# BlinkFlow: A Dataset to Push the Limits of Event-based Optical Flow Estimation

Yijin Li[1*], Zhaoyang Huang[2*], Shuo Chen[1], Xiaoyu Shi[2], Hongsheng Li[2], Hujun Bao[1],
Zhaopeng Cui[1], Guofeng Zhang[1†]
[1]State Key Lab of CAD&CG, Zhejiang University
[2]Multimedia Laboratory, The Chinese University of Hong Kong

*Abstract*— Event cameras provide high temporal precision, low data rates, and high dynamic range visual perception, which are well-suited for optical flow estimation. While data-driven optical flow estimation has obtained great success in RGB cameras, its generalization performance is seriously hindered in event cameras mainly due to the limited and biased training data. In this paper, we present a novel simulator, BlinkSim, for the fast generation of large-scale data for event-based optical flow. BlinkSim incorporates a configurable rendering engine alongside an event simulation suite. By leveraging the wealth of current 3D assets, the rendering engine enables us to automatically build up thousands of scenes with different objects, textures, and motion patterns and render very high-frequency images for realistic event data simulation. Based on BlinkSim, we construct a large training dataset and evaluation benchmark BlinkFlow that contains sufficient, diversiform, and challenging event data with optical flow ground truth. Experiments show that BlinkFlow improves the generalization performance of state-of-the-art methods by more than 40% on average and up to 90%. Moreover, we further propose an Event-based optical Flow transFormer (E-FlowFormer) architecture. Powered by our BlinkFlow, E-FlowFormer outperforms the SOTA methods by up to 91% on the MVSEC dataset and 14% on the DSEC dataset and presents the best generalization performance. The source code and data are available at `https://zju3dv.github.io/blinkflow/`.

## I. INTRODUCTION

Event cameras detect changes in intensity at each pixel of the image as a stream of asynchronous events, which naturally depicts pixel motions. Besides, their high temporal precision (on the order of microseconds), low data rates, and high dynamic ranges ($> 120$ dB) far exceed traditional RGB cameras. Such significant characteristics make event cameras popular in robotics [37], [13] and particularly suitable for tasks such as optical flow [18], [35] estimation, which requires accurate and efficient tracking of pixel motion between frames.

The success of data-driven optical flow estimation for RGB cameras motivates researchers to develop learning-based methods for event-based optical flow estimation. However, previous works [43], [22], [18] highly overfit their dataset (Fig. 1) and their performance is far from satisfactory in practical scenarios. One major reason is the lack of large-scale diversiform training data along with a benchmark to evaluate their generalization performance. Due to the novelty
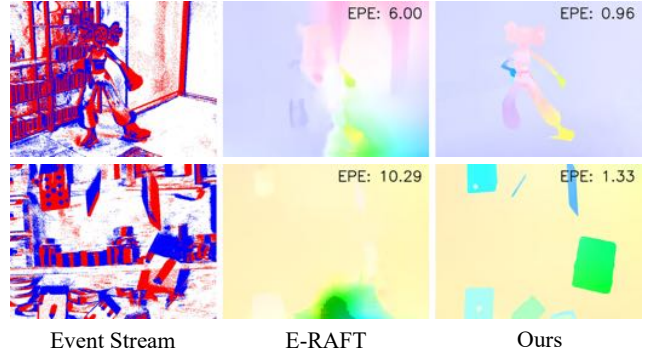


Fig. 1. Comparison of generalization. E-RAFT [18] overfits the DSEC dataset [17] and generalizes poorly to unseen environments. In contrast, our E-FlowFormer powered by the proposed BlinkFlow dataset presents good generalization performance and recovers complex and flexible optical flows.

and complexity of the event camera, there are currently only four event camera datasets that contain the ground truth data of optical flow. These datasets cover limited scenarios, simple motion patterns and only contain sparse optical flow ground truth, as shown in Table I. To alleviate the problem, one potential strategy [15] is converting the existing RGB frame-based dataset into event data based on video interpolation, but it inevitably produces artifacts and introduces data bias. Such issues may be acceptable in object recognition [24] and semantic segmentation [15], which focus on high-level semantics, but can severely disturb low-level tasks such as optical flow especially when encountering violent motions and occlusion, as shown in Fig. 2.

We observe that the favorable outcome of frame-based optical flow learning may largely be attributed to the synthetic datasets [28], [11]. Motivated by them, we propose BlinkSim, an event simulator that renders event data sequences and obtains their optical flow ground truth. BlinkSim consists of two modules including a configurable rendering engine and a simulation suite with integrated multiple event emulators. Our effort is similar in spirit to the FlyingChair [11] and FlyingThings [28] for RGB image-based optical flow estimation, but it is customized for event cameras. Specifically, BlinkSim allows rendering very high frame rate images and is coupled with the event data simulator suite, which ensures realistic event data generation. Based on BlinkSim, we generate a large training dataset and evaluation benchmark BlinkFlow. As shown in Fig. 3, BlinkFlow covers diverse

---

Fig. 2. Video interpolation artifacts (pointed by red arrows).

scene categories and complex motion patterns with sufficient realism. To boost the research in event camera communities, we also open-source the BlinkSim besides the release of BlinkFlow data. BlinkSim can be extended to alleviate the data problem in other tasks related to event cameras, such as feature tracking [16], [5] and depth estimation [29].

Aside from data issues, the unfamiliar way in which events encode visual information also poses challenges for optical flow estimation. Recent methods [18], [10] have demonstrated that the methods designed for images, i.e., usage of correlation volume, can be adopted and work quite well on event data. However, as event cameras provide neither absolute brightness information nor spatially continuous data and the event data suffers from severe noise compared to image data, directly applying correlation to event data cannot produce a discriminative matching cost. To address this problem, we propose an Event optical Flow transFormer (E-FlowFormer) architecture, which enhances the event feature extraction by fully encoding information from input events. Empowered by the diversiform and high-quality training data of our BlinkFlow, E-FlowFormer learns effective and generalizable event features to build a distinct correlation volume for the following flow refinement.

