

# Perturbed State Space Feature Encoders for Optical Flow with Event Cameras

Gokul Raju Govinda Raju\* Nikola Zubić\* Marco Cannici Davide Scaramuzza  
Robotics and Perception Group, University of Zurich

ggovinda@student.ethz.ch, {zubic, cannici, sdavide}@ifi.uzh.ch

\*Equal contribution

## Abstract

*With their motion-responsive nature, event-based cameras offer significant advantages over traditional cameras for optical flow estimation. While deep learning has improved upon traditional methods, current neural networks adopted for event-based optical flow still face temporal and spatial reasoning limitations. We propose Perturbed State Space Feature Encoders (P-SSE) for multi-frame optical flow with event cameras to address these challenges. P-SSE adaptively processes spatiotemporal features with a large receptive field akin to Transformer-based methods, while maintaining the linear computational complexity characteristic of SSMs. However, the key innovation that enables the state-of-the-art performance of our model lies in our perturbation technique applied to the state dynamics matrix governing the SSM system. This approach significantly improves the stability and performance of our model. We integrate P-SSE into a framework that leverages bi-directional flows and recurrent connections, expanding the temporal context of flow prediction. Evaluations on DSEC-Flow and MVSEC datasets showcase P-SSE's superiority, with 8.48% and 11.86% improvements in EPE performance, respectively.*

## 1. Introduction

Optical flow estimation, a cornerstone of computer vision, plays a crucial role in diverse applications such as object tracking [26, 55], robotic navigation [35, 41], video interpolation [27, 52, 53], and action recognition [40, 48]. Recent deep learning approaches [25, 43, 50] have significantly outperformed traditional energy minimization techniques [24], advancing the field substantially.

However, optical flow estimation with traditional frame-based cameras faces challenges because images are captured at fixed time intervals, leading to temporal gaps, which could affect motion information. To overcome these

limitations, researchers have turned to alternative imaging devices like event-based cameras [10], which provide a continuous, asynchronous stream of data describing changes in relative pixel brightness with high temporal precision.

Despite having access to rich and continuous motion information, existing learning-based optical flow estimation methods that leverage event-based data, [14, 15, 34, 62] are still limited to pair-wise network designs, where only the local neighborhood of events surrounding two consecutive time instants is used to predict the corresponding optical flow. As a result, valuable temporal cues from preceding and successive instants are ignored. On the other hand, recent state-of-the-art methods in frame-based optical flow, such as VideoFlow [43], have demonstrated that a multi-frame approach that optimizes bidirectional flows across multiple frames leads to superior performance, exploiting richer latent information within extended frame sequences.

While VideoFlow [43] demonstrates improved temporal reasoning, it faces two significant limitations. First, the inherent temporal resolution constraints of traditional cameras hinder their ability to capture fine-grained motion details. Second, its reliance on the Twins Transformer [4] architecture for encoding introduces spatial limitations. The lack of a global effective receptive field in this encoder restricts the network's capacity to capture extensive spatial dependencies. Consequently, these constraints potentially impair VideoFlow's ability to effectively handle scenarios involving occluded objects or large-scale motions, highlighting areas where further advancements could yield substantial improvements in optical flow estimation.

Drawing inspiration from recent advancements in frame-based optical flow techniques, we present a novel multi-event optical flow estimation approach. Our model departs significantly from conventional event-based methods through two key innovations. First, we implement bi-directional optical flow prediction and optimization across up to five consecutive time instants. This allows us to recursively propagate motion information between adjacent

frames, substantially expanding the network’s temporal receptive field. To address the spatial limitations of existing methods, we introduce the Perturbed State Space Feature Encoders (P-SSE) to optical flow estimation—a first in this domain. This novel encoder is designed to capture more comprehensive spatial dependencies, improving the model’s ability to handle complex scenarios. These innovations combine to leverage extended temporal context and enriched spatial information, significantly improving optical flow estimation accuracy and robustness.

The key innovation of P-SSE lies in the novel perturbation technique applied to the state dynamics matrix. While P-SSE builds upon the strengths of Transformers and Convolutional Neural Networks (CNNs), incorporating dynamic weights and global receptive fields characteristic of models like DeiT [51], its distinguishing feature is the perturbation method that crucially improves its performance and stability. The perturbation technique is the key idea of P-SSE’s capabilities, making the state dynamics matrix of the State Space Model (SSM) more robust to noise and significantly stabilizing the entire system. By applying a carefully designed perturbation followed by diagonalization, we address the inherent instability issues often associated with SSMs, particularly when processing complex, asynchronous event data. P-SSE utilizes the scanning procedure of the VMamba model [33]. Still, our perturbation approach regularizes the state matrix and improves stability, allowing the model to capture the spatial context in event data more effectively. This improved spatial reasoning, combined with temporal insights from our VideoFlow-inspired [43] multi-event strategy, results in a comprehensive spatiotemporal approach to optical flow estimation. This innovative combination of perturbed spatial processing and multi-event temporal processing allows our model to capture complex motion patterns with strong accuracy and robustness, setting new standards in event-based optical flow estimation.

We conducted a comprehensive evaluation of our model on the DSEC-Flow [13] and MVSEC [63] datasets, revealing several contributions:

- Our experiments demonstrate the successful adaptation of multi-frame optical flow strategies to event-based vision, resulting in substantial performance improvements.
- The global receptive field and perturbation technique provided by our proposed P-SSE, combined with the model’s capacity to leverage rich temporal motion information, significantly improves performance in challenging scenarios characterized by motion blur and occlusions.
- On the DSEC-Flow benchmark [13], our method achieves a top EPE score of 0.680, representing an 8.48% improvement over the previous best TMA model [32], and a 9.33% improvement compared to the BFlow [15] model. On the MVSEC dataset [63], we observe an 11.86% increase in EPE performance.

## 2. Related Work

**Event-based optical flow** Event-based optical flow estimation has evolved through several distinct phases, each adapting to the unique characteristics of event data. Initial approaches focused on modifying classical computer vision techniques, including adaptations of the Lucas-Kanade algorithm [1, 12, 38], local plane fitting strategies [2, 36], and methods recovering pixel trajectories via contrast maximization on the image plane [9]. The success of learning-based approaches in frame-based optical flow catalyzed a shift towards adapting these architectures for event processing. This transition necessitated the development of structured representations of raw event data [14, 32, 67], capable of encapsulating rich motion details. Concurrently, unsupervised methods emerged, estimating sparse optical flow by minimizing photometric loss over warped images [62] or event-frames [59], often incorporating additional smoothness constraints [61] or adversarial losses [28]. Recent advancements have prioritized dense predictions, aiming to provide reliable estimates even in event-free regions. ERAFT [14] pioneered this approach by adapting the RAFT architecture [50], utilizing correlation volumes and iterative refinement for forward optical flow computation in event data. Building on this foundation, TMA [32] introduced a pattern aggregation mechanism to enhance motion features and improve convergence. Concurrently, ADMFlow [34] addressed the challenge of varying event densities through an adaptive event representation, significantly improving robustness. Despite these advancements, current event-based optical flow methods face limitations in both temporal and spatial reasoning. The prevalent reliance on pairwise predictions restricts the temporal context. In contrast, using encoders with limited spatial receptive fields that cannot stably encode rich contexts constrains the ability to capture and encode global motion patterns. Our research diverges from previous approaches by adopting a VideoFlow-inspired [43] architecture to exploit the rich temporal information inherent in multi-event sequences, enabling precise bidirectional optical flow predictions. This capability is further supported by integrating P-SSE, which efficiently captures long-range spatial dependencies via global context. Most recent event-based vision works include IDNet [57], which employs a lightweight event representation and recurrent refinement for robust flow predictions. Also, ECDDP [58] applies a cross-scale dynamic propagation approach to capture finer spatiotemporal details. Although these methods were released after our initial submission, we report better results and include a discussion and comparison in our results section to ensure a comprehensive evaluation of recent state-of-the-art.

