

# EvHandPose: Event-based 3D Hand Pose Estimation with Sparse Supervision

Jianping Jiang<sup>†</sup>, Jiahe Li<sup>†</sup>, Baowen Zhang,  
Xiaoming Deng<sup>‡</sup>, *Member, IEEE*, Boxin Shi<sup>‡</sup>, *Senior Member, IEEE*

**Abstract**—Event camera shows great potential in 3D hand pose estimation, especially addressing the challenges of fast motion and high dynamic range in a low-power way. However, due to the asynchronous differential imaging mechanism, it is challenging to design event representation to encode hand motion information especially when the hands are not moving (causing motion ambiguity), and it is infeasible to fully annotate the temporally dense event stream. In this paper, we propose *EvHandPose* with novel hand flow representations in Event-to-Pose module for accurate hand pose estimation and alleviating the motion ambiguity issue. To solve the problem under sparse annotation, we design contrast maximization and hand-edge constraints in Pose-to-IWE (Image with Warped Events) module and formulate *EvHandPose* in a weakly-supervision framework. We further build *EvRealHands*, a large-scale real-world event-based hand pose dataset on several challenging scenes to bridge the real-synthetic domain gap. Experiments on *EvRealHands* demonstrate that *EvHandPose* outperforms previous event-based methods under all evaluation scenes, achieves accurate and stable hand pose estimation with high temporal resolution in fast motion and strong light scenes compared with RGB-based methods, generalizes well to outdoor scenes and another type of event camera, and shows the potential for the hand gesture recognition task.

**Index Terms**—3D hand pose estimation, event camera.

## 1 INTRODUCTION

HAND pose estimation plays a key role in computer vision, human-computer interaction, virtual reality, and robotics [1], [2]. Most of the existing vision-based hand pose estimation approaches adopt conventional imaging sensors such as RGB or RGB-D cameras. Although these sensors are widely available and effective for various hand pose relevant applications, due to the limitations from their imaging mechanisms, it is still challenging to achieve robust and accurate hand pose results in many real scenarios, especially for those with fast motions and abrupt changes of lighting conditions. Event cameras [3] that generate neuromorphic events by measuring per-pixel intensity changes asynchronously have shown great potential in many vision tasks such as detection [4], tracking [5], 3D reconstruction [6], flow estimation [7], simultaneous localization and mapping [8], [9], due to the advantages of high dynamic range (HDR, 120 dB), low power consumption, and high temporal resolution (up to 1  $\mu$ s). Recently, event cameras have also demonstrated their potential in robust 3D hand pose estimation for scenes with fast motions and broad dynamic range, because of their rich temporal information and HDR property [10].

Due to the lack of appearance and texture information in event streams, estimating hand poses using events faces three key challenges as shown in Figure 1 (C1)-(C3). First, from the data inherent properties perspective, asynchronous event streams

are spatially sparse and temporally dense, so existing image-based hand pose estimation solutions can not be directly applied to events. A straightforward solution is to reconstruct intensity frames from events via an “integration” process as [11] and [12]. However, such strategies would sacrifice the advantage of temporally fast and dense events, and it is still an open problem to design effective event representation and feature extractor for event-to-image conversion without losing these advantages. Moreover, event cameras suffer from the motion ambiguity issue, since the static hand parts can not generate events under the assumption of brightness consistency, which may result in that different hand poses generate the same event stream (only relative motion is recorded). Second, it is generally infeasible to annotate the asynchronous event stream using existing hand pose annotation techniques. For annotators, they could only provide sparsely annotated hand poses by sampling the key frames along the temporal dimension of event streams, so the event-based hand pose model is expected to be learned with sparse hand supervision in the temporal dimension. Recently, a pilot study [10] was conducted to learn 3D hand pose from event streams, which uses event simulators to generate synthetic dataset and evaluates the performance on real data. Although promising results are achieved, the third challenge is that the existing method [10] suffers from domain gap between synthetic and real data, since a small real dataset that can not support strong model generalization ability desired for different challenging scenarios.

To address the above challenges, we propose *EvHandPose* – a weakly-supervision framework to learn event-based 3D hand pose estimation under sparse supervision, as shown in Figure 1 (S1)-(S3). First, in order to deal with the asynchronous data format and motion ambiguity issue, we design specific hand flow representation and apply temporal information to regress the accurate hand pose. And we propose to approximate the hand optical flow in a sequence of event stream by embedding the interpolated hand models. Second, to address the sparse annotation challenge, we apply contrast

- J. Jiang, B. Shi are with National Key Laboratory for Multimedia Information Processing, National Engineering Research Center of Visual Technology, and AI Innovation Center, School of Computer Science, Peking University, Beijing 100871, China.
- J. Li, B. Zhang, X. Deng are with Beijing Key Laboratory of Human Computer Interactions, Institute of Software, Chinese Academy of Sciences, Beijing 100190, China, and University of Chinese Academy of Sciences, Beijing, China.
- <sup>†</sup> Equal contribution. <sup>‡</sup> Corresponding authors: shiboxin@pku.edu.cn, xiaoming@iscas.ac.cn

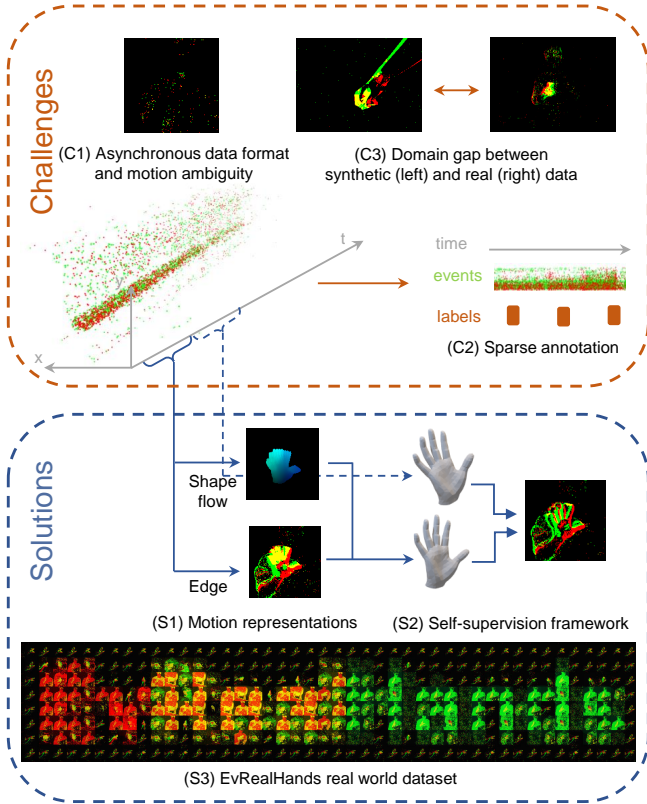


Fig. 1: Hand pose estimation from event data faces several challenges: (C1) asynchronous data format and motion ambiguity issue caused by inherent data properties of events, (C2) sparse annotation with current annotation techniques, (C3) domain gap between synthetic and real data. In order to tackle these issues, we propose *EvHandPose* by extracting novel motion representations including edge and shape flow from event streams with temporal information to predict hand poses (S1). We design a weakly-supervision framework to tackle the sparse annotation issue (S2). Furthermore, we present the first large real-world event-based hand pose dataset *EvRealHands* (S3) (zoom-in for details).

maximization and hand-edge constraints in a weakly-supervision framework. We first train our Event-to-Pose module on labeled data to get prior knowledge and then fine-tune the model by proposing a weakly-supervision framework, which is supervised by a Pose-to-IWE (Image with Warped Events) module during training. Finally, to address the domain gap issue limited by synthetic datasets, we construct *EvRealHands*, the first large-scale real event-based hand pose dataset under various challenging scenarios. Experiments demonstrate that our event camera based solution outperforms conventional RGB camera based solutions in challenging scenarios such as fast motion and strong light with high temporal resolution, and generalize well to outdoor scenes and other event cameras. Our dataset, code and models will be made public after acceptance.