Our contributions can be summarized in three folds: 1) We develop a novel simulator, BlinkSim, to synthesize events along with related ground truth. 2) We build a dataset BlinkFlow based on BlinkSim, for training and evaluating event-based optical flow estimation methods. BlinkFlow exceeds previous datasets in quantity, quality, and diversity. 3) We propose an E-FlowFormer to fully unleash the power of the large-scale training data provided by BlinkFlow. E-FlowFormer ranks 1st on mainstream benchmarks DSEC and MVSEC and the proposed BlinkFlow benchmark.

## II. RELATED WORK

### A. Optical Flow Dataset and Simulator for Event Camera

As far as we know, there are currently only four public event camera datasets that include optical flow ground truth data. which can be categorized based on how the ground truth optical flow is obtained. In DVSFLOW16 [32] and DVS-MOTION20 [1], the ground truth optical flow was deduced from the camera's rotational movement, as measured by the inertial measurement unit. DSEC [17] and MVSEC [43] computed the ground truth through LiDAR SLAM. Both ways limit the ground truth to the static elements of the scene and the latter one only offers sparse data up to a specific distance and height. Besides, these datasets have a very limited motion pattern. For example, DVSFLOW16 [32] and DVSMOTION20 [1] contain only rotational camera motion, while in DSEC [17], the camera primarily moves straight ahead, interspersed with turns to the left or right. A comprehensive overview and comparison are listed in Table I. As for event simulators, most of the current works [19], [31], [25] focus on how to generate realistic events from high-frequency image sequences. However, obtaining these sequences is not easy either. While some works [15] propose to convert existing frame-based datasets into event data by firstly interpolating the frames densely, they are limited by the interpolation artifacts and are not suitable for low-level applications like optical flow. In this paper, we propose BlinkSim to alleviate the problem of generating large-scale diverse scenarios with corresponding event data and optical flow labels. Both the BlinkSim and the generated data BlinkFlow are released.

### B. Event-based Optical Flow Estimation

Extensive research has been carried out on the use of event cameras for estimating optical flow, given their identified advantages. Previous studies have proposed adaptations of frame-based techniques (block matching [27] and Lucas-Kanade [4]), time surface matching [1], variational optimization [2], spatio-temporal plane-fitting [3], contrast maximization [35] and so on. These methods usually assume constant brightness. Alternatively, another stream of research is exploring learning-based methods. In these methods, event streams are typically converted into a grid-like representation in a pre-processing step. Then frame-based designs and frameworks such as FlowNet [11], cost volume [36] and recurrent refinement [38], [20], [33], [34] can be applied. In total, the learning-based methods outperform the former hand-designed methods in current benchmarks [43], [17]. In this paper, we propose E-FlowFormer to further improve the capacity of the learning-based method by enhancing event features through transformers.

## III. BLINKSIM

BlinkSim is a simulator designed to automate the generation of large amounts of realistic event data and high-quality optical flow ground truth. It is mainly composed of two separate components: a configurable and photo-realistic rendering engine in a high frame rate built with the Blender[1] and a simulation suite with integrated multiple event emulators. The two modules are decoupled and can run independently from each other. Users can simulate event data from rendered images of thousands of 3D scenes that are automatedly built from scratch using our rendering engine, or generate events from existing high-speed video data by selecting an appropriate event data simulator from our suite. In this section, we first introduce the event generation model, then introduce the rendering engine and the event simulation module, respectively. Finally, we describe how we render the diversiform event optical flow dataset, BlinkFlow, based on the proposed simulator.

---

[1]https://www.blender.org

TABLE I

COMPARISON OF AVAILABLE DATASETS.

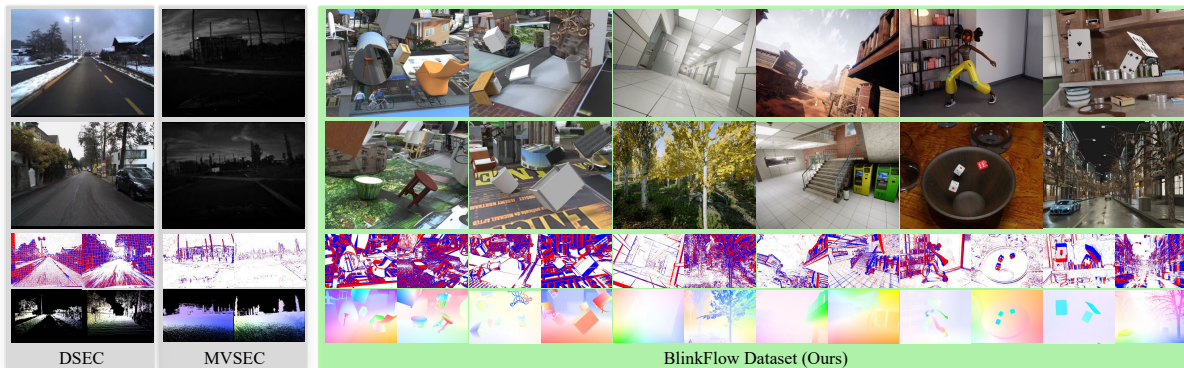| Dataset | Motion Pattern | Dynamic Object | Occlusion | Training Frames | Training Scenes | Test Scenes | Resolution |
|---|---|---|---|---|---|---|---|
| DVSFLOW16 [32] | Rotation | ✗ | ✗ | - | - | 5 | $180 \times 240$ |
| DVSMOTION20 [1] | Rotation | ✗ | ✗ | - | - | 4 | $260 \times 346$ |
| MVSEC [42] | Drone | ✗ | ✗ | 3k | 1 | 4 | $260 \times 346$ |
| DSEC [17] | Car | ✗ | ✗ | 8k | 18 | 7 | $640 \times 480$ |
| BlinkFlow (Ours) | Random | ✔ | ✔ | 33k | 3362 | 225 | $640 \times 480$ |



Fig. 3. Example scenes from DSEC, MVSEC and the proposed BlinkFlow dataset. **3rd row: Event data**, **4th row: Optical flow images**. Our BlinkFlow dataset contains complex object/camera motions and various scenarios which significantly outperform DSEC and MVSEC in quantity, quality, and diversity. The 1-2, 3-4, and 5-6 columns of the BlinkFlow Dataset correspond to the sequences of FlyingObjects, E-Tartan and E-Blender, respectively. Best viewed on a color screen in high resolution.