**State Space Models** State Space Models (SSMs) represent a novel class of sequence models in deep learning, bridging RNNs, CNNs, and classical state space models [18, 19, 45, 46, 69, 70]. Inspired by continuous state space models from control systems and utilizing HiPPO initialization [17], LSSL [18] showed promise in addressing long-range dependencies. However, LSSL’s practical application was limited by its computational and memory intensity. S4 [19] addressed this by introducing parameter normalization into a diagonal format, enabling the development of various deep SSM architectures. These include complex-diagonal models [20, 21], multi-input multi-output support [45], diagonal plus low-rank decompositions [23], and selective mechanisms like Mamba [16]. This evolution of SSMs has significantly expanded their applicability and efficiency in handling complex sequential data, making them particularly relevant for our work in event-based optical flow estimation.

In the domain of visual applications, S4ND [37] pioneered using SSMs. However, its direct extension of the S4 framework fell short in dynamically capturing image-specific details, largely due to its time-invariant nature. This limitation also manifests in the initial implementation of SSMs for event-based vision [68], which employed a similar recurrent, time-invariant SSM approach for object detection. While drawing inspiration from Vim [65], and VMamba [33], our work goes beyond their capabilities by introducing a more stable system state through the novel perturbation technique. This advancement enables P-SSE to encode the spatial context of motion patterns more effectively and over longer sequences.

### 3. Methodology

Our approach is influenced by the latest advances in multi-frame optical flow estimation for image datasets [8, 42], particularly the VideoFlow model [43], which we adapt for event-based vision. Explained in Section 3.2, we exploit VideoFlow’s TRi-frame Optical Flow (TROF) module, which operates on three frames, and its extension, the MOtion Propagation (MOP) module, which interconnects multiple TROF over five frames. While traditional models have focused on analyzing pairs of consecutive event representations to compute optical flows, leveraging the high temporal resolution of event data, they often overlook the valuable temporal insights available from adjacent events. Inspired by VideoFlow, our architecture employs a longer temporal context spanning up to five consecutive event representations, exploiting the rich motion information inherent in event data, thereby improving the precision of our optical flow estimations. Moreover, as explained in Section 3.3, the spatial encoding in our architecture benefits from a global effective receptive field at linear computational expense thanks to the P-SSE. With its novel perturbation technique, the P-SSE plays a crucial role in our methodology

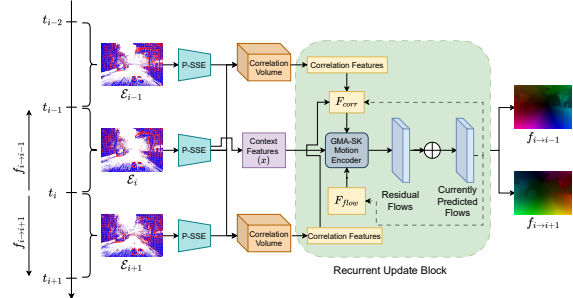


Figure 1. Illustrative diagram of our E-TROF model. Utilizing three successive event representations, E-TROF employs bidirectional correlation, context, and flow features for recurrently estimating bidirectional optical flows.

by significantly improving the stability and performance of the state dynamics matrix, enabling our model to capture complex spatial-temporal dependencies in event data with consistent improvements in accuracy and efficiency.

#### 3.1. Problem Formulation for Event-based Vision

Optical flow estimation aims at computing a displacement field  $\mathbf{f}_{i \rightarrow i+1} : \mathbb{I}^{H \times W \times 2} \rightarrow \mathbb{R}^{H \times W \times 2}$  mapping each pixel’s coordinates from an initial timestamp  $t_i$  to their new positions at a subsequent timestamp  $t_{i+1}$ , which we refer to as the forward optical flow. Conversely, the backward optical flow,  $\mathbf{f}_{i+1 \rightarrow i}$ , maps pixel positions from  $t_{i+1}$  back to  $t_i$ .

#### 3.2. Model Architecture

Denote the stream of events captured by an event-based camera between two instants  $t_h$  and  $t_k$  as  $\mathcal{S}(t_h, t_k) = \langle \langle \mathbf{e}_i | t_i \in [t_h, t_k] \rangle \rangle$ . Each event is represented by a tuple  $\mathbf{e}_i = (x_i, y_i, t_i, p_i)$ , where  $(x_i, y_i)$  specifies the pixel’s location, the time of the event  $t_i$ , and the polarity  $p_i$ , indicating whether the brightness increased or decreased.

Given a sequence of successive timestamps  $\{t_{i-2}, t_{i-1}, t_i, t_{i+1}, t_{i+2}\}$  among which to estimate the optical flow, we first aggregate the events triggered in between these timestamps into a sequence of image-like representations  $\mathcal{E}_i \in \mathbb{R}^{H \times W \times C}$ , each one aggregating the events  $\mathcal{S}(t_{i-1}, t_i)$  into  $C$ -dimensional pixel-wise features. In this work, we employ ERGO-12 [67] and time surfaces [44] due to their favorable balance between complexity and performance [67]. Then, following [43], we perform hierarchical recurrent processing, where we first process these representations in overlapping triplets  $\{\mathcal{E}_{i-2}, \mathcal{E}_{i-1}, \mathcal{E}_i\}$ ,  $\{\mathcal{E}_{i-1}, \mathcal{E}_i, \mathcal{E}_{i+1}\}$  and  $\{\mathcal{E}_i, \mathcal{E}_{i+1}, \mathcal{E}_{i+2}\}$  using the Event-Triplet Optical Flow (E-TROF) module, and then interconnect them through the Event-Motion Propagation (E-MOP) module.