Our main contributions are threefold:

- Novel hand flow representations: Asynchronous event data are effectively processed in Event-to-Pose module for accurate hand pose estimation and alleviating the motion ambiguity issue.
- Weakly-supervision framework: Pose-to-IWE module is complementarily designed to tackle the sparse annotation

issue with contrast maximization and hand-edge constraints specially designed for the hand.

- Dataset: The first large-scale real dataset *EvRealHands*, consisting of 79 minutes event sequences in indoor and outdoor scenes, 425 K RGB images, with accurate 3D pose and shape annotations is built to tackle the domain gap issue and facilitate future research.

## 2 RELATED WORK

### 2.1 3D Hand Pose Estimation

Prior arts of 3D hand pose estimation methods consist of appearance-based methods [13], [14], [15], [16], [17] and model-based approaches [18], [19], [20], [21], [22]. Appearance-based methods directly learn the mapping from image space to hand pose space by training on large-scale datasets, which often require large-scale annotated dataset. Model-based methods are designed to fit a prior hand parametric model such as MANO [23] to the image observation. However, these methods require a good initialization for optimization. In order to reduce the expensive annotation cost of large scale dataset, several efforts are adopted to learn hand pose from unlabeled data in weakly-supervised [24], [25], [26], [27] or self-supervised [28], [29], [30] manners. These methods use additional constraints such as segmentation mask, depth, appearance, and biomechanics consistency to ensure the quality of the predicted hand pose. When dealing with hand pose estimation from video sequences, learning from sparse-labeled data is effective to reduce the burden of full annotations for video sequences, and the existing methods often leverage motion or photometric consistency as constraints between labeled and unlabeled frames [31], [32] to train models. However, there is no weakly-supervised or self-supervised framework specially designed for event-based hand pose estimation.

### 2.2 Event-based Pose Estimation

The first challenge for event-based 3D hand pose estimation is the real-synthetic domain gap. Previous synthetic event-based dataset [10] faces two issues: 1) lack of hand pose diversity and 2) domain gap of synthetic data due to the event simulator [33]. Although domain adaptation methods, such as NGA [34] and EvTransfer [35], tend to utilize the shared information between events and images to bridge the gap, they can hardly achieve superior results compared with supervised methods. Since there is no large-scale real event-based hand dataset with 3D pose annotation, it motivates us to build a new real event-based hand dataset. In this work, we select multiple ubiquitous color cameras to get 3D hand pose annotations of event camera.

Since the events are mainly generated from the edge of the moving objects, they can preserve the information of hand movement. Existing methods use different representations of event streams, such as event count image (ECI) [36], time surface [37], voxel [38], time-ordered recent event volumes (TORE) [39], event spike tensor (EST) [40], Matrix-LSTM [41] *etc.* for learning models, while they are not specifically designed to represent hand motions for 3D hand pose estimation. Therefore, a key problem for event-based hand pose estimation is to leverage edge and movement features for pose estimation.

Optical flow is an effective feature in event-based vision tasks, and it can serve as a useful approximation to hand motion fields when events are used for hand pose estimation. Although optical

flow from event streams has been extensively studied [7], [38], [42], to the best of our knowledge, no existing work is specifically designed for hand motion. EV-FlowNet [42] learns self-supervised optical flow by photometric consistency of gray-scale images from the DAVIS camera [3], but it fails when capturing gray-scale images from fast motion or HDR scenes. To tackle this, Zhu *et al.* [38] learn an effective optical flow model in an unsupervised manner based on the contrast maximization method. However, it is still an open problem to estimate the hand specific optical flow (*hand flow*) by leveraging the hand shape model, and how to use hand flow to boost the hand pose estimation performance.

In the realm of pose estimation, event camera is firstly applied to human pose estimation. EventCap [43] utilizes initial pose estimation via gray-scale images to reconstruct high-frequency human pose motion by event trajectories. However, EventCap [43] requires additional gray-scale images (usually with generated qualities for scenes that events demonstrate advantages) as input, and could lead to unstable pose results under challenging illumination conditions or fast motions. EventHPE [44] estimates human pose from only event data, but predicting relative poses of neighboring event frames will lead to accumulated errors as time goes on.

When events are used for hand pose estimation, it is more challenging due to high degree-of-freedom hand motion, heavy self-occlusion and significant similarity between different hand parts. EventHands [10] is proposed to estimate 3D hand pose estimation from event streams. It estimates MANO parameters from Locally-Normalised Event Surfaces (LNES), which assigns higher weights to events close to the target timestamp and forms an event frame. However, it fails to resolve the motion ambiguity in pose estimation, *i.e.*, static hand parts will not trigger events, which leads to spatial uncertainty of these static parts. Jalees *et al.* [45] also propose a model-based approach to estimate 3D hand pose from event streams, which fits a hand model via nonlinear optimization and enforces event stream simulator to generate similar event frames to the real event camera. However, the hand tracking by optimization requires good initialization, and it is prone to be stuck at local minimum due to the non-convex optimization of the object function.

Besides, due to the asynchronous imaging mechanism between color cameras and event cameras, hand pose annotations from color cameras are very sparse compared to the temporally dense event streams. To deal with the sparse annotation issue, we need to design effective constraints to enforce the predicted hand poses of the unlabeled frames to be plausible and well aligned with events.

### 3 PRELIMINARIES

#### 3.1 Hand Model Representation

We adopt a differentiable hand parametric MANO model [23], which can be formulated as:

$$M(\beta, \theta) = W(T_P(\beta, \theta), J(\beta), \theta, \mathcal{W}), \quad (1)$$

where a skinning function  $W$  (linear blend skinning [46]) is applied to an articulated rigged body mesh with shape  $T_P$ , joint locations  $J$  defining a kinematic tree, pose parameter  $\theta$ , shape parameter  $\beta$ , and blend weights  $\mathcal{W}$ . Given pose and shape parameters, MANO model outputs a hand mesh with 778 vertices  $\mathbf{v} = M(\beta, \theta) \in \mathbb{R}^{778 \times 3}$ , and 3D joints  $\mathbf{J}_{3D} \in \mathbb{R}^{21 \times 3}$  can be recovered by  $\mathbf{J}_{3D} = J_{\text{reg}}(\beta, \theta)$ , where  $J_{\text{reg}}$  is the regression function from the hand mesh and pose parameters to 3D joints. Hereafter, we denote  $\varphi = (\beta, \theta)$  to be *hand parameters* for easy reference.