## A. Event Generation Model

Event cameras operate by responding to changes in the logarithmic brightness signal in an asynchronous and independent manner for each event pixel [14]. An event is triggered when the brightness change (either increment or decrease) since the last event at that pixel reaches a threshold $\pm C$ (with $C > 0$):

$$p_k(L(u_k, t_k) - L(u_k, t_k - \Delta t)) \geq C, \quad (1)$$

where $\Delta t$ is the time since the last event triggered at pixel location $u_k$ and at time $t_k$. $p_k \in \{-1, 1\}$ is the polarity of the brightness change. To simulate the event generation process, we need access to the continuous representation of the visual signal for each pixel, which is practically unattainable. In practice, we approximate the continuous signal changes by rendering images with a high frame rate and then synthesize events according to frame-wise pixel differences.

## B. Rendering Engine

The rendering engine for BlinkSim is developed with the open-source 3D creation suite Blender, based on which we can render diverse objects with complex motions under flexible scene configurations.

**3D Environment Setup** Data-driven models benefit greatly from having access to abundant and diverse training data. To this end, we have developed a scalable pool of rendering assets based on ShapeNet [6], which contains thousands of real-world 3D objects. The pool also contains various basic shapes like cuboids and cylinders which are randomly textured with real-world images from ADE20K dataset [41].

By default, each scene is bounded by a large textured sky box and objects are randomly selected from the pool, randomly scaled, rotated, textured, and placed on the ground plane. The related materials can be further extended with other available resources. For example, we can collect objects in Google Scanned Objects [12] and images in COCO [26]. In addition, BlinkSim provides a user-friendly interface for importing customized scenes downloaded from websites such as BlenderKit[2] and Blender Store[3].

**Trajectory Sampling** To generate diverse motion patterns, we need to sample various trajectories, including object-moving and camera-moving trajectories. The two trajectories are similar: for a timestamp $t$, determining a 6-degree-of-free (DoF) pose $\mathbf{P}(t)$ including 3-DoF rotation and 3-DoF translation. The difference is that we put objects in the scene according to the object poses while we send camera poses to the virtual camera in Blender for rendering data. We randomly sample N points in the free space with random rotation and fit a spline. The spline curve determines that for any timestamp $t$ we have a unique 6-DoF pose. We also detect collisions between the objects and the camera and between multiple objects by checking the intersection of the trajectories [21]. Once a collision is detected, we cut the trajectory of the collided objects to avoid the model breakthrough. The next step is to sample the timestamps for rendering. A simple strategy is sampling $t$ evenly but may fail to faithfully simulate events when the brightness signal varies more quickly than the chosen rendering frame rate.

[2]https://www.blenderkit.com/
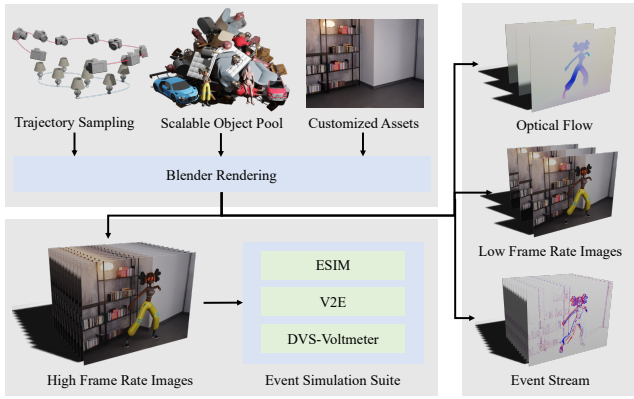[3]https://store.blender.org/

Fig. 4. Framework of BlinkSim. BlickSim consists of a configurable rendering engine and an event data simulation suite. It allows the fast generation of large-scale data with realistic event data and related optical flow ground truth.

Another choice is to adaptively increase $t$ by limiting the maximum pixel displacement or brightness change between two sampled timestamps [31]. However, it requires generating high frame rate optical flows for pixel displacement estimation, which hinders rendering efficiency, and it is unnecessary when the motion is slow-varying. BlinkSim supports both image sequence rendering strategies so that users can choose a suitable strategy according to their time-consuming requirements and motion settings.

**Data Generation** We render the data with two passes. As shown in Fig. 4, we first render images of the densely-sampled viewpoint which are used for event data generation. In the second pass, we sample the viewpoint sparsely and render the simulated RGB camera's view. We use Blender-Proc [9] which modifies the pipeline of Blender's internal render engine and allows us to obtain the corresponding optical flow ground truth.

### C. Event Data Simulator

Eq. 1 describes the ideal process for event generation, but in practice, the triggering threshold is not constant. Furthermore, there are complicated noise effects such as refractory periods and temporal noises at low illumination. BlinkSim integrates three state-of-the-art (SOTA) event data simulators inside and provides flexible choices for users. The comparison of these data simulators is listed in Table II. ESIM [31] is an early-stage simulator that only considers spatially varying thresholds but it runs the fastest on a CPU. DVS-Voltmeter [25] generates the most realistic event data by considering spatiotemporally varying thresholds, modeling the event triggering timestamp as Brownian motion instead of linearly distributed. Besides, it considers the influence of temperature and parasitic photocurrent. But at the same time, the DVS-Voltmeter runs the slowest and does not have GPU acceleration. V2E works best when simulating the scenes in low illuminations and it supports accelerating rendering with GPUs, whose speed is on par with ESIM. Users can choose suitable emulators according to the scenes and the requirement of time consumption.

### D. BlinkFlow Dataset

Based on BlinkSim, we create BlinkFlow, which consists of a large-scale, diversiform training dataset and a challenging evaluation benchmark. We split the object pool into two disjoint parts and generated 3362 and 200 scenes for training and testing, respectively. Besides, we selected 20 sequences from the TartanAir [40] dataset and converted them to event data using our event simulator suite. The sequences were carefully chosen to ensure that no apparent artifacts from frame interpolation. Finally, we downloaded some well-designed assets from Blender Store and Mixamo[4] and created 5 customized scenes including "dancing women", "falling dice" and so on, which are partly shown in Fig. 3. The resulting 200, 20, and 5 test sequences (called FlyingObjects, E-Tartan, and E-Blender, respectively) contain a total of over 4000 frames. We choose V2E [19] for the event simulation when generating most of the training data and DVS-Voltmeter [25] for other training data and all the testing data. When simulating events for each sequence, the corresponding simulation parameters are randomly selected within a range given in advance, so that the trained network do not overfit to specific parameters. For example, the contrast threshold $C$ is randomly selected in the range of [0.18, 0.25]. As shown in Table I, BlinkFlow contains complex camera motions and dynamic objects based on our trajectory sampling, produces events with ground truth flows for occluded objects, provides individual training/test splits, covers more training scenes (3362 v.s. 18) and test scenes (225 v.s. 7), and generates more training frames (33k v.s. 8k). Such a rich dataset can push the limits of event-based optical flow estimation and encourage researchers to explore better neural networks for event cameras.

