-net: Convolutional networks for biomedical image segmentation. In *International Conference on Medical image computing and computer-assisted intervention*, pages 234–241. Springer, 2015. [3](#)
- [27] Shintaro Shiba, Yoshimitsu Aoki, and Guillermo Gallego. Secrets of event-based optical flow. In *European Conference on Computer Vision*, pages 628–645. Springer, 2022. [2](#), [5](#), [6](#), [7](#)
- [28] Leslie N Smith and Nicholay Topin. Super-convergence: Very fast training of neural networks using large learning rates. In *Artificial Intelligence and Machine Learning for Multi-Domain Operations Applications*, volume 11006, page 1100612. International Society for Optics and Photonics, 2019. [6](#)
- [29] Deqing Sun, Xiaodong Yang, Ming-Yu Liu, and Jan Kautz. Pwc-net: Cnns for optical flow using pyramid, warping, and cost volume. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 8934–8943, 2018. [1](#), [2](#)
- [30] Zachary Teed and Jia Deng. Raft: Recurrent all-pairs field transforms for optical flow. In *European conference on computer vision*, pages 402–419. Springer, 2020. [1](#), [2](#), [3](#), [5](#), [8](#)
- [31] Stepan Tulyakov, Alfredo Bochicchio, Daniel Gehrig, Stamatios Georgoulis, Yuanyou Li, and Davide Scaramuzza. Time lens++: Event-based frame interpolation with parametric non-linear flow and multi-scale fusion. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 17755–17764, 2022. [1](#)
- [32] Stepan Tulyakov, Daniel Gehrig, Stamatios Georgoulis, Julius Erbach, Mathias Gehrig, Yuanyou Li, and Davide Scaramuzza. Time lens: Event-based video frame interpolation. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 16155–16164, 2021. [1](#)
- [33] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. Attention is all you need. *Advances in neural information processing systems*, 30, 2017. [5](#)
- [34] Zhexiong Wan, Yuchao Dai, and Yuxin Mao. Learning dense and continuous optical flow from an event camera. *IEEE Transactions on Image Processing*, 31:7237–7251, 2022. [1](#), [3](#), [4](#), [6](#), [7](#)
- [35] Haofei Xu, Jiaolong Yang, Jianfei Cai, Juyong Zhang, and Xin Tong. High-resolution optical flow from 1d attention and correlation. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 10498–10507, 2021. [2](#)
- [36] Haofei Xu, Jing Zhang, Jianfei Cai, Hamid Rezaatoughi, and Dacheng Tao. Gmflow: Learning optical flow via global matching. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 8121–8130, 2022. [2](#), [6](#)
- [37] Gengshan Yang and Deva Ramanan. Volumetric correspondence networks for optical flow. *Advances in neural information processing systems*, 32:794–805, 2019. [2](#)
- [38] Chengxi Ye, Anton Mitrokhin, Cornelia Fermüller, James A Yorke, and Yiannis Aloimonos. Unsupervised learning of dense optical flow, depth and egomotion with event-based sensors. In *2020 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 5831–5838. IEEE, 2020. [1](#), [2](#), [7](#)
- [39] Feihu Zhang, Victor Prisacariu, Ruigang Yang, and Philip HS Torr. Ga-net: Guided aggregation net for end-to-end stereo matching. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 185–194, 2019. [2](#)
- [40] Feihu Zhang, Xiaojuan Qi, Ruigang Yang, Victor Prisacariu, Benjamin Wah, and Philip Torr. Domain-invariant stereo matching networks. In *European Conference on Computer Vision*, pages 420–439. Springer, 2020. [2](#)
- [41] Feihu Zhang, Oliver J Woodford, Victor Adrian Prisacariu, and Philip HS Torr. Separable flow: Learning motion cost volumes for optical flow estimation. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 10807–10817, 2021. [2](#)
- [42] Xu Zheng, Yexin Liu, Yunfan Lu, Tongyan Hua, Tianbo Pan, Weiming Zhang, Dacheng Tao, and Lin Wang. Deep learning for event-based vision: A comprehensive survey and benchmarks. *arXiv preprint arXiv:2302.08890*, 2023. [1](#)
- [43] Alex Zhu, Liangzhe Yuan, Kenneth Chaney, and Kostas Daniilidis. Ev-flownet: Self-supervised optical flow estimation for event-based cameras. In *Proceedings of Robotics: Science and Systems*, Pittsburgh, Pennsylvania, June 2018. [2](#), [5](#), [6](#), [7](#)
- [44] Alex Zihao Zhu, Dinesh Thakur, Tolga Özaslan, Bernd Pfrommer, Vijay Kumar, and Kostas Daniilidis. The multivehicle stereo event camera dataset: An event camera dataset for 3d perception. *IEEE Robotics and Automation Letters*, 3(3):2032–2039, 2018. [2](#), [3](#), [5](#), [7](#)
- [45] Alex Zihao Zhu, Liangzhe Yuan, Kenneth Chaney, and Kostas Daniilidis. Unsupervised event-based learning of optical flow, depth, and egomotion. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 989–997, 2019. [1](#), [2](#), [3](#), [4](#), [5](#), [7](#)