### 3.2.1. E-TROF (Event-Triplet Optical Flow)

As depicted in Figure 1, the E-TROF module takes a triplet  $\mathcal{E}_{i-1}$ ,  $\mathcal{E}_i$ , and  $\mathcal{E}_{i+1}$  and predicts the bidirectional optical flows  $\mathbf{f}_{i \rightarrow i+1}$  and  $\mathbf{f}_{i \rightarrow i-1}$  centered around  $t_i$ . The process starts with a feature encoding phase, where event representations are transformed into feature embeddings through the proposed P-SSE model, explained in detail in Section 3.3, which extracts features at 1/8th the original resolution of the event data. Following this, we construct forward and backward correlation volumes,  $\mathbf{Corr}_{i,i+1}$  and  $\mathbf{Corr}_{i,i-1}$ , encoding the similarities between pixels across frames, to serve as the basis for deducing both forward and backward optical flows. The forward correlation volume is derived from  $\mathcal{E}_i$  and  $\mathcal{E}_{i+1}$ , while the backward volume from  $\mathcal{E}_i$  and  $\mathcal{E}_{i-1}$ . Context features  $\mathbf{x}$ , encoded from  $\mathcal{E}_i$  using the same exact encoder, complement these correlation volumes.

Subsequently, we utilize the forward and backward correlation features to derive the combined bidirectional correlation features. This is achieved through an interactive process, similar to RAFT [50], where forward and backward flows are initialized as zero-displacements,  $\mathbf{f}_{i \rightarrow i+1}^0 = \mathbf{0}$  and  $\mathbf{f}_{i \rightarrow i-1}^0 = \mathbf{0}$ , and then progressively refined over  $K$  iterations:

$$\mathbf{c}_{i \rightarrow i-1}^k = \mathbf{Corr}_{i,i-1}(\mathbf{f}_{i \rightarrow i-1}^k), \quad (1)$$

$$\mathbf{c}_{i \rightarrow i+1}^k = \mathbf{Corr}_{i,i+1}(\mathbf{f}_{i \rightarrow i+1}^k), \quad (2)$$

$$\mathbf{F}_{corr}^k = \text{CorrelationFusion}(\mathbf{c}_{i \rightarrow i-1}^k, \mathbf{c}_{i \rightarrow i+1}^k), \quad (3)$$

$$\mathbf{F}_{flow}^k = \text{FlowFusion}(\mathbf{f}_{i \rightarrow i-1}^k, \mathbf{f}_{i \rightarrow i+1}^k). \quad (4)$$

where correlation values  $\mathbf{c}^k$  are obtained by sampling the corresponding correlation volumes based on the predicted flows, while flow features  $\mathbf{F}_{corr}^k$  and  $\mathbf{F}_{flow}^k$  through learnable encoders. Following the fusion step, we compute the motion features using the Super Kernel Motion Encoder [43], which takes as input the fused bidirectional correlation  $\mathbf{F}_{corr}^k$ , the flow features  $\mathbf{F}_{flow}^k$ , and an additional feature map  $\mathbf{F}_{map}^k$ , described later in Section 3.2.2, which carries long-range motion features from adjacent triplets:

$$\mathbf{F}_{motion}^k, \mathbf{M}_i^{k+1} = \text{SK-MotionEncoder}(\mathbf{F}_{corr}^k, \mathbf{F}_{flow}^k, \mathbf{F}_{map}^k), \quad (5)$$

where  $\mathbf{M}_i^{k+1}$  is a feature encoding the motion state at the center instant  $t_i$ .

The flow refinement process is analogous to RAFT [50], but carried over for both the forward and backward flow simultaneously, leveraging a hidden state feature  $\mathbf{h}^k$ , initialized with the context feature  $\mathbf{x}$ :

$$\mathbf{h}^{k+1} = \text{SK-UpdateBlock}(\mathbf{F}_{motion}, \mathbf{x}, \mathbf{h}^k), \quad (6)$$

$$\Delta \mathbf{f}_{i \rightarrow i-1, i \rightarrow i+1}^k = \text{FlowDecoder}(\mathbf{h}^{k+1}), \quad (7)$$

$$\mathbf{f}_{i \rightarrow i-1, i \rightarrow i+1}^{k+1} = \mathbf{f}_{i \rightarrow i-1, i \rightarrow i+1}^k + \Delta \mathbf{f}_{i \rightarrow i-1, i \rightarrow i+1}^k \quad (8)$$

The flow estimates are similarly scaled because the feature embeddings were initially computed at a resolution reduced to 1/8th of the original event data. The final step involves upscaling the bidirectional flow predictions to match the original resolution of the event data.

### 3.2.2. E-MOP (Event-Motion Propagation)

The E-MOP propagates motion information across triplets of event frames, enabling the exchange of motion state information between them to collaboratively refine the bidirectional optical flows estimated for each E-TROF segment, as illustrated in Figure 2. To do so, a motion state feature  $\mathbf{M}_i^k$ , initially set to a random value, is computed from each triplet of event frames and then progressively refined. A bidirectional temporal motion information is extracted by warping the motion state features from adjacent TROFs,  $\mathbf{M}_{i-1}$  and  $\mathbf{M}_{i+1}$ , according to the predicted optical flows and the current TROF's motion state feature.

$$\mathbf{m}_f^k = \text{Warp}(\mathbf{M}_{i+1}^k, \mathbf{f}_{i \rightarrow i+1}^k), \quad (9)$$

$$\mathbf{m}_b^k = \text{Warp}(\mathbf{M}_{i-1}^k, \mathbf{f}_{i \rightarrow i-1}^k), \quad (10)$$

$$\mathbf{F}_{mop}^k = \text{Concatenate}(\mathbf{M}_i, \mathbf{m}_f, \mathbf{m}_b). \quad (11)$$

Following this, the estimated motion features are first used to extract motion features as in Equation (5) and then used to refine the bidirectional flows iteratively through Equation (8). The continuous update of the motion state feature allows each event triplet to benefit from the temporal insights of adjacent triplets, enhancing the temporal depth and exploiting evolving temporal cues through iterative refinement.

### 3.3. Perturbed State Space Feature Encoder (P-SSE)

Our P-SSE encoder leverages state space models to achieve an effective global receptive field while maintaining linear complexity. At the heart of our encoder is a linear system that maps a 1D input signal of length  $L$ ,  $x(t) \in \mathbb{R}^L$ , to an output response  $y(t) \in \mathbb{R}^L$  through a hidden state  $h(t) \in \mathbb{R}^N$ :

$$\begin{aligned} h'(t) &= \mathbf{A}h(t) + \mathbf{B}x(t), \\ y(t) &= \mathbf{C}h(t) + \mathbf{D}x(t), \end{aligned} \quad (12)$$

with evolution parameter matrix  $\mathbf{A} \in \mathbb{C}^{N \times N}$ , projection parameters  $\mathbf{B}, \mathbf{C} \in \mathbb{C}^N$ , defined for a state dimension  $N$ , alongside a skip connection parameter  $\mathbf{D} \in \mathbb{C}^1$ .