### IV. E-FLOWFORMER

E-FlowFormer is built upon E-RAFT [18], a state-of-the-art neural network architecture for learning event-based optical flow estimation. It encodes event features with a shallow Siamese CNN to construct a 4D correlation volume, which measures event similarity. However, events are noisier than RGB images and the corresponding locations derived from ground truth flows do not always present the same event patterns. Inspired by recent progress in deep learning [39], [20], [23], [30], we propose to enhance the event feature encoding modules with transformers. We call the enhanced neural network as E-FlowFormer. An overview of the proposed method E-FlowFormer is presented in Fig. 5. In this section, we first review E-RAFT and then describe our transformer-based event feature enhancement module.

### A. E-RAFT Brief Review

E-RAFT consists of three stages: feature extraction, correlation volume computation, and motion refinement. Given two consecutive packets of events, E-RAFT converts them to tensor-like event representations and then encodes local features of events $F_A, F_B \in \mathbb{R}^{H \times W \times D}$ with a siamese CNN.

---

[4] https://www.mixamo.com/

| | Threshold | Timestamp Modeling | Temperature Noise | Low Illumination | Refractory Periods | Speed (10k ev./s) CPU / GPU |
|---|---|---|---|---|---|---|
| ESIM [31] | Spatial | Linear | ✗ | ✗ | ✗ | 170 / - |
| V2E [19] | Spatial | Linear | ✗ | ✔ | ✔ | 26 / 141 |
| DVS-Voltmeter[25] | Spaitio-temperoal | Brownian | ✔ | ✗ | ✗ | 29 / - |

With the event feature tensors $F_A$ and $F_B$, a correlation volume $C \in \mathbb{R}^{H \times W \times H \times W}$ which encodes the feature similarity for each pixel in $F_A$ with respect to all pixels in $F_B$ is computed through:

$$C = \frac{F_A F_B^T}{\sqrt{D}} \in \mathbb{R}^{H \times W \times H \times W}, \qquad (2)$$

where $\frac{1}{\sqrt{D}}$ is a normalization factor that can avoid large values after the dot-product operation. Then, E-RAFT iteratively refines the flow estimation by taking the current motion estimation, the correlation volume, and the context features with a ConvGRU module. The flow is usually initialized to zero in previous implementation [18]. E-RAFT supervises all flow predictions using a sequence loss:

$$L = \sum_{i=1}^{N} \gamma^{N-1} \|V_i - V_{gt}\|_1, \qquad (3)$$

where $N$ is the number of flow predictions. $\gamma$ is the weight that exponentially increases to give higher weights for later predictions. $V_i$ is the flow prediction at the $i$-th iteration and $V_{gt}$ is the ground truth flow.

### B. Event Feature Enhancement Through Transformers

The key to E-RAFT lies in obtaining high-quality discriminative event features. As previously mentioned, the features $F_A$ and $F_B$ are encoded independently from a weight-sharing shallow CNN. We enhance the event feature encoding with transformers to incorporate their interdependence better. Cooperated with positional encoding, transformers can effectively capture both the relevant spatial relationships and feature similarity between the two sets of events and thus generate refined and discriminative event features. A remaining issue is that the standard Transformer architecture introduces quadratic computational complexity. To avoid this issue, we adopt the spatially separable self-attention proposed in Twins [7] instead of the vanilla version. The enhanced event feature encoder improves the distinctiveness of event features and thus brings a better correlation volume, based on which, we can identify the correspondence through a differentiable matching layer [30]. The resulting correspondence will be added and supervised in the sequence loss, effectively suppressing the correlation strength far from the ground truth locations. Besides, we use the resulting correspondence to better initialize the following refinement module.

## V. EXPERIMENT

In this section, we first show that the models trained on BlinkFlow gain better generalization performance than those
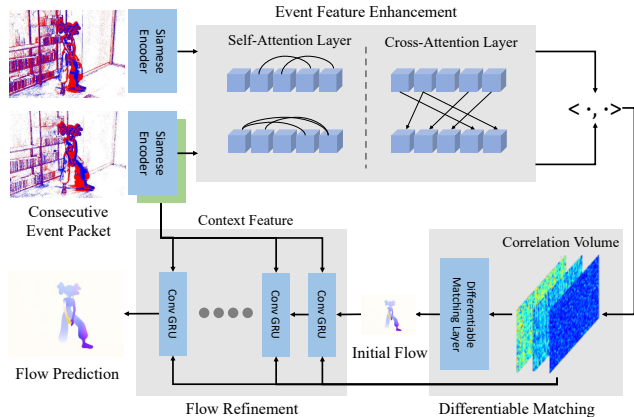


Fig. 5. Network architecture of E-FlowFormer. The proposed feature enhancement module brings a better correlation volume through the transformer design. Besides, it provides a better initialization of optical flow for the following refinement module.

trained on DSEC and MVSEC, which reveals the superiority of training data in BlinkFlow. Then, we demonstrate that the evaluation benchmark in BlinkFlow is more challenging and comprehensive. Finally, we evaluate the proposed E-FlowFormer which ranks 1st on the MVSEC, DSEC and BlinkFlow benchmarks. This indicates that the transformer-based event feature enhancement module effectively learns the event features.

**Dataset** There are two commonly used datasets for event-based optical flow training and evaluation: MVSEC [42] and DSEC [17]. MVSEC contains two scenarios, i.e., sequences recorded in the indoor room and outdoor sequences recorded while driving on public roads. The ground truth of optical flow is approximated by multiplying the motion field by the time interval (denoted as "dt") where the motion field is computed according to the camera movement and scene depth. DSEC [17] contains data in a higher resolution and improves the quality of the ground truth optical flow by removing the occlusions and moving objects. Both of MVSEC and DSEC collect data in rigid and static scenes and compute optical flows via post-processing, resulting in highly biased data and inaccurate ground truth. In contrast, our BlinkFlow dataset covers various scenes and objects and directly obtains the ground truth flows from the renderer. We denote DSEC, MVSEC, and BlinkFlow datasets by "D", "M", and "B".