#### 3.2 Image Formation and Motion Ambiguity

Event camera encodes the changes of per-pixel intensity  $I(x, y, t)$  into event streams. An event  $e_i = (x_i, y_i, t_i, p_i)$  is triggered at pixel  $(x_i, y_i)$  at time  $t_i$  when the logarithmic radiance change reaches the threshold:

$$\log I(x_i, y_i, t_i) - \log I(x_i, y_i, t_p) = p_i C, \quad (2)$$

where  $t_p$  is the timestamp of the last event fired at pixel  $(x_i, y_i)$ ,  $p_i \in \{-1, 1\}$  is the polarity of intensity change,  $C$  is the contrast threshold.

Due to the differential imaging mechanism of the event camera, the static hand part will not generate events under constant lighting. Events  $E_{t_{n-1}, t_n}$  within time interval  $[t_{n-1}, t_n]$  may correspond to multiple hand poses (we call it “motion ambiguity” in event-based hand pose problem), because the static hand parts can have multiple pose states and their relative pose changes reflected by events could look the same. Several efforts such as EventHPE [44] are conducted to predict relative pose  $\Delta\theta_{t_n}$  between the current frame and the previous frame  $\varphi_{t_{n-1}}$ . However, these approaches can not resolve the motion ambiguity issue, because the relative poses are still affected by the absolute hand poses even if they generate the same event stream.

One way to alleviate the challenge is to encode the previous event data to assist in estimating the current hand pose, because event camera is more likely to record complete hand movements over a long interval. In this paper, we use a Conv-GRU [47] recurrent model to embed the previous sequence features. We divide the asynchronous event sequence into several sub-segments for hand pose prediction. Given an event stream segment ranging from  $t_0$  to  $t_N$ , we split the segment into  $N$  sub-segments  $\{[t_{i-1}, t_i]\}_{i=1}^N$  with the same duration. Each sub-segment of the event stream will be represented as event frame and event volume for hand pose estimation.

## 4 METHODS

As shown in Figure 2, EvHandPose pipeline consists of two main parts: Event-to-Pose module and Pose-to-IWE (Image with Warped Events) module. The Event-to-Pose module is designed to predict hand pose from the asynchronous event stream in Section 4.1. It represents the edge feature as LNES [10] and predicts the shape flow (*i.e.*, a new optical flow specially designed for hands) with the proposed FlowNet to explore motion information from the asynchronous event stream. After fusing the edge feature and shape flow by a feature encoder (ResNet34 [48] backbone), a Conv-GRU model is applied to alleviate the motion ambiguity issue by involving the temporal correlation between the neighboring event frames. The Pose-to-IWE module is then designed to supervise Event-to-Pose module during training. The key idea is to use the sparse annotations of hand shapes and hand joints as supervision when they are available, use a contrast maximization loss and hand-edge loss on the unannotated event frames to enhance hand pose estimation performance, and use hand flow loss for all event frames. In Section 4.2, we elaborate on the details of these loss functions for annotated and unannotated event frames, respectively.

#### 4.1 Motion Representations

Events in time interval  $[t_{n-1}, t_n]$  contains two types of hand motion aware information (*i.e.*, optical flow and edge) under photometric consistency assumption. By recording the brightness

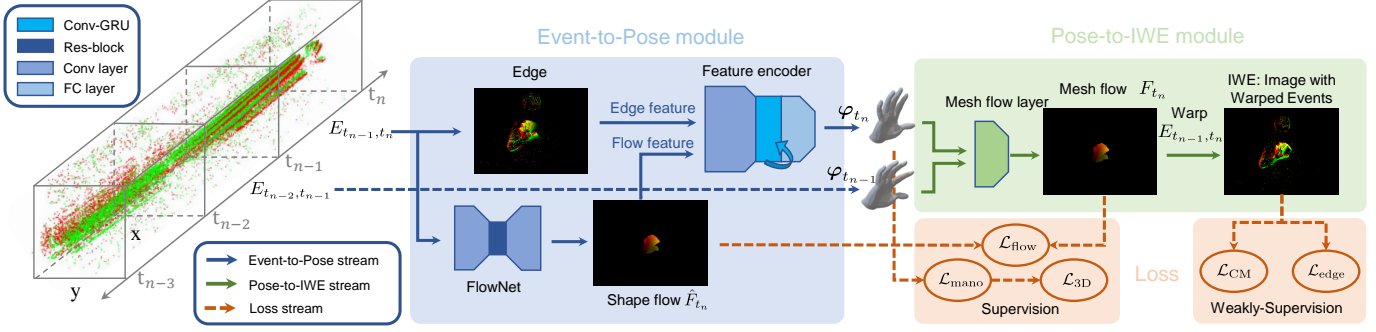


Fig. 2: Overview of EvHandPose pipeline. Event streams are firstly split into several sub-segments  $\{E_{t_{k-1}, t_k}\}$ , and they are represented as event frames and volumes. Event-to-Pose module extracts edge of each frame and shape flow  $\hat{F}_{t_n}$  of each volume by FlowNet, and predicts the hand parameters  $\varphi_{t_n}$  of the current event timestamp. Given sparsely annotated and unannotated event frames, we learn the hand pose estimation model by designing several weakly-supervision loss functions. Specially, we design a differential mesh flow layer (details in Figure 3) in Pose-to-IWE module to compute mesh flow  $F_{t_n}$  with the neighboring hand parameters  $\varphi_{t_n}$  and  $\varphi_{t_{n-1}}$ . In order to address the issue of sparse annotation, we warp the events in  $E_{t_{n-1}, t_n}$  with estimated hand flow to get IWE, and adopt contrast maximization loss and hand-edge loss upon IWE for weakly-supervision. In the 2D visualization of events in this paper, red pixels represent positive events, green pixels represent negative events, and yellow pixels represent both types of events.

difference between the hand and background, events in  $[t_{n-1}, t_n]$  describe the hand movement and can be used to resolve the motion ambiguity issue of hand pose estimation by including observations from previous status, and events at a specific timestamp  $t_n$  indicates the accurate hand edges. Therefore, we propose new motion representations that include both hand flow and edge information for event-based hand pose estimation.

As edges are naturally encoded in events, we utilize the edge information in Event-to-Pose module for accurate pose estimation. Specifically, we follow EventHands [10] to use the effective 2D representation LNES. The LNES image  $I$  is updated by iterating each event  $e_i$  in  $E_{t_{n-1}, t_n}$ :

$$I(x_i, y_i, p_i) = \frac{t_i - t_{n-1}}{t_n - t_{n-1}}. \quad (3)$$

The LNES can reserve the latest event information because it will be overwritten when an event with the same polarity occurs at that location.

#### 4.1.1 Hand Flow from Events

We use two types of representation of hand flow from events. The first type is the model-based hand flow  $F$  (denoted as *mesh flow*) by designing differential hand flow layer with the neighboring hand parameters  $\varphi_{t_n}$  and  $\varphi_{t_{n-1}}$ , shown in Figure 3. The second type of hand flow  $\hat{F}$  (denoted as *shape flow*) is obtained by the optical flow predicted by proposing a FlowNet (bottom right of Event-to-Pose module in Figure 2), which is learned using mesh flow as supervision.