Analogous to previous works [16, 65], our encoder implements a discretized version of this continuous system that aligns it to the digital nature of the input data [18]. Given a discrete input  $x_k \in \mathbb{R}^{L \times D}$ , the discretization of Eq. 12 under the zeroth-order hold (ZOH) method using a

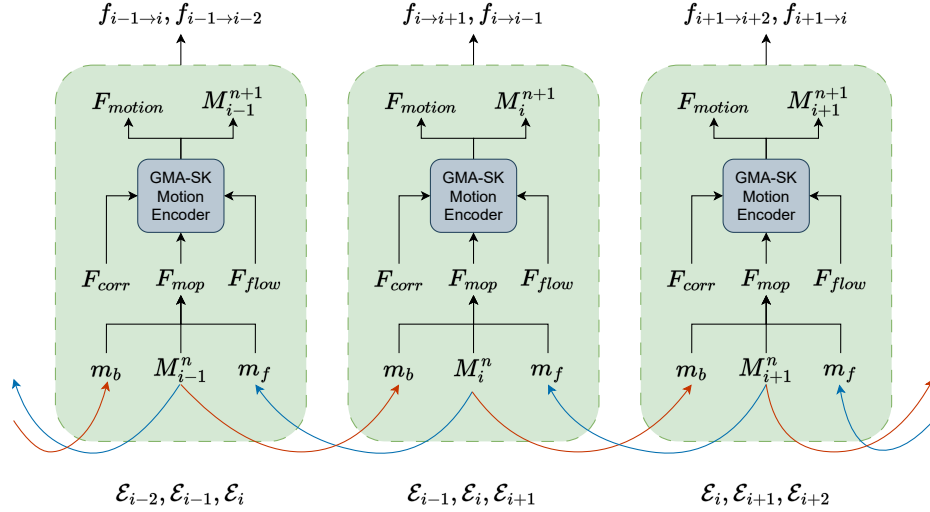


Figure 2. Diagram of the E-MOP model, which integrates 3 E-TROFs for 5 consecutive event representations to predict and refine bidirectional optical flows by sharing dynamic temporal motion information among adjacent TROFs.

step size  $\Delta$  yields:

$$\begin{aligned}
 \bar{\mathbf{A}} &= e^{\Delta \mathbf{A}}, \quad \bar{\mathbf{B}} \approx (\Delta \mathbf{A})(\Delta \mathbf{A})^{-1} \Delta \mathbf{B} = \Delta \mathbf{B}, \quad \bar{\mathbf{C}} = \mathbf{C} \\
 h_k &= \bar{\mathbf{A}} h_{k-1} + \bar{\mathbf{B}} x_k, \\
 y_k &= \bar{\mathbf{C}} h_k + \bar{\mathbf{D}} x_k,
 \end{aligned} \tag{13}$$

where we consider the first-order Taylor approximation for  $\bar{\mathbf{B}}$  defined in [16]. In S6 [16] and our case, the matrices  $\mathbf{B} \in \mathbb{R}^{B \times L \times N}$ ,  $\mathbf{C} \in \mathbb{R}^{B \times L \times N}$ , and  $\Delta \in \mathbb{R}^{B \times L \times D}$  are dynamically derived from the input data  $x \in \mathbb{R}^{B \times L \times D}$ .

While classical SSM formulations handle 1D sequences, we follow the approach of Vim [65] (and VMamba [33]) to process 2D feature maps by scanning row-wise (or column-wise). Specifically, at each encoder block, we reshape the  $H \times W \times C$  feature map into a sequence of length  $H \times W$  and channel dimension  $C$ , apply the SSM transformation, and then reshape back. Repeated blocks with positional embeddings ensure the model can capture 2D global context despite operating on 1D scans.

P-SSE is constructed by repeating the block structure in a ViT fashion, similar to [68]. These features constitute the output of our P-SSE encoders, which we use to extract correlation and context features as described in the Sec. 3.2. To ensure the stability of the training, we first learn a perturbation matrix for the state matrix on the ImageNet dataset [5], and then use it, along with a perturb-then-diagonalize (PTD) technique, during training in the same fashion as [60]. However, they tested the model only on synthetic tasks. We use the  $L1$  loss between the predicted and ground truth bidirectional optical flows to supervise the training of our model [43, 66].

## 4. Explanation of Perturb-then-Diagonalize (PTD) Technique

### 4.1. Preliminaries

The Perturb-then-Diagonalize (PTD) technique [60] aims to enhance the performance and stability of SSMs by making the state dynamics matrix  $\mathbf{A}$  more robust to data noise, surpassing even the HiPPO theory-based matrix initialization [17]. This method is particularly effective for high-dimensional data and complex dynamics, as encountered in our case. The PTD method comprises two primary steps applied to the state dynamics matrix  $\mathbf{A}$  of an SSM. In the first step, the system matrix  $\mathbf{A}$  is perturbed by a small matrix  $\mathbf{E}$ , yielding a new matrix  $\mathbf{A}^*$ . The perturbation matrix  $\mathbf{E}$  is carefully chosen to enhance system properties like stability and robustness without significantly altering its original behavior. Following [60], we set the magnitude of  $\mathbf{E}$  to approximately 10% of  $\mathbf{A}$ 's magnitude. This process is mathematically represented as:

$$\mathbf{A}^* = \mathbf{A} + \mathbf{E}, \tag{14}$$

where  $\mathbf{E}$  is typically small in magnitude compared to  $\mathbf{A}$ . We constrain the size of the perturbation matrix to 10% of the size of the HiPPO [17] matrix  $\mathbf{A}$ . Following perturbation, the perturbed matrix  $\mathbf{A}^*$  is diagonalized.

### 4.2. Application of the Technique to Our Problem

Incorporating pre-training in deep learning frameworks improved the model performance [67], particularly in tasks requiring high-level feature recognition. Our work leverages the PTD [60] technique to pre-train a model on Ima-

geNet [5], a large-scale image classification dataset, focusing specifically on the learned  $\mathbf{A}$  matrix.

We initiate by pre-training our model using the PTD [60] technique for SSMs on ImageNet [5]. PTD introduces controlled perturbations in the training dynamics, enhancing the model’s ability to escape local minima and improving generalization. We modify the loss function  $L_{\text{PTD}}(A)$  for the  $\mathbf{A}$  matrix as follows:

$$L_{\text{PTD}}(A) = L(A) + \lambda \cdot |\mathbf{E}|, \quad (15)$$

where  $|\mathbf{E}|$  represents the magnitude of perturbations in the model’s parameters  $\theta$ , and  $\lambda$  controls these perturbations’ strength. As previously stated, the perturbation magnitude is approximately 10% of the  $\mathbf{A}$  matrix magnitude. Perturbations  $\mathbf{E}$  are sampled from a Gaussian distribution [60], encouraging exploration of a wider parameter space during training.

Post-ImageNet [5] pre-training, we utilize the state dynamics matrix  $\mathbf{A}^*$  learned within this PTD framework to initialize our state dynamics matrix. This initialization is then used to train the model on the DSEC dataset [13] and perform a zero-shot evaluation on the MVSEC dataset [63].