**Experiment Setup.** We measure the accuracy of the optical flow using several metrics, namely the Average Endpoint Error (AEE), outlier rate, Angular Error (AE), and N-pixel errors (N-PE). The AEE measures the Euclidean distance between the predicted flow and the ground truth. The outlier

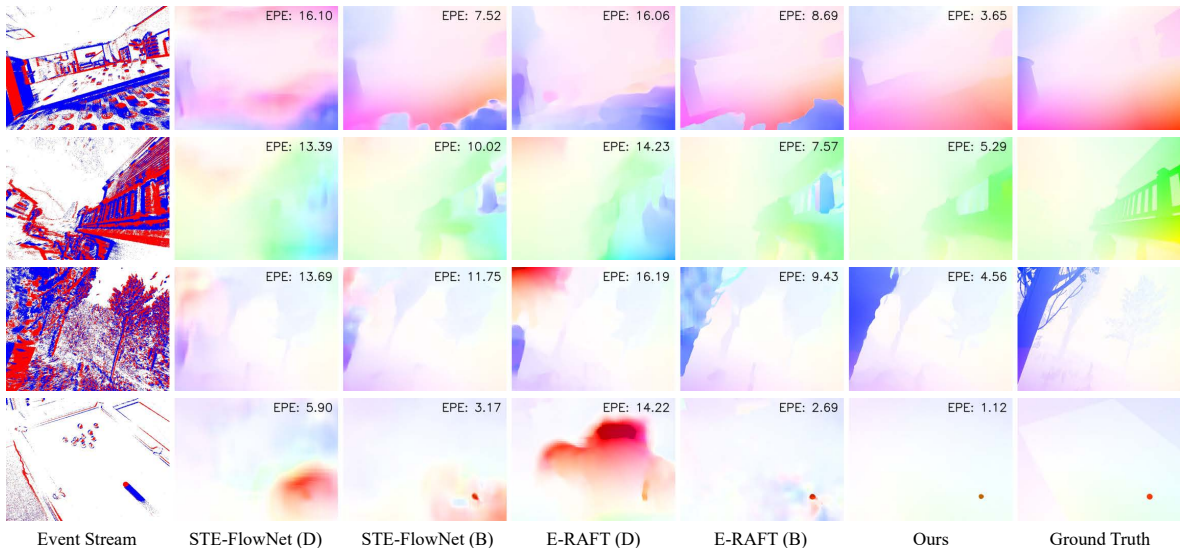| Event Stream | STE-FlowNet (D) | STE-FlowNet (B) | E-RAFT (D) | E-RAFT (B) | Ours | Ground Truth |

Fig. 6. Qualitative comparison on BlinkFlow. BlinkFlow contains challenging data that makes previous methods trained on DSEC almost fail. Models trained on BlinkFlow consistently outperform models trained on DSEC and E-FlowFormer achieves the lowest AEE among all the methods, producing accurate and smooth optical flow prediction.

| Methods | Training Data | DSEC (train) | | Training Data | MVSEC (indoor) | | MVSEC (outdoor) | |
|---|---|---|---|---|---|---|---|---|
| | | AEE | Outlier | | AEE | Outlier | AEE | Outlier |
| STE-FlowNet [10] | M | 9.91 | 77.59 | D | 2.17 | 19.77 | 1.34 | 8.52 |
| | B | 1.45 | 9.19 | B | 1.72 | 12.21 | 1.37 | 8.59 |
| | Impro. | +85% | +88% | Impro. | +22% | +38% | -2% | -1% |
| E-RAFT [18] | M | 11.24 | 84.89 | D | 2.29 | 18.12 | 1.47 | 8.98 |
| | B | 1.34 | 8.12 | B | 1.85 | 12.03 | 1.35 | 8.56 |
| | Impro. | +88% | +90% | Impro. | +19% | +34% | +8% | +5% |
| E-FlowFormer (Ours) | M | 7.92 | 65.85 | D | 2.44 | 22.35 | 1.45 | 9.21 |
| | B | 1.25 | 6.73 | B | 1.62 | 11.18 | 1.29 | 8.08 |
| | Impro. | +84% | +90% | Impro. | +34% | +50% | +11% | +12% |

rate represents the percentage of points where the endpoint error exceeds 3 pixels and 5% of the magnitude. The AE calculates the angle between normalized optical flow predictions and the normalized ground truth. N-PE denotes the percentage of flow errors greater than N pixels in magnitude.

**Training Data in BlinkFlow Improves Generalization.** We set up two experiments with three network architectures, i.e., STE-FlowNet [10], E-RAFT [18], and our E-FlowFormer, to present that models trained on BlinkFlow (B) gain the best generalization performance. We train models on MVSEC (M) and BlinkFlow (B) and evaluate the models on DSEC. As shown in Table III, the models trained on BlinkFlow achieve at least 84% error reduction. Then, we further train the models on DSEC and BlinkFlow and evaluate them on MVSEC (Table III), which also reveals significant performance improvement for E-RAFT and our E-FlowFormer, up to 34% AEE reduction on MVSEC (indoor). BlinkFlow does not outperform DSEC much when evaluated on the outdoor sequence of MVSEC because DSEC and MVSEC (outdoor) share a similar scene pattern, i.e., both are outdoor driving scenarios. Totally speaking, both experiments present the extraordinary superiority of our BlinkFlow. Moreover, our

proposed E-FlowFormer makes the best use of BlinkFlow and outperforms the other two methods all-sided.

**Test Data in BlinkFlow Is More Challenging.** All of DSEC, MVSEC, and our proposed BlinkFlow provide both training data and test data. We can train and test models with data belonging to the same dataset to present the easiness of datasets. We select E-RAFT and STE-FlowNet and use outlier metrics in this experiment. As shown in Table V, E-RAFT and STE-FlowNet obtain overall less than 4 percent in terms of outliers, which indicates that DSEC and MVSEC are so simple that they can be easily fitted. However, BlinkFlow is still challenging for them, which reveals that BlinkFlow is a more comprehensive evaluation benchmark to push the limits of event-based optical flow estimation.