**Mesh Flow.** Desired motion representations for event-based hand modeling should precisely describe the hand movement and ignore disturbance from the background. To achieve this, we calculate the dense hand flow by projecting two sequential MANO meshes on the image, and compute the displacements of corresponding vertices like [44] and [32].

Figure 3 illustrates hand flow computation process. Given two MANO models  $M(\beta_{t_{n-1}}, \theta_{t_{n-1}})$  and  $M(\beta_{t_n}, \theta_{t_n})$  at two event moments, we get  $K + 1$  hand shapes  $\{M_k\}_{k=0}^K$  with hand

parameters  $\{\varphi_k\}_{k=0}^K$  by linear interpolation of the hand parameters of the two MANO models as follows:

$$t_k = (1 - \frac{k}{K}) \cdot t_{n-1} + \frac{k}{K} \cdot t_n, \quad k = 0, 1, \dots, K. \quad (4)$$

$$\beta_k = (1 - \frac{k}{K}) \cdot \beta_{t_{n-1}} + \frac{k}{K} \cdot \beta_{t_n}, \quad k = 0, 1, \dots, K. \quad (5)$$

$$\theta_k = \text{slerp}(\theta_{t_{n-1}}, \theta_{t_n}, \frac{k}{K}), \quad k = 0, 1, \dots, K. \quad (6)$$

where spherical linear interpolation (*i.e.*,  $\text{slerp}(\cdot)$ ) [49] is used for pose parameters  $\theta$ , and  $\varphi_k = (\theta_k, \beta_k, t_k)$  is the interpolated MANO parameters at time  $t_k$ .

For each interpolated hand mesh  $M_k$ , we can compute the speed  $\mathbf{S}_k \in \mathbb{R}^{778 \times 2}$  of all vertices of the hand mesh on the image as follows:

$$\mathbf{S}_{k,i} = \frac{K \cdot \pi(\mathbf{v}_{K,i} - \mathbf{v}_{k,i})}{(K - k) \cdot (t_n - t_{n-1})}, \quad k = 0, 1, 2, \dots, K - 1, \quad (7)$$

where  $\mathbf{v}_{k,i}$  is the  $i$ -th vertex on the hand mesh  $M_k$ , and  $\pi$  is the camera projection matrix.

Since the hand mesh represented by MANO model is a triangular mesh with 778 vertices, the displacements of the corresponding vertices of neighboring interpolated hand mesh  $M_k$  can not get dense optical flow. In order to deal with this issue, we get the corresponding triangle of each pixel  $(x, y)$  by the first intersection of the ray from the camera optical center to the pixel  $(x, y)$  and the hand mesh, and calculate the barycentric coordinates  $b_{\mathbf{v}_j}(x, y)$  of the pixel  $(x, y)$ . Then we can compute the dense optical flow by the weighted sum of the vertex speed  $\mathbf{S}_k$  with barycentric weights  $b_{\mathbf{v}_j}$  as follows:

$$F_k(x, y) = \sum_{\mathbf{v}_j \in V_{\text{bary}}(x, y)} b_{\mathbf{v}_j}(x, y) \mathbf{S}_{k,j}, \quad (8)$$

where  $V_{\text{bary}}(x, y)$  is the vertices of the corresponding triangle of pixel  $(x, y)$ . Then our mesh flow representation  $F(x, y)$ , which is differentiable with respect to hand parameters, is calculated as:

$$F(x, y) = \frac{\sum_{k=0}^{K-1} F_k(x, y)}{\sum_{k=0}^{K-1} \mathbb{1}(F_k(x, y))}, \quad (9)$$

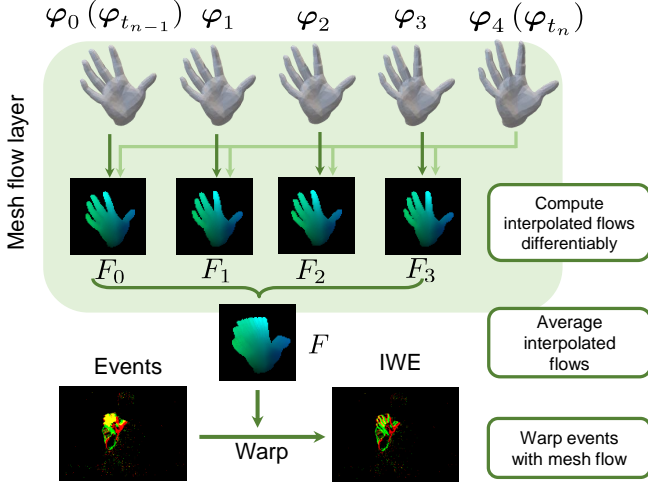


Fig. 3: Illustration of Pose-to-IWE module. Given MANO parameters  $\varphi_{t_{n-1}}$  and  $\varphi_{t_n}$ , we first conduct interpolation to get interpolated MANO parameters  $\{\varphi_k | k = 0, 1, \dots, K\}$ ,  $K = 4$  here. Then we compute the interpolated flows  $\{F_k\}$  by projecting the movements between each hand mesh  $M_k (k = 0, 1, \dots, K - 1)$  and the last hand mesh  $M_K$  to the image plane. Lastly, we average all the interpolated flows to get the final hand mesh flow  $F$ .

where  $\mathbb{1}$  is the function to indicate whether there is flow at pixel  $(x, y)$ . Events are mainly triggered from the edge of the hand under constant lighting condition. If we warp a period of events to a certain timestamp, events from the same position on the hand will gather together to form sharp hand edges and accumulated an image with warped events. Figure 3 illustrates that mesh flow layer and IWEs using our mesh flow have sharp edges, which will be used as the contrast maximization and hand-edge constraints in our weakly-supervision framework.

**Shape Flow.** While mesh flow is obtained from the mesh flow layer and serves for weakly-supervision losses, Event-to-Pose module needs to directly extract motion information from event streams for accurate hand pose estimation. Thus Event-to-Pose module uses FlowNet which is a U-Net like CNN backbone to predict hand shape flow from event volumes [38]. Being different from [38], our FlowNet is supervised by mesh flow and only learns hand movement information, not background movement information, thus reserving high quality hand movement information. We apply the endpoint error (EPE) between shape flow  $\hat{F}_{t_n}$  and mesh flow  $F_{t_n}$  as loss function, which is the normal measurement for optical flow estimation:

$$\mathcal{L}_{\text{flow}} = \sum_n \text{EPE}(\hat{F}_{t_n}, F_{t_n}). \quad (10)$$

## 4.2 Weakly-supervision Learning Framework

The annotations are only available for sparse subset of event frames, because the annotations are provided with sparse frames of multiple color cameras (details in Section 5). In this section, we formulate the weakly-supervision framework that can learn the hand pose model using both the sparse annotated and unannotated event frames.