## 5. Experiments

**Datasets and Evaluation Metric.** We train our models on the DSEC-Flow dataset [13] and evaluate them on both DSEC and MVSEC [63] (zero-shot) datasets. DSEC-Flow provides 7800 training instances, each with bidirectional optical flow ground-truth. In Table 1, we report results obtained from the official DSEC-Flow server. Notably, the raw training event data is initially distorted and not rectified, necessitating pre-processing to rectify events before model training, ensuring compatibility with the rectified event representations used in benchmark tests.

MVSEC [63] comprises 5 driving sequences and 4 indoor drone-collected sequences, with ground truth derived from lidar odometry and mapping. Following [56], we focus primarily on the *indoor\_flying1*, *indoor\_flying2*, and *indoor\_flying3* sequences to demonstrate our model’s generalization capabilities on out-of-distribution data. In accordance with benchmark protocols [13], we primarily use End-Point Error (EPE) to evaluate optical flow accuracy, calculating the mean L2 distance between predicted and actual flow vectors. The DSEC benchmark additionally employs Angular Error (AE) and N-pixel Error (NPE), the latter quantifying the proportion of pixels where optical flow error exceeds a given threshold  $N \in \{1, 2, 3\}$ .

**Implementation Details.** Our models are implemented using PyTorch. We primarily utilize the ERGO-12 representation [67] with 12 bins. Feature and context embeddings are generated using our proposed P-SSE. For motion feature estimation and bidirectional flow learning, we employ the GMA-SK update block [49]. Both E-TROF and

Table 1. Performance comparison on the DSEC-Flow Public Benchmark [13]. 1PE, 2PE, 3PE are the N-pixel errors. The top result in each column is in bold.

Method	1PE	2PE	3PE	EPE	AE
E-RAFT [14]	12.742	4.740	2.684	0.788	2.851
ADMFlow [34]	12.522	4.673	2.647	0.779	2.838
E-FlowFormer [30]	11.225	4.102	2.446	0.759	2.676
BFlow [15]	11.901	4.411	2.440	0.750	2.680
TMA [32]	10.863	3.972	2.301	0.743	2.684
IDNet [57]	10.111	3.523	2.018	0.723	2.724
ECDDP [58]	<b>8.887</b>	<b>3.199</b>	1.958	0.697	2.575
<b>Ours (P-SSE)</b>	9.144	3.232	<b>1.816</b>	<b>0.680</b>	<b>2.560</b>

E-MOP architectures undergo 250k training iterations with batch size 4. We use the AdamW optimizer with a  $4 \times 10^{-4}$  learning rate and a OneCycle LR Scheduler. We apply data augmentation techniques (horizontal/vertical flips, random cropping) for robustness. Unless otherwise stated, training and inference are conducted on either A100 or V100 GPUs.

### 5.1. DSEC-Flow Results

In Table 1, we compare our approach to a range of existing event-based optical flow methods, including TMA [32], BFlow [15], IDNet [57], and ECDDP [58]. Although IDNet and ECDDP were published concurrently or shortly after our initial development, we highlight their performance for completeness. As shown, our P-SSE achieves an EPE of 0.68, which is lower than IDNet’s 0.723 and ECDDP’s 0.697, marking an improvement of approximately 6% compared to IDNet and about 2.4% compared to ECDDP. Additionally, our method outperforms TMA (0.743 EPE) and BFlow (0.750 EPE). Our approach is also the only one in Table 1 that integrates a multi-event design within a perturbed state space model.

Fig. 3 demonstrates how our model excels in predicting flow for a street sign moving partially out of frame, where E-RAFT [14] struggles. Similarly, Fig. 4 shows that P-SSE reliably tracks a car even when it is partially occluded. We attribute these improvements to the stability of our state matrix (enabled by perturbation), which helps effectively leverage temporal information across multiple events. Additionally, Section 5.3 discusses the resource efficiency of our model in more detail, demonstrating that our multi-event design does not significantly inflate training/inference costs compared to lighter models such as IDNet.

### 5.2. MVSEC Results

Table 2 presents performance on the MVSEC [63] *indoor\_flying* sequences under the protocol of [56] (with  $dt = 1$ ). To test generalization, we train on DSEC and evaluate on MVSEC without fine-tuning, following [56]. This set-

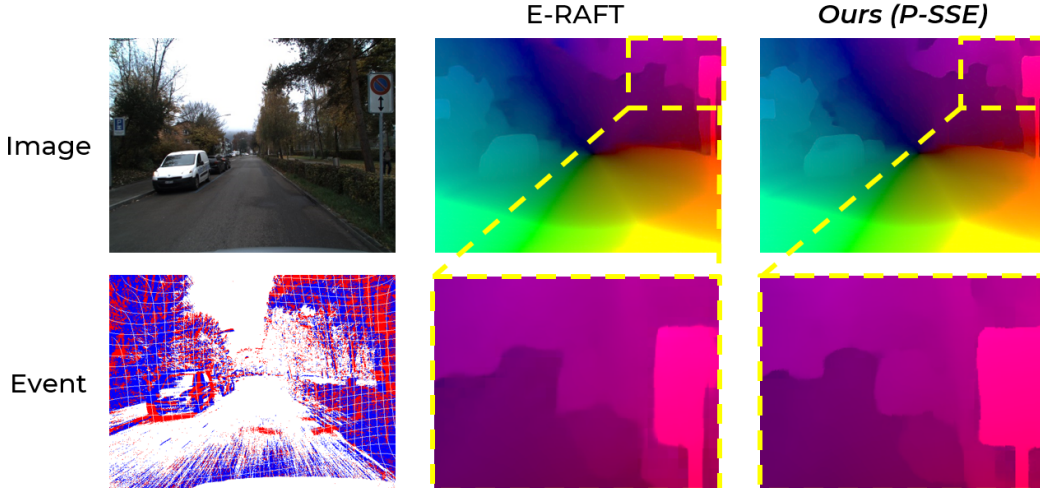


Figure 3. Illustration of P-SSE’s efficacy in handling out-of-boundary regions, in comparison with E-RAFT [14].