**E-FlowFormer Ranks 1st on MVSEC.** In Table VI, we compare with recent SOTA methods on the MVSEC dataset. E-FlowFormer, only trained on the proposed BlinkFlow dataset achieves comparable performances with the SOTA methods leveraging the MVSEC training data. Especially, E-FlowFormer outperforms other methods by a significant margin in almost all the indoor sequences, because the training set of MVSEC is only the outdoor sequences. It is easy to overfit the training data and thus brings severe per-

TABLE IV

COMPARISON ON BLINKFLOW DATASET.

| Methods | FlyingObjects | | | E-Tartan | | | E-Blender | | |
|---|---|---|---|---|---|---|---|---|---|
| | AEE | Outlier | AE | AEE | Outlier | AE | AEE | Outlier | AE |
| Spike-FlowNet [22] | 5.39 | 32.16 | 11.67 | 4.13 | 31.31 | 15.26 | 3.82 | 34.72 | 28.03 |
| EV-FlowNet [43] | 4.02 | 16.64 | 7.14 | 3.41 | 22.17 | 11.31 | 3.24 | 26.95 | 22.46 |
| E-RAFT [18] | 3.19 | 9.27 | 5.48 | 2.83 | 16.43 | 10.13 | 2.66 | 19.90 | 18.54 |
| STE-FlowNet [10] | 3.22 | 10.44 | 5.97 | 2.79 | 16.80 | 9.83 | 2.59 | 18.87 | 17.43 |
| E-FlowFormer (Ours) | **2.94** | **9.03** | **4.92** | **2.48** | **15.98** | **8.42** | **2.38** | **17.56** | **14.78** |

TABLE V

DATASET COMPARISON BY TRAINING AND TESTING MODELS ON THE SAME DATASET IN TERMS OF OUTLIER.

| Methods | DSEC | MVSEC | BlinkFlow |
|---|---|---|---|
| E-RAFT [18] | 2.68 | 0.92 | 15.20 |
| STE-FlowNet [10] | 4.06 | 0.50 | 15.37 |

TABLE VI

COMPARISON ON MVSEC DATASET.

| dt=1 frame | Training Data | indoor | | outdoor | |
|---|---|---|---|---|---|
| | | AEE | Outlier | AEE | Outlier |
| EV-FlowNet [43] | M | 1.43 | 9.73 | 0.49 | 0.20 |
| Spike-FlowNet [22] | M | 1.08 | 3.90 | 0.47 | **0.00** |
| E-RAFT [18] | M | 0.97 | 2.76 | 0.65 | 2.19 |
| Zhu et al. [44] | M | 0.82 | 2.33 | 0.32 | **0.00** |
| STE-FlowNet [10] | M | 0.69 | 1.00 | 0.42 | **0.00** |
| E-FlowFormer (Ours) | B | 0.71 | 0.94 | 0.67 | 1.68 |
| E-FlowFormer (Ours) | B+M | **0.51** | **0.09** | **0.29** | 0.05 |
| dt=4 frame | Training Data | indoor | | outdoor | |
| | | AEE | Outlier | AEE | Outlier |
| EV-FlowNet [43] | M | 3.25 | 36.57 | 1.23 | 7.30 |
| Spike-FlowNet [22] | M | 3.08 | 33.43 | 1.09 | 5.60 |
| E-RAFT [18] | M | 2.46 | 23.15 | 1.43 | 9.17 |
| Zhu et al. [44] | M | 3.07 | 39.60 | 1.30 | 9.70 |
| STE-FlowNet [10] | M | 2.17 | 20.97 | 0.99 | 3.90 |
| E-FlowFormer (Ours) | B | 1.62 | 11.18 | 1.29 | 8.08 |
| E-FlowFormer (Ours) | B+M | **1.56** | **10.25** | **0.83** | **3.44** |

TABLE VII

COMPARISON ON DSEC DATASET.

| Methods | Training Data | 1PE | 2PE | 3PE | AEE | AE |
|---|---|---|---|---|---|---|
| MultiCM [35] | D | 76.57 | 48.48 | 30.86 | 3.47 | 13.98 |
| OF-EV-SNN [8] | D | 53.67 | 20.24 | 10.31 | 1.71 | 6.34 |
| E-RAFT [18] | D | 12.74 | 4.74 | 2.68 | 0.79 | 2.85 |
| E-FlowFormer(Ours) | B | 36.39 | 13.02 | 6.05 | 1.33 | 4.65 |
| E-FlowFormer(Ours) | B+D | **11.23** | **4.10** | **2.45** | **0.76** | **2.68** |

formance degradation on indoor sequences. On the contrary, BlinkFlow covers complex scene structures so our method can generalize well on both indoor and outdoor sequences. We also fine-tune E-FlowFormer on the data combined from BlinkFlow and MVSEC (denoted as "B+M"). Fine-tuning further improves the performance of E-FlowFormer, especially on the sequence of outdoor day1. Finally, E-FlowFormer outperforms all the published methods in almost all the sequences, up to 28% in terms of AEE metric and 91% in terms of outlier metric.

**E-FlowFormer Ranks 1st on DSEC.** We show the comparisons on the DSEC dataset in Table VII. Only trained on the proposed BlinkFlow dataset, E-FlowFormer surpasses most published methods except for E-RAFT. After fine-tuning on the data combined from BlinkFlow and DSEC (denoted as "B+D"), E-FlowFormer outperforms all other methods in terms of all metrics, and up to 14% in terms of 2PE metric. Notably, the margin between our approach and SOTA methods is smaller on DSEC than on MVSEC, because DSEC provides much more training data than MVSEC. Meanwhile, both the training and test data of DSEC are collected in outdoor driving scenarios, making it easy to overfit the testing data.

**E-FlowFormer Ranks 1st on BlinkFlow.** We train several state-of-the-art event-based optical flow estimation methods with the training data in BlinkFlow and test them on the test split in BlinkFlow, as shown in Table IV. E-FlowFormer outperforms previous methods all-sided. E-Blender, which contains flexible objects, is the most challenging subset while our E-FlowFormer obtains the most significant AE improvement in this subset. It denotes that E-FlowFormer best exploits the sufficient training data provided in BlinkFlow. More qualitative results can be found in Fig. 6.