### 4.2.1 Learning under Sparse Annotations

Event-to-Pose module of EvHandPose extracts edge feature and hand flow from event streams, and predicts the hand MANO

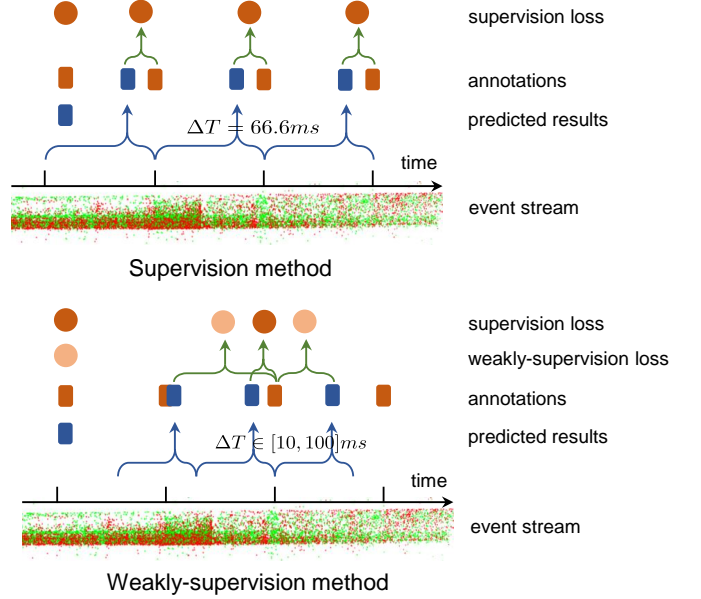


Fig. 4: Comparison of supervision and weakly-supervision methods. In supervision method, the model can only learn hand poses at annotated timestamps and the interval of each event sub-segment has to be strictly overlapped with the annotated period. In weakly-supervision method, we only need partial labels and the outputs of unlabeled sub-segments are constrained by weakly-supervision loss. Therefore, EvHandPose can learn hand poses at any timestamp near the annotations with dynamic time window for event sub-segment representation.

parameters. When annotations of 3D joints and hand MANO parameters are available, we use 3D joint loss  $\mathcal{L}_{3D}$  and hand MANO parameter loss  $\mathcal{L}_{\text{mano}}$  for model learning:

$$\mathcal{L}_{3D} = \sum_n \|J_{\text{reg}}(\hat{\beta}_n, \hat{\theta}_n) - \mathbf{J}_{3D_n}\|^2, \quad (11)$$

$$\mathcal{L}_{\text{mano}} = \sum_n (\lambda_\beta \|\hat{\beta}_n - \beta_n\|^2 + \lambda_\theta \|\hat{\theta}_n - \theta_n\|^2), \quad (12)$$

where  $\mathbf{J}_{3D_n}$  is the ground-truth 3D joints vector, MANO parameters with hat ( $\hat{\beta}, \hat{\theta}$ ) are prediction results and those without hat ( $\beta, \theta$ ) are ground-truth,  $\lambda_\beta$  and  $\lambda_\theta$  are loss weights and they are set to 0.25 and 5.0, respectively.

Therefore, the loss function  $\mathcal{L}_{\text{super}}$  to learn under supervision can be formulated as follows:

$$\mathcal{L}_{\text{super}} = \lambda_{3D} \mathcal{L}_{3D} + \lambda_{\text{mano}} \mathcal{L}_{\text{mano}}, \quad (13)$$

where  $\lambda_{3D}$  and  $\lambda_{\text{mano}}$  are loss weights, and they are set to 1000 and 0.1 in our experiments, respectively.

### 4.2.2 Learning with Weakly-supervision

Since we get the annotated hand pose dataset EvRealHands using multiple RGB cameras, hand pose can be annotated at a pre-defined fps (e.g., 15 fps in EvRealHands). Due to the asynchronous image mechanism of event camera, the event streams can only have sparse hand pose annotations. In order to supervise the unannotated event sub-segment and leverage the dense event streams, we propose a weakly-supervision framework to effectively utilize the event with sparse annotations as shown in Figure 4. Specially, we design a contrast maximization loss to constrain the movement of the hand

over time, and an hand-edge loss to learn hand pose at predicted time, and adopt smooth loss to encourage the smoothness of the sequential predicted hand poses.

**Contrast Maximization Loss.** Inspired by the contrast maximization method that was originally proposed to handle motion compensation problem [50], we design a new contrast maximization loss (denoted as CM loss) for weakly-supervision. The key idea of CM loss is to enforce the IWE is sharp and has good contrast, an effective global statistical priors of IWE.

Event-to-Pose module can predict sequential hand MANO parameters  $\{\varphi_{t_n} | n = 1, \dots, N\}$ . For each neighboring pair of parameters  $\{\varphi_{t_{n-1}}, \varphi_{t_n}\}$ , we can compute the mesh flow  $F_n$ , and warp the events in  $[t_{n-1}, t_n]$  to get the image with warped events  $\text{IWE}_{n,\hat{t}}$  at timestamp  $\hat{t}$  as follows:

$$\text{IWE}_{n,\hat{t}} = \text{Warp}(E_{t_{n-1},t_n}; F_n; \hat{t}), \quad (14)$$

where  $\text{Warp}(\cdot, \cdot, \cdot)$  is composed of two steps: We adopt commonly-used flow-constancy hypothesis [51], transfer each event  $e_i$  in  $E_{t_{n-1},t_n}$  to target time  $\hat{t}_i$ , and get the warped event  $\hat{e}_i$  in new position:

$$\hat{x}_i = x_i + F_n(x_i) \cdot (\hat{t}_i - t_i). \quad (15)$$

We warp events form the IWE by counting events at pixel  $\mathbf{x}$ :

$$\text{IWE}_{n,\hat{t}}(\mathbf{x}, p) = \sum \delta(\mathbf{x} - \hat{x}_i) \delta(p - p_i), \quad (16)$$

where each polarity  $p$  corresponds to one channel of the IWE.

Then we use the focus loss functions using variance measure in [52] to evaluate how well events are aligned along the hand flow, *i.e.*, the contrast of the image with warped event, as follow:

$$\text{Var}(I) = \frac{1}{N_I} \sum_{i,j} (h_{ij} - \mu_I)^2, \quad (17)$$

where  $N_I$  is the number of pixels of image  $I$ , and  $\mu_I = \frac{1}{N_I} \sum_{i,j} h_{ij}$  is the mean value of image  $I$ .

The CM loss is finally defined by computing the contrasts in both forward and backward direction as follows:

$$\mathcal{L}_{\text{CM}} = - \sum_n (\text{Var}(\text{IWE}_{n,t_n}) + \text{Var}(\text{IWE}_{n,t_{n-1}})), \quad (18)$$

where  $\text{IWE}_{n,t_n}$  and  $\text{IWE}_{n,t_{n-1}}$  are the forward and backward IWEs at timestamp  $t_n$  and  $t_{n-1}$ , respectively.

**Hand-edge Loss.** As shown in Figure 3, IWE should demonstrate sharp hand edges, and this inspires us to design an edge loss to enforce the alignment of the projection of predicted hand shape and the edge of IWE. Intuitively, such a constraint can bring a benefit that the local details of the reconstructed hand shape are well consistent to the event observation. However, it is challenging to find reliable correspondences of events and hand mesh vertices due to two issues: 1) Events are not aligned at the same time with the hand shape and 2) fingers will gather together in optimization because an event may have several corresponding vertices. To solve the above misalignment issues, we first warp the events in  $[t_{n-1}, t_n]$  by mesh flow and get IWE at timestamp  $t_n$ , then find the correspondence between pixels of IWE and hand mesh vertices by searching the proper hand vertex for each IWE pixel. In order to reduce computation cost and make our method robust to noisy events, we only select  $M$  IWE key pixels with the most accumulated events to calculate the edge loss. For each key pixel  $\text{IWE}(x_i, y_i)$ , we get a candidate set of hand mesh vertices  $P_i$  that can be projected within a local neighborhood of the pixel  $(x_i, y_i)$

(*e.g.*,  $12 \times 12$ ). For these vertices, we select a candidate vertex using both orientation weight  $w_j^o$  and motion weight  $w_j^m$ .