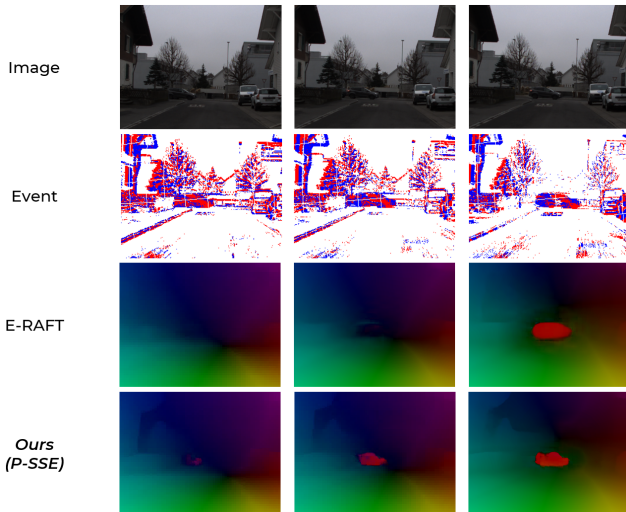


Figure 4. Demonstration of P-SSE’s capability in managing partially occluded scenes, compared with E-RAFT [14]. The sequence progresses from left to right, showcasing frames from a DSEC-Flow test sequence.

ting is challenging due to substantial domain differences: DSEC covers driving scenarios with a DAVIS346 sensor, while MVSEC’s indoor sequences use a DAVIS240. Despite these disparities, our network outperforms existing methods and achieves a remarkable 11.86% improvement over the previous best results.

### 5.3. Case Study and Analysis

**Train and Inference Efficiency.** Table 3 compares our optical flow prediction network against other state-of-the-art architectures, including both image-based (RAFT [50]) and

Table 2. EPE results on MVSEC. “in\_fly” = “indoor\_flying”. We report results for both *Base* and *Small* variants of P-SSE using either ERGO-12 or Time Surface representations.

Methods	<i>in_fly1</i>	<i>in_fly2</i>	<i>in_fly3</i>	<i>EPE</i>	
EV-FlowNet [62]	1.03	1.72	1.53	1.43	
Zhu et al. [64]	0.58	1.02	0.87	0.82	
EST [11]	0.97	1.38	1.43	1.26	
Matrix-LSTM [3]	0.82	1.19	1.08	1.03	
Spike-FlowNet [29]	0.84	1.28	1.11	1.08	
Stoffregen et al. [47]	0.56	0.66	0.59	0.60	
Paredes et al. [39]	0.79	1.40	1.18	1.12	
LIF-EV-FlowNet [22]	0.71	1.44	1.16	1.10	
Deng et al. [6]	0.89	0.66	1.13	0.89	
Li et al. [31]	0.59	0.64	-	0.62	
STE-FlowNet [7]	0.57	0.79	0.72	0.69	
DCEIFlow [56]	0.56	0.64	0.57	0.59	
<b>Ours (P-SSE):</b>					
ERGO-12	Base	<b>0.47</b>	0.59	<b>0.50</b>	<b>0.52</b>
	Small	0.52	0.63	0.56	0.57
Time Surf	Base	<b>0.47</b>	0.58	<b>0.50</b>	<b>0.52</b>
	Small	0.48	<b>0.57</b>	0.51	<b>0.52</b>

event-based (TMA [32], BFlow [15]) approaches. All measurements are taken on an NVIDIA Quadro RTX 8000 with batch size 1, averaged over 100 runs. Our full model (*P-SSE<sub>base</sub>*) achieves an inference time of 111.2 ms, while a *small* (4.12 M parameters) and *nano* (2.99 M) variant run even faster at 100.4 ms and 95.1 ms respectively, maintaining competitive accuracy.

**Perturb-Then-Diagonalize Strategy.** In Table 4, we evaluate the impact of the PTD [60] technique versus the standard HiPPO [17] initialization. We observe that PTD yields both higher training stability (fewer spikes in validation loss) and

Table 3. Efficiency Comparison on Quadro RTX 8000 GPU.

Methods	Inference Time [ms]	# of Params [M]	VRAM [GB]
RAFT [50]	226.4	5.26	<b>4.35</b>
TMA [32]	201.4	6.88	10.03
BFlow [15]	133.2	5.64	7.93
<b>(Ours) P-SSE<sub>base</sub></b>	111.2	6.01	12.0
<b>(Ours) P-SSE<sub>small</sub></b>	100.4	4.12	10.0
<b>(Ours) P-SSE<sub>nano</sub></b>	<b>95.1</b>	<b>2.99</b>	8.3

Table 4. Comparison of methods using PTD (vs. HiPPO) for EPE on DSEC-Flow.

Methods	Initialization	EPE
P-SSE <sub>Small</sub>	PTD [60]	0.691
	HiPPO [17]	0.713
P-SSE <sub>Base</sub>	PTD [60]	0.680
	HiPPO [17]	0.705

a consistent reduction in final EPE (3.09% improvement for *small*, 2.41% for *base*). These results suggest that small perturbations to the state matrix improve the convergence properties of our SSM-based encoder and are especially beneficial for event-based optical flow, which typically exhibits sparse and highly asynchronous inputs.

## 6. Conclusion

This work introduces a novel approach to optical flow estimation using event-based cameras, significantly improving the temporal and spatial reasoning crucial for accurate and efficient flow computation. P-SSE improves spatial feature encoding by combining the global receptive field of Transformers [54] with the efficiency of SSMs. We implement the PTD technique [60] in our P-SSE to ensure model stability, effectively addressing the challenges of training deep networks on complex event-based data. Inspired by recent advances in frame-based optical flow [43], we adopt a multi-frame approach that utilizes bi-directional flows and propagates motion information across multiple time instants. This extension of the temporal receptive field significantly enriches the model’s motion understanding.

Empirical evaluations on DSEC-Flow [13] and MVSEC [63] datasets demonstrate the superiority of our approach. We not only establish new benchmarks in optical flow estimation accuracy but also achieve substantial improvements in computational efficiency. The notable reductions in training time and resource requirements, coupled with faster inference, underscore the practicality and applicability of our method in real-world scenarios, particularly in domains demanding rapid and reliable motion estimation.

## 7. Acknowledgment

This work was supported by the European Research Council (ERC) under grant agreement No. 864042 (AGILE-FLIGHT).

## References

- [1] 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
- [2] 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
- [3] Marco Cannici, Marco Ciccone, Andrea Romanoni, and Matteo Matteucci. A differentiable recurrent surface for asynchronous event-based data. In *ECCV*, pages 136–152, 2020. 7
- [4] Xiangxiang Chu, Zhi Tian, Yuqing Wang, Bo Zhang, Haibing Ren, Xiaolin Wei, Huaxia Xia, and Chunhua Shen. Twins: Revisiting the design of spatial attention in vision transformers. In *Neural Inf. Process. Syst.*, 2021. 1
- [5] Jia Deng, Wei Dong, Richard Socher, Li-Jia Li, Kai Li, and Li Fei-Fei. Imagenet: A large-scale hierarchical image database. In *CVPR*, pages 248–255, 2009. 5, 6
- [6] Yongjian Deng, Hao Chen, Huiying Chen, and Youfu Li. Learning from images: A distillation learning framework for event cameras. *IEEE TIP*, 2021. 7
- [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 *AAAI*, 2022. 7
- [8] Fisseha Admasu Ferede and Madhusudhanan Balasubramanian. Sstm: Spatiotemporal recurrent transformers for multi-frame optical flow estimation. *Neurocomputing*, 558: 126705, 2023. 3
- [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 *CVPR*, pages 3867–3876, 2018. 2
- [10] Guillermo Gallego, Tobi Delbrück, Garrick Orchard, Chiara Bartolozzi, Brian Taba, Andrea Censi, Stefan Leutenegger, Andrew J. Davison, Jörg Conradt, Kostas Daniilidis, and Davide Scaramuzza. Event-based vision: A survey. *IEEE TPAMI*, 2020. 1
- [11] Daniel Gehrig, Antonio Loquercio, Konstantinos G Derpanis, and Davide Scaramuzza. End-to-end learning of representations for asynchronous event-based data. In *CVPR*, pages 5633–5643, 2019. 7
- [12] Daniel Gehrig, Henri Rebecq, Guillermo Gallego, and Davide Scaramuzza. Eklt: Asynchronous photometric feature tracking using events and frames. *IJCV*, 128(3):601–618, 2020. 2
- [13] Mathias Gehrig, Willem Aarents, Daniel Gehrig, and Davide Scaramuzza. Dsec: A stereo event camera dataset for driving scenarios. *IEEE Robot. Autom. Lett.*, 2021. 2, 6, 8