## VI. CONCLUSION

In this paper, we have proposed a novel simulator, BlinkSim, and a large-scale and diversiform dataset, Blink-Flow. Thanks to the well-integrated BlinkSim, BlinkFlow is two orders of magnitude ahead of other available datasets in diversity, so it provides a much better training and test set. We believe that BlinkFlow will push the limits of event-based optical flow estimation and BlinkSim can also benefit other event-based tasks. Moreover, we propose a novel transformer-based event flow estimation network, E-FlowFormer, which effectively learns from the training set of BlinkFlow and ranks 1st on mainstream benchmarks DSEC and MVSEC and the proposed BlinkFlow benchmark.

REFERENCES

[1] M. Almatrafi, R. Baldwin, K. Aizawa, and K. Hirakawa, "Distance Surface for Event-based Optical Flow," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 42, no. 7, pp. 1547–1556, 2020.

[2] P. Bardow, A. J. Davison, and S. Leutenegger, "Simultaneous Optical Flow and Intensity Estimation from An Event Camera," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2016, pp. 884–892.

[3] R. Benosman, C. Clercq, X. Lagorce, S.-H. Ieng, and C. Bartolozzi, "Event-based Visual Flow," *IEEE Transactions on Neural Networks and Learning Systems*, vol. 25, no. 2, pp. 407–417, 2013.

[4] R. Benosman, S.-H. Ieng, C. Clercq, C. Bartolozzi, and M. Srinivasan, "Asynchronous Frameless Event-based Optical Flow," *Neural Networks*, vol. 27, pp. 32–37, 2012.

[5] W. Bian, Z. Huang, X. Shi, Y. Dong, Y. Li, and H. Li, "Context-TAP: Tracking Any Point Demands Spatial Context Features," *arXiv preprint arXiv:2306.02000*, 2023.

[6] A. X. Chang, T. Funkhouser, L. Guibas, P. Hanrahan, Q. Huang, Z. Li, S. Savarese, M. Savva, S. Song, H. Su, *et al.*, "ShapeNet: An Information-rich 3D Model Repository," *arXiv preprint arXiv:1512.03012*, 2015.

[7] X. Chu, Z. Tian, Y. Wang, B. Zhang, H. Ren, X. Wei, H. Xia, and C. Shen, "Twins: Revisiting the Design of Spatial Attention in Vision Transformers," *Advances in Neural Information Processing Systems*, vol. 34, pp. 9355–9366, 2021.

[8] J. Cuadrado, U. Rançon, B. Cottereau, F. Barranco, and T. Masquelier, "Optical Flow Estimation with Event-based Cameras and Spiking Neural Networks," *Frontiers in Neuroscience*, 2023.

[9] M. Denninger, D. Winkelbauer, M. Sundermeyer, W. Boerdijk, M. Knauer, K. H. Strobl, M. Humt, and R. Triebel, "BlenderProc2: A Procedural Pipeline for Photorealistic Rendering," *Journal of Open Source Software*, vol. 8, no. 82, p. 4901, 2023. [Online]. Available: https://doi.org/10.21105/joss.04901

[10] Z. Ding, R. Zhao, J. Zhang, T. Gao, R. Xiong, Z. Yu, and T. Huang, "Spatio-temporal Recurrent Networks for Event-based Optical Flow Estimation," in *Proceedings of the AAAI Conference on Artificial Intelligence*, vol. 36, no. 1, 2022, pp. 525–533.

[11] A. Dosovitskiy, P. Fischer, E. Ilg, P. Hausser, C. Hazirbas, V. Golkov, P. Van Der Smagt, D. Cremers, and T. Brox, "FlowNet: Learning Optical Flow with Convolutional Networks," in *Proceedings of the IEEE International Conference on Computer Vision*, 2015, pp. 2758–2766.

[12] L. Downs, A. Francis, N. Koenig, B. Kinman, R. Hickman, K. Reymann, T. B. McHugh, and V. Vanhoucke, "Google Scanned Objects: A High-quality Dataset of 3D Scanned Household Items," in *Proceedings of the International Conference on Robotics and Automation*. IEEE, 2022, pp. 2553–2560.

[13] D. Falanga, K. Kleber, and D. Scaramuzza, "Dynamic Obstacle Avoidance for Quadrotors with Event Cameras," *Science Robotics*, vol. 5, no. 40, p. eaaz9712, 2020.

[14] G. Gallego, T. Delbrück, G. Orchard, C. Bartolozzi, B. Taba, A. Censi, S. Leutenegger, A. J. Davison, J. Conradt, K. Daniilidis, *et al.*, "Event-based Vision: A Survey," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 44, no. 1, pp. 154–180, 2020.

[15] D. Gehrig, M. Gehrig, J. Hidalgo-Carrió, and D. Scaramuzza, "Video to Events: Recycling Video Datasets for Event Cameras," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2020, pp. 3586–3595.

[16] D. Gehrig, H. Rebecq, G. Gallego, and D. Scaramuzza, "Asynchronous, Photometric Feature Tracking Using Events and Frames," in *Proceedings of the European Conference on Computer Vision*, 2018, pp. 750–765.

[17] M. Gehrig, W. Aarents, D. Gehrig, and D. Scaramuzza, "DSEC: A Stereo Event Camera Dataset for Driving Scenarios," *IEEE Robotics and Automation Letters*, vol. 6, no. 3, pp. 4947–4954, 2021.

[18] M. Gehrig, M. Millhäusler, D. Gehrig, and D. Scaramuzza, "E-RAFT: Dense Optical Flow from Event Cameras," in *Proceedings of the International Conference on 3D Vision*. IEEE, 2021, pp. 197–206.

[19] Y. Hu, S.-C. Liu, and T. Delbruck, "V2E: From Video Frames to Realistic DVS Events," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2021, pp. 1312–1321.

[20] Z. Huang, X. Shi, C. Zhang, Q. Wang, K. C. Cheung, H. Qin, J. Dai, and H. Li, "Flowformer: A Transformer Architecture for Optical Flow," in *Proceedings of the European Conference on Computer Vision*. Springer, 2022, pp. 668–685.

[21] P. Jiménez, F. Thomas, and C. Torras, "3D Collision Detection: A Survey," *Computers & Graphics*, vol. 25, no. 2, pp. 269–285, 2001.