Since hand mesh vertices with the normal vectors perpendicular to the camera ray will generate more events, we design the orientation weight to each vertex  $\mathbf{v}_j$  as follows:

$$w_j^o = b_{\text{orient}} - \|\cos(\mathbf{n}_j, \mathbf{d}_j)\|, \quad (19)$$

where  $\mathbf{n}_j$  is the normal vector at  $\mathbf{v}_j$ ,  $\mathbf{d}_j$  is the camera ray pointing from the camera center to  $\mathbf{v}_j$ , constant and  $b_{\text{orient}}$  can be any value slightly over 1 (set as 1.2 in our experiments).

We further observe that vertices with large motion displacements often produce more events, so we design the motion weight of vertex  $\mathbf{v}_j$  as follows:

$$w_j^m = \sqrt{\|\pi(\mathbf{v}_{j,t_{n-1}}) - \pi(\mathbf{v}_{j,t_n})\|^2 + b_{\text{motion}}}, \quad (20)$$

where  $\pi(\cdot)$  is the projection function,  $\mathbf{v}_{j,t_{n-1}}$  is the vertex at timestamp  $t_{n-1}$ , and  $b_{\text{motion}}$  can be any constant slightly over 0 (set as 4 in our experiments).

Next, we design a metric  $\text{metric}_{i,j}$  with the orientation weight  $w_j^o$  and the motion weight  $w_j^m$ , and select the corresponding vertex  $\mathbf{v}_{P_i}$  of  $\text{IWE}(x_i, y_i)$  with the highest metric value. Such a metric can be defined as:

$$\text{metric}_{i,j} = \frac{w_j^o \cdot w_j^m}{\|\pi(\mathbf{v}_{j,t_n}) - \text{IWE}(x_i, y_i)\|^2 + b_d}, \quad (21)$$

where  $b_d = 4.0$  is the constant shift of distance.

The hand-edge loss is finally defined as :

$$\mathcal{L}_{\text{hand-edge}} = \sum_n \frac{\sum_i^M w_{P_i}^o \cdot w_{P_i}^m \cdot \|\pi(\mathbf{v}_{P_i,t_n}) - \text{IWE}(x_i, y_i)\|^2}{\sum_i^M w_{P_i}^o \cdot w_{P_i}^m}. \quad (22)$$

**Smooth Loss.** To ensure the smoothness of the predicted hand parameters of event streams, we use a smooth loss  $\mathcal{L}_{\text{smooth}}$ :

$$\mathcal{L}_{\text{smooth}} = \sum_n \text{sigmoid}(\max(0, \lambda_\beta \|\hat{\beta}_n - \hat{\beta}_{n-1}\|^2 + \lambda_\theta \|\hat{\theta}_n - \hat{\theta}_{n-1}\|^2 - b_{\text{smooth}})), \quad (23)$$

where the sigmoid function is used to avoid the gradients of  $\mathcal{L}_{\text{smooth}}$  to be large,  $b_{\text{smooth}} = 0.5$  is the empirical margin of the smooth loss, and the hyper-parameters  $\lambda_\beta$  and  $\lambda_\theta$  are set to 0.2 and 1.0.

In summary, the predicted results from the Event-to-Pose module are fed into the Pose-to-IWE module. Those results of the unlabeled segments will be constrained by the weakly-supervision loss  $\mathcal{L}_{\text{weakly}}$ :

$$\mathcal{L}_{\text{weakly}} = \lambda_{\text{CM}} \mathcal{L}_{\text{CM}} + \lambda_{\text{hand-edge}} \mathcal{L}_{\text{hand-edge}} + \lambda_{\text{smooth}} \mathcal{L}_{\text{smooth}}, \quad (24)$$

where  $\lambda_{\text{CM}}$ ,  $\lambda_{\text{smooth}}$ , and  $\lambda_{\text{hand-edge}}$  are loss weights, and they are set to 3, 0.1, 0.2, respectively.

#### 4.2.3 Training Strategy


Our training strategy consists of three steps. First, we train FlowNet under supervision with  $\mathcal{L}_{\text{flow}}$  and fix its parameters in the following steps. Then, we train EvHandPose under supervision with  $\mathcal{L}_{\text{super}}$  to get reasonable good parameter initialization. Finally, we train EvHandPose with both labeled data and unlabeled data with  $\mathcal{L}_{\text{all}}$ , which contains supervision loss and weakly-supervision loss and can be formulated as follows:

$$\mathcal{L}_{\text{all}} = \lambda_{\text{super}} \mathcal{L}_{\text{super}} + \lambda_{\text{weakly}} \mathcal{L}_{\text{weakly}}, \quad (25)$$


where  $\lambda_{\text{super}}$  and  $\lambda_{\text{weakly}}$  are weighting factors, and set to 2.0, 1.0.

TABLE 1: Comparison of existing event-based hand pose estimation dataset and EvRealHands. EvRealHands involves multiple challenging real-world scenes. We measure the amount of event data by their sequences’ lengths.

Dataset	Source	Type	Resolution	Subjects	Annotation	Frames	Scenes		
							Strong light	Fast motion	Flash
NYU [53]	Depth	Real	640 × 480	2	Tracking	243 K	None	None	None
STB [54]	RGB + Depth	Real	640 × 480	1	Human	36 K	None	None	None
BigHand2.2M [55]	Depth	Real	640 × 480	10	Marker	2.2 M	None	None	None
FreiHAND [56]	RGB	Real	224 × 224	32	Semi-automatic	134 K	None	None	None
InterHand2.6M [57]	RGB	Real	512 × 334	27	Human+machine	2.6 M	None	None	None
ContactPose [58]	RGB + Depth	Real	1920 × 1080	50	Human+machine	2.9 M	None	None	None
EventHands [10]	Event	Synthetic	240 × 180	1	Synthetic	1.24 M	None	12.6 s	None
<b>EvRealHands</b>	<b>Event + RGB*</b>	<b>Real</b>	<b>346 × 260</b>	<b>10</b>	<b>Human+machine</b>	<b>425 K</b>	<b>558.5 s</b>	<b>122.4 s</b>	<b>316.7 s</b>



(a) Multi-camera system



(b) DAVIS346 & FLIR

(c) PROPHESEE & FLIR

Fig. 5: Capture system for creating EvRealHands dataset. In the multi-camera system (a), seven FLIR BFS-U3-51S5 cameras (yellow) are synchronized to capture multi-view RGB images and the event camera DAVIS346 (red) captures the event streams in indoor scenes. We also attach an Azure Kinect DK (white) as an auxiliary camera to use the depth image only for accurate geometric calibration of the camera system. In the hybrid camera system for capturing outdoor sequences, we use the beam-splitter settings composed of the RGB camera FLIR BFS-U3-51S5 and the event camera, DAVIS346 Mono in (b) or PROPHESEE GEN 4.0 in (c).