- [14] Mathias Gehrig, Mario Millhäusler, Daniel Gehrig, and Davide Scaramuzza. E-raft: Dense optical flow from event cameras. In *Int. Conf. 3D Vision*, 2021. 1, 2, 6, 7
- [15] Mathias Gehrig, Manasi Muglikar, and Davide Scaramuzza. Dense continuous-time optical flow from event cameras. *IEEE TPAMI*, 2024. 1, 2, 6, 7, 8
- [16] Albert Gu and Tri Dao. Mamba: Linear-time sequence modeling with selective state spaces. *arXiv preprint arXiv:2312.00752*, 2023. 3, 4, 5
- [17] Albert Gu, Tri Dao, Stefano Ermon, Atri Rudra, and Christopher Ré. Hippo: Recurrent memory with optimal polynomial projections. *Neural Inf. Process. Syst.*, 33, 2020. 3, 5, 7, 8
- [18] Albert Gu, Isys Johnson, Karan Goel, Khaled Saab, Tri Dao, Atri Rudra, and Christopher Ré. Combining recurrent, convolutional, and continuous-time models with linear state-space layers. *Neural Inf. Process. Syst.*, 34, 2021. 3, 4
- [19] Albert Gu, Karan Goel, and Christopher Ré. Efficiently modeling long sequences with structured state spaces. In *ICLR*, 2022. 3
- [20] Albert Gu, Ankit Gupta, Karan Goel, and Christopher Ré. On the parameterization and initialization of diagonal state space models. *Neural Inf. Process. Syst.*, 35, 2022. 3
- [21] Ankit Gupta, Albert Gu, and Jonathan Berant. Diagonal state spaces are as effective as structured state spaces. *Neural Inf. Process. Syst.*, 35:22982–22994, 2022. 3
- [22] Jesse Hagenaars, Federico Paredes-Vallés, and Guido De Croon. Self-supervised learning of event-based optical flow with spiking neural networks. *Neural Inf. Process. Syst.*, 34:7167–7179, 2021. 7
- [23] Ramin Hasani, Mathias Lechner, Tsun-Hsuan Wang, Makram Chahine, Alexander Amini, and Daniela Rus. Liquid structural state-space models. In *ICLR*, 2022. 3
- [24] Berthold K.P. Horn and Brian G. Schunck. Determining optical flow. *Artificial Intelligence*, 1981. 1
- [25] Azin Jahedi, Lukas Mehl, Marc Rivinius, and Andrés Bruhn. Multi-scale RAFT: Combining hierarchical concepts for learning-based optical flow estimation. In *ICIP*, pages 1236–1240, 2022. 1
- [26] Kiran Kale, Sushant Pawar, and Pravin Dhulekar. Moving object tracking using optical flow and motion vector estimation. In *Int. Conf. Reliability, Infocom Technol. Optim.*, pages 1–6, 2015. 1
- [27] Taewoo Kim, Yujeong Chae, Hyun-Kurl Jang, and Kuk-Jin Yoon. Event-based video frame interpolation with cross-modal asymmetric bidirectional motion fields. In *CVPR*, pages 18032–18042, 2023. 1
- [28] Wei-Sheng Lai, Jia-Bin Huang, and Ming-Hsuan Yang. Semi-supervised learning for optical flow with generative adversarial networks. *Neural Inf. Process. Syst.*, 30, 2017. 2
- [29] 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 *ECCV*, pages 366–382, 2020. 7
- [30] Yijin Li, Zhaoyang Huang, Shuo Chen, Xiaoyu Shi, Hongsheng Li, Hujun Bao, Zhaopeng Cui, and Guofeng Zhang. Blinkflow: A dataset to push the limits of event-based optical flow estimation. In *IEEE/RSJ Int. Conf. Intell. Robot. Syst.*, 2023. 6
- [31] Zhuoyan Li, Jiawei Shen, and Ruitao Liu. A lightweight network to learn optical flow from event data. In *ICPR*. IEEE, 2021. 7
- [32] Haotian Liu, Guang Chen, Sanqing Qu, Yanping Zhang, Zhi-jun Li, Alois Knoll, and Changjun Jiang. Tma: Temporal motion aggregation for event-based optical flow. In *ICCV*, pages 9685–9694, 2023. 2, 6, 7, 8
- [33] Yue Liu, Yunjie Tian, Yuzhong Zhao, Hongtian Yu, Lingxi Xie, Yaowei Wang, Qixiang Ye, and Yunfan Liu. Vmamba: Visual state space model. *arXiv preprint arXiv:2401.10166*, 2024. 2, 3, 5
- [34] Xinglong Luo, Kunming Luo, Ao Luo, Zhengning Wang, Ping Tan, and Shuaicheng Liu. Learning optical flow from event camera with rendered dataset. In *ICCV*, pages 9847–9857, 2023. 1, 2, 6
- [35] Zhixiang Min, Yiding Yang, and Enrique Dunn. Voldor: Visual odometry from log-logistic dense optical flow residuals. In *CVPR*, 2020. 1
- [36] Elias Mueggler, Christian Forster, Nathan Baumli, Guillermo Gallego, and Davide Scaramuzza. Lifetime estimation of events from dynamic vision sensors. In *IEEE Int. Conf. Robot. Autom.*, pages 4874–4881, 2015. 2
- [37] Eric Nguyen, Karan Goel, Albert Gu, Gordon Downs, Preey Shah, Tri Dao, Stephen Baccus, and Christopher Ré. S4nd: Modeling images and videos as multidimensional signals with state spaces. *Neural Inf. Process. Syst.*, 35:2846–2861, 2022. 3
- [38] Garrick Orchard, Ryad Benosman, Ralph Etienne-Cummings, and Nitish V Thakor. A spiking neural network architecture for visual motion estimation. In *IEEE Biomedical Circuits and Systems Conference*, pages 298–301. IEEE, 2013. 2
- [39] 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 *CVPR*, pages 3446–3455, 2021. 7
- [40] AJ Piergiovanni and Michael S. Ryoo. Representation flow for action recognition. In *CVPR*, 2019. 1
- [41] Tong Qin, Peiliang Li, and Shaojie Shen. Vins-mono: A robust and versatile monocular visual-inertial state estimator. *IEEE Trans. Robot.*, 34(4):1004–1020, 2018. 1
- [42] Zhile Ren, Orazio Gallo, Deqing Sun, Ming-Hsuan Yang, Erik B. Sudderth, and Jan Kautz. A fusion approach for multi-frame optical flow estimation. In *Winter Conf. Appl. Comput. Vis.*, pages 2077–2086, 2019. 3
- [43] Xiaoyu Shi, Zhaoyang Huang, Weikang Bian, Dasong Li, Manyuan Zhang, Ka Chun Cheung, Simon See, Hongwei Qin, Jifeng Dai, and Hongsheng Li. Videoflow: Exploiting temporal cues for multi-frame optical flow estimation. In *ICCV*, pages 12469–12480, 2023. 1, 2, 3, 4, 5, 8
- [44] Amos Sironi, Manuele Brambilla, Nicolas Bourdis, Xavier Lagorce, and Ryad Benosman. HATS: Histograms of averaged time surfaces for robust event-based object classification. In *CVPR*, pages 1731–1740, 2018. 3