[22] C. Lee, A. K. Kosta, A. Z. Zhu, K. Chaney, K. Daniilidis, and K. Roy, "Spike-FlowNet: Event-based Optical Flow Estimation with Energy-efficient Hybrid Neural Networks," in *Proceedings of the European Conference on Computer Vision*. Springer, 2020, pp. 366–382.

[23] Y. Li, X. Liu, W. Dong, H. Zhou, H. Bao, G. Zhang, Y. Zhang, and Z. Cui, "DELTAR: Depth Estimation from a Light-Weight ToF Sensor and RGB Image," in *Proceedings of the European Conference on Computer Vision*. Springer, 2022, pp. 619–636.

[24] Y. Li, H. Zhou, B. Yang, Y. Zhang, Z. Cui, H. Bao, and G. Zhang, "Graph-based Asynchronous Event Processing for Rapid Object Recognition," in *Proceedings of the IEEE International Conference on Computer Vision*, 2021, pp. 934–943.

[25] S. Lin, Y. Ma, Z. Guo, and B. Wen, "DVS-Voltmeter: Stochastic Process-Based Event Simulator for Dynamic Vision Sensors," in *Proceedings of the European Conference on Computer Vision*. Springer, 2022, pp. 578–593.

[26] T.-Y. Lin, M. Maire, S. Belongie, J. Hays, P. Perona, D. Ramanan, P. Dollár, and C. L. Zitnick, "Microsoft COCO: Common Objects in Context," in *Proceedings of the European Conference on Computer Vision*. Springer, 2014, pp. 740–755.

[27] M. Liu and T. Delbrück, "Adaptive Time-Slice Block-Matching Optical Flow Algorithm for Dynamic Vision Sensors," in *Proceedings of the British Machine Vision Conference*. BMVA Press, 2018, p. 88.

[28] N. Mayer, E. Ilg, P. Hausser, P. Fischer, D. Cremers, A. Dosovitskiy, and T. Brox, "A Large Dataset to Train Convolutional Networks for Disparity, Optical Flow, and Scene Flow Estimation," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2016, pp. 4040–4048.

[29] M. Muglikar, D. P. Moeys, and D. Scaramuzza, "Event Guided Depth Sensing," in *International Conference on 3D Vision*. IEEE, 2021, pp. 385–393.

[30] J. Ni, Y. Li, Z. Huang, H. Li, H. Bao, Z. Cui, and G. Zhang, "PATS: Patch Area Transportation with Subdivision for Local Feature Matching," in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2023, pp. 17776–17786.

[31] H. Rebecq, D. Gehrig, and D. Scaramuzza, "ESIM: An Open Event Camera Simulator," in *Proceedings of the Conference on Robot Learning*. PMLR, 2018, pp. 969–982.

[32] B. Rueckauer and T. Delbruck, "Evaluation of Event-based Algorithms for Optical Flow with Ground-truth from Inertial Measurement Sensor," *Frontiers in neuroscience*, vol. 10, p. 176, 2016.

[33] X. Shi, Z. Huang, W. Bian, D. Li, M. Zhang, K. C. Cheung, S. See, H. Qin, J. Dai, and H. Li, "Videoflow: Exploiting Temporal Cues for Multi-frame Optical Flow Estimation," *arXiv preprint arXiv:2303.08340*, 2023.

[34] X. Shi, Z. Huang, D. Li, M. Zhang, K. C. Cheung, S. See, H. Qin, J. Dai, and H. Li, "Flowformer++: Masked Cost Volume Autoencoding for Pretraining Optical Flow Estimation," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2023, pp. 1599–1610.

[35] S. Shiba, Y. Aoki, and G. Gallego, "Secrets of Event-based Optical Flow," in *Proceedings of the European Conference on Computer Vision*. Springer, 2022, pp. 628–645.

[36] D. Sun, X. Yang, M.-Y. Liu, and J. Kautz, "PWC-Net: Cnns for Optical Flow using Pyramid, Warping, and Cost Volume," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2018, pp. 8934–8943.

[37] S. Sun, G. Cioffi, C. De Visser, and D. Scaramuzza, "Autonomous Quadrotor Flight Despite Rotor Failure with Onboard Vision Sensors: Frames vs. Events," *IEEE Robotics and Automation Letters*, vol. 6, no. 2, pp. 580–587, 2021.

[38] Z. Teed and J. Deng, "RAFT: Recurrent All-pairs Field Transforms for Optical Flow," in *Proceedings of the European Conference on Computer Vision*. Springer, 2020, pp. 402–419.

[39] A. Vaswani, N. Shazeer, N. Parmar, J. Uszkoreit, L. Jones, A. N. Gomez, Ł. Kaiser, and I. Polosukhin, "Attention is all you need," *Advances in Neural Information Processing Systems*, vol. 30, 2017.

[40] W. Wang, D. Zhu, X. Wang, Y. Hu, Y. Qiu, C. Wang, Y. Hu, A. Kapoor, and S. Scherer, "TartanAir: A Dataset to Push the Limits of Visual Slam," in *Proceedings of the IEEE/RSJ International Conference on Intelligent Robots and Systems*. IEEE, 2020, pp. 4909–4916.

[41] B. Zhou, H. Zhao, X. Puig, S. Fidler, A. Barriuso, and A. Torralba, "Scene Parsing Through Ade20k Dataset," in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2017, pp. 633–641.

[42] A. Z. Zhu, D. Thakur, T. Özaslan, B. Pfrommer, V. Kumar, and K. Daniilidis, "The Multivehicle Stereo Event Camera Dataset: An Event Camera Dataset for 3D Perception," *IEEE Robotics and Automation Letters*, vol. 3, no. 3, pp. 2032–2039, 2018.

[43] A. Z. Zhu, L. Yuan, K. Chaney, and K. Daniilidis, "Ev-flownet: Self-supervised optical flow estimation for event-based cameras," in *Robotics: Science and Systems*, H. Kress-Gazit, S. S. Srinivasa, T. Howard, and N. Atanasov, Eds., 2018.

[44] Z. Zhu, L. Yuan, K. Chaney, and K. Daniilidis, "Unsupervised Event-based Learning of Optical Flow, Depth, and Egomotion," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2019, pp. 989–997.