## 5 DATASET

We present the first real-world event-based hand pose dataset EvRealHands. Our dataset contains event streams and RGB image with annotations of 3D hand pose and hand shape parameter, and it contains hand poses of commonly used hand gestures. EvRealHands is collected under common scenarios (*e.g.*, normal lighting and hand movements) and challenging scenarios such as strong light, flash and fast motion scenes, indoors and outdoors, which are overlooked in previous hand pose datasets as shown in

Table 1. Please refer to Appendix B in the supplemental material for information about detailed attributes, adequacy of data, scanned hand meshes, and hand gesture visualization and annotation.

### 5.1 Data Acquisition

**Capture System Setup.** For indoor scenes, we employ a multi-camera fashion similar to [59] and [60]. As shown in Figure 5 (a), the multi-camera system consists of 7 RGB cameras (FLIR, 2660×2300 pixels) and an event camera (DAVIS346, 346×260 pixels). We synchronize all the cameras with an external TTL signal of 15 Hz. For calibration, we use a moving calibration chessboard and calculate the intrinsic and extrinsic parameters following [61] with RGB frames from FLIR, APS frames from DAVIS346 and depth frames from an auxiliary RGB-D camera (Azure Kinect DK). The average pixel root mean square error on all the images in calibration reaches 0.184. Since building a multi-camera system outdoor faces several challenges, such as power supply, we build a hybrid camera system in a compact form, as shown in Figure 5 (b) and (c), to collect outdoor sequences for evaluation. The system is composed of an event camera (DAVIS346 Mono or PROPHESEE GEN 4.0) and an RGB camera (FLIR BFS-U3-51S5) via a beam-splitter (Thorlabs CCM1-BS013) with 50% optical splitting [62]. The beam-splitter settings can assure the event streams and RGB images are of approximately the same view towards hand motions.

**Hand Gestures.** We collect 102 sequences from 10 subjects in total with duration ranging from 20 s to 60 s. For each subject, we collect 3 sequences for each hand under normal lighting, which consists of 15 pre-defined hand poses similar to [63] about 40 s (fixed poses), random hand poses about 20 s (random short), random hand poses about 60 s (random long). As shown in Figure 6, the hand pose variance in random pose sequences is larger than that in fixed pose sequences.

**Scenes with Challenging Illuminations and Motions.** We capture 12 sequences for strong light scenes and 8 sequences for flash scenes from 2 subjects in indoor scenes. We utilize two glare flashlights with 2000 lumen. We keep the flashlights on in strong light scenes and make the flashlights flash at a rate of 1 Hz in flash scenes to capture the event streams. Simultaneously, we keep the annotation RGB cameras from overexposure by reducing their exposure time. We also collect 10 fast motion sequences from 4 subjects. For outdoor scenes, we collect 6 DAVIS346 sequences and 6 PROPHESEE sequences for evaluation, including 6 scenes with fast motions (3 by DAVIS346 and 3 by PROPHESEE).




Fig. 6: Visualization of hand pose distribution among EvRealHands, FreiHAND [56], and InterHand2.6M [57] using t-SNE. In our visualization setup, each vector fed into t-SNE is the Euler angles of MANO hand pose parameters with 20 dimension as in [57]. Compared with FreiHAND [56] (green) and InterHand2.6M [57] (blue), EvRealHands (red for fixed poses, yellow for random poses) has more diversity in hand poses.

## 5.2 Annotation

### 5.2.1 3D Joint Annotation

Our dataset is of single hands and we apply 21 keypoints scheme [13] to annotate each hand. Inspired by Interhand2.6M [57] and FreiHand [56], we use multi-view RGB images for 3D annotation and apply a two-stage process consisting of machine annotation and human annotation. First, we use Mediapipe [64] to detect 2D hand keypoints on all the RGB images and triangulate 2D keypoints to obtain 3D keypoints with RANSAC method. Then we manually verify all the keypoints re-projected by 3D keypoints and select the unqualified views for human annotation. Since our annotations have been checked manually, we do not perform a bootstrapping procedure as [59]. Due to the cost of human annotation, manual annotation is only applied to fixed pose and random short sequences and machine annotation is then applied to random long sequences. For fast motion sequences, we can not get accurate 3D joint annotations due to the severe motion blur in the captured images. To quantitatively evaluate our method, we manually annotate the 2D joints on the ECI [36] frames from the event sequences. For outdoor sequences, we only annotate 2D bounding boxes for qualitative evaluation.

### 5.2.2 Shape Annotation

We use MANO shape parameter  $\beta$  to represent the hand shape. We first scan the hands of each subject with an Artec Eva 3D Scanner, and get meshes of static hands. Then we fit a hand MANO model to each mesh to get the prior shape parameter. For each frame of multiple RGB cameras, we fit hand MANO parameters with prior shape as initialization under the constraints from 2D keypoints, 3D joints, and regularization terms similar to FreiHand [56].

## 6 EXPERIMENTS

In this section, we will first introduce the evaluation metrics in Section 6.1. Then we will compare our method with baselines

on different scenes in Section 6.2. Finally, we will perform ablation study to demonstrate the claimed effect of each component we propose in Section 6.3. For implementation details, 120 fps inference results, different methods of CM loss and hand-edge loss, and downstream hand gesture recognition task, please refer to Appendix A in the supplemental material and the video.

### 6.1 Evaluation Metrics

Similar to previous evaluation metrics [10] and [17], we apply root-aligned mean per joint position error (MPJPE, mm), Procrustes Analysis (PA) [65] based MPJPE (PA-MPJPE, mm), and the area under the curve (AUC) of PCK (percentage of correct keypoints) with thresholds ranging from 0 to 10 cm for 3D annotated sequences. For fast motion sequences, we do not have 3D annotations. We compute root-aligned 2D-MPJPE (pixels) on keypoints which have been normalized by the palm length (distance from wrist to MCP joint) as evaluation metric. Since our 2D joint annotations are on the imaging plane of the event camera, the 3D joints predicted by the RGB-based approach are projected onto the imaging plane of the event camera for evaluation.

### 6.2 Comparison with State-of-the-art Methods

#### 6.2.1 Baselines

In order to compare the hand pose estimation performance of event camera and RGB camera under challenging scenes, we select MobRecon [17] as the baseline, because MobRecon [17] can efficiently achieve the state-of-the-art results on popular RGB datasets such as FreiHAND [56], *etc.* Since pixel resolution of the RGB camera is much higher than event camera DAVIS346, the scale of cropped images for MobRecon [17] are much larger than the event representation in EvHandPose. Thus we resize the RGB images to make the hand occupy similar pixels as events for fair comparison. For MobRecon [17], we select the data of one viewpoint from our camera acquisition system and divide the dataset according to the above convention for the sake of fairness. In strong light scenes, a direct approach is to preprocess the overexposed images using HDR techniques and then perform hand pose estimation using RGB-based methods. For comparing this approach, we choose a representative image HDR method, SingleHDR [66], to preprocess overexposed images, followed by hand pose estimation using MobRecon [17]. This approach is designed to compare the performance between the HDR-preprocessing method and the event-based methods. To compare with the most relevant solutions, *i.e.*, event-based hand pose estimation methods, we select EventHands [10] and the method proposed by Jalees *et al.* [45]. EventHands [10] is the first neural event-based 3D hand pose estimation approach, which demonstrates better performance than RGB-based methods on their proposed synthetic dataset. Jalees *et al.*'s method [45] is the model-based approach for hand pose tracking. We use their officially released codes for effective evaluation.