- [45] Jimmy T.H. Smith, Andrew Warrington, and Scott Linderman. Simplified state space layers for sequence modeling. In *ICLR*, 2023. 3
- [46] Taylan Soydan, Nikola Zubić, Nico Messikommer, Sidhartha Mishra, and Davide Scaramuzza. S7: Selective and simplified state space layers for sequence modeling. *arXiv*, 2024. 3
- [47] Timo Stoffregen, Cedric Scheerlinck, Davide Scaramuzza, Tom Drummond, Nick Barnes, Lindsay Kleeman, and Robert Mahony. Reducing the sim-to-real gap for event cameras. In *ECCV*, pages 534–549, 2020. 7
- [48] Shuyang Sun, Zhanghui Kuang, Lu Sheng, Wanli Ouyang, and Wei Zhang. Optical flow guided feature: A fast and robust motion representation for video action recognition. In *CVPR*, 2018. 1
- [49] SHANGKUN SUN, Yuanqi Chen, Yu Zhu, Guodong Guo, and Ge Li. SKFlow: Learning optical flow with super kernels. In *Neural Inf. Process. Syst.*, 2022. 6
- [50] Zachary Teed and Jia Deng. Raft: Recurrent all-pairs field transforms for optical flow. In *ECCV*, 2020. 1, 2, 4, 7, 8
- [51] Hugo Touvron, Matthieu Cord, Matthijs Douze, Francisco Massa, Alexandre Sablayrolles, and Herve Jegou. Training data-efficient image transformers and distillation through attention. In *Int. Conf. Mach. Learn.*, pages 10347–10357, 2021. 2
- [52] Stepan Tulyakov, Daniel Gehrig, Stamatios Georgoulis, Julius Erbach, Mathias Gehrig, Yuanyou Li, and Davide Scaramuzza. Time lens: Event-based video frame interpolation. In *CVPR*, pages 16155–16164, 2021. 1
- [53] 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 *CVPR*, pages 17755–17764, 2022. 1
- [54] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. Attention is all you need. In *Neural Inf. Process. Syst.*, 2017. 8
- [55] Jonáš Šerých and Jiří Matas. Planar object tracking via weighted optical flow. In *Winter Conf. Appl. Comput. Vis.*, pages 1593–1602, 2023. 1
- [56] Zhexiong Wan, Yuchao Dai, and Yuxin Mao. Learning dense and continuous optical flow from an event camera. *IEEE TIP*, 2022. 6, 7
- [57] Yilun Wu, Federico Paredes-Vallés, and Guido C. H. E. de Croon. Lightweight event-based optical flow estimation via iterative deblurring. In *IEEE Int. Conf. Robot. Autom.*, pages 14708–14715, 2024. 2, 6
- [58] Yan Yang, Liyuan Pan, and Liu Liu. Event camera data dense pre-training. In *ECCV*, 2024. 2, 6
- [59] 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 *IEEE/RSJ Int. Conf. Intell. Robot. Syst.*, pages 5831–5838, 2020. 2
- [60] Annan Yu, Arnur Nigmatov, Dmitriy Morozov, Michael W. Mahoney, and N. Benjamin Erichson. Robustifying state-space models for long sequences via approximate diagonalization. In *ICLR*, 2024. 5, 6, 7, 8
- [61] Jason J Yu, Adam W Harley, and Konstantinos G Derpanis. Back to basics: Unsupervised learning of optical flow via brightness constancy and motion smoothness. In *ECCV*, 2016. 2
- [62] Alex Zhu, Liangzhe Yuan, Kenneth Chaney, and Kostas Daniilidis. Ev-flownet: Self-supervised optical flow estimation for event-based cameras. In *Robot. Sci. Syst.*, 2018. 1, 2, 7
- [63] Alex Zihao Zhu, Dinesh Thakur, Tolga Özaslan, Bernd Pfrommer, Vijay Kumar, and Kostas Daniilidis. The multi-vehicle stereo event camera dataset: An event camera dataset for 3d perception. *IEEE Robot. Autom. Lett.*, 3(3):2032–2039, 2018. 2, 6, 8
- [64] Alex Zihao Zhu, Liangzhe Yuan, Kenneth Chaney, and Kostas Daniilidis. Unsupervised event-based learning of optical flow, depth, and egomotion. In *CVPR*, 2019. 7
- [65] Lianghui Zhu, Bencheng Liao, Qian Zhang, Xinlong Wang, Wenyu Liu, and Xinggang Wang. Vision mamba: Efficient visual representation learning with bidirectional state space model. *arXiv preprint arXiv:2401.09417*, 2024. 3, 4, 5
- [66] Nikola Zubić and Pietro Liò. An effective loss function for generating 3d models from single 2d image without rendering. In *Artif. Intell. Appl. Innov.*, pages 309–322. Springer International Publishing, 2021. 5
- [67] Nikola Zubić, Daniel Gehrig, Mathias Gehrig, and Davide Scaramuzza. From chaos comes order: Ordering event representations for object recognition and detection. In *ICCV*, pages 12846–12856, 2023. 2, 3, 5, 6
- [68] Nikola Zubić, Mathias Gehrig, and Davide Scaramuzza. State space models for event cameras. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2024. 3, 5
- [69] Nikola Zubić and Davide Scaramuzza. Gg-ssms: Graph-generating state space models. *CVPR*, 2025. 3
- [70] Nikola Zubić, Federico Soldá, Aurelio Sulser, and Davide Scaramuzza. Limits of deep learning: Sequence modeling through the lens of complexity theory. *ICLR*, 2025. 3