#### 6.2.2 Results

In this section, we compare the state-of-the-art color-based and event-based hand pose methods, *i.e.*, MobRecon [17], EventHands [10], and Jalees *et al.* [45], under different typical hand scenarios including normal scenes, strong light scenes, and fast motion scenes. We show the results of different methods qualitatively and quantitatively in Figure 7 and Table 2, respectively. Quantitative results show that EvHandPose predicts more accurate hand pose



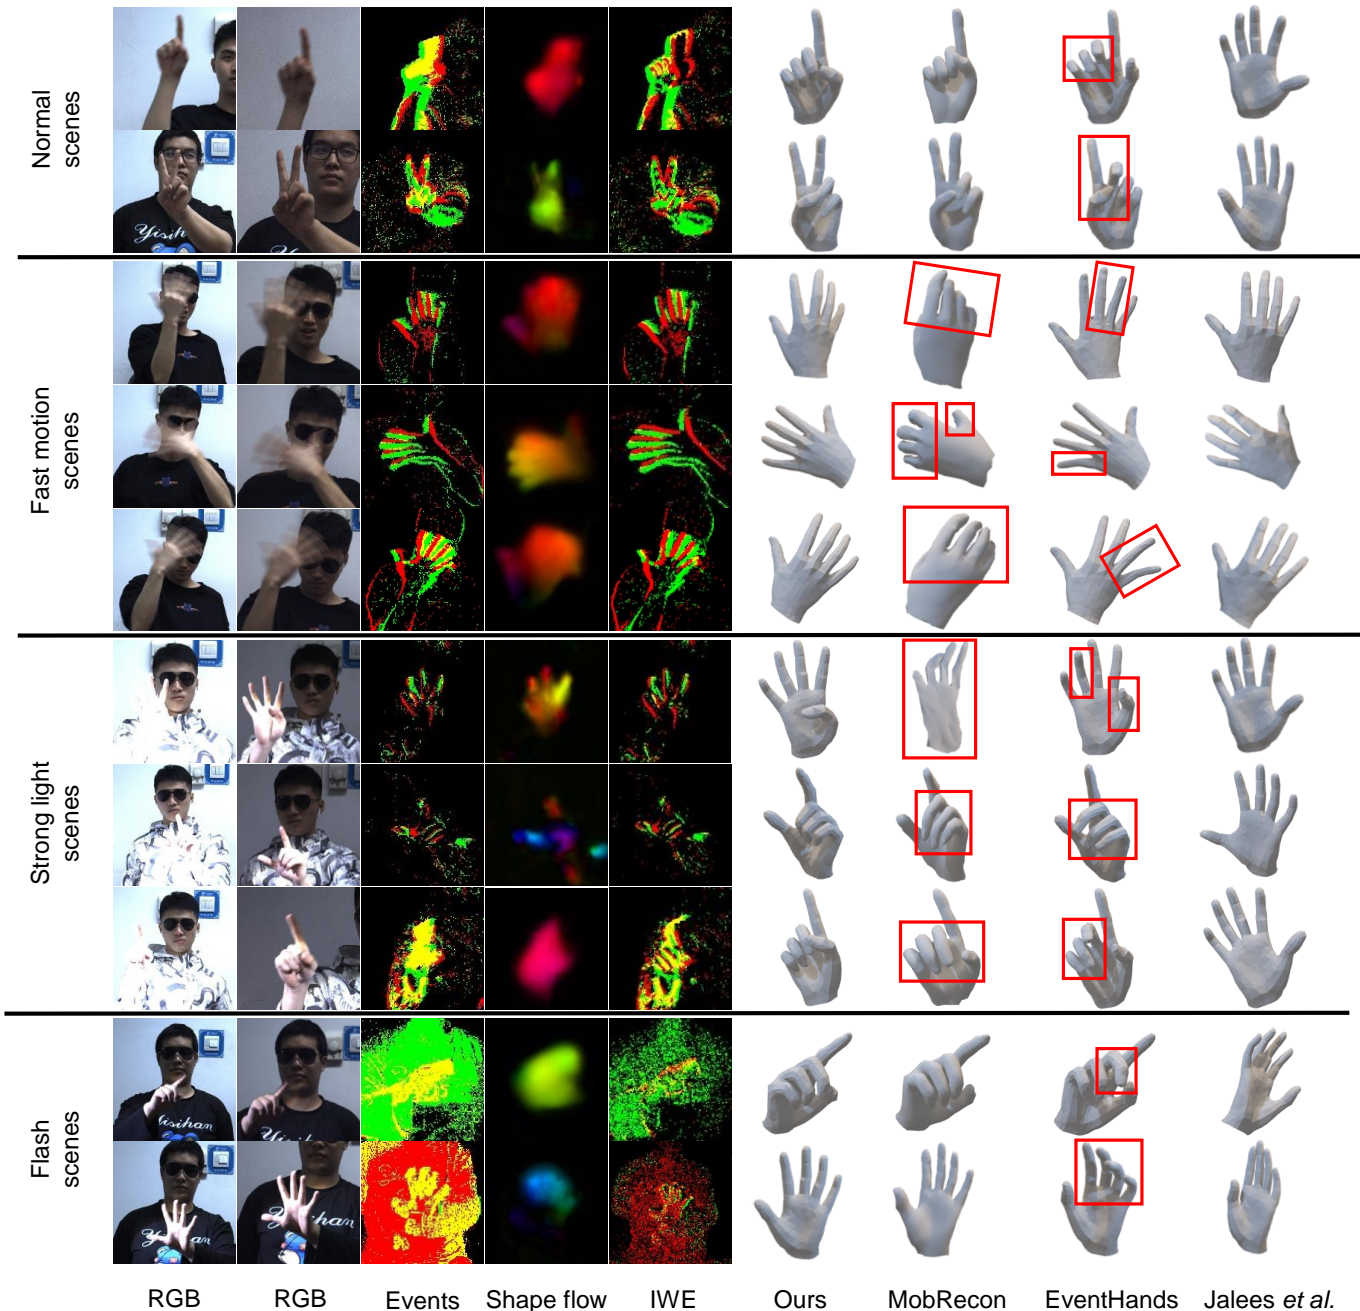


Fig. 7: Qualitative analysis. Columns from left to right: Two views of RGB images, events, predicted shape flows, IWEs, poses from EvHandPose (ours), poses from MobRecon [17], poses from EventHands [10], and poses from the method of Jalees *et al.* [45]. EvHandPose achieves accurate and robust hand pose estimation compared with EventHands [10], Jalees *et al.* [45], and MobRecon [17] as highlighted in red boxes.

than EventHands [10] and Jalees *et al.* [45] by MPJPE over 20 mm lower. This is because our method leverages the temporal and spatial features of events effectively, and addresses the issues of motion ambiguity and sparse annotations. Compared with MobRecon [17], our method achieves better performance in challenging scenarios such as strong light and fast motions.

**Normal Scenes.** In normal scenes, our method outperforms EventHands [10] on MPJPE by 19 ~ 26 mm. As shown in Table 2, MPJPE of our method in the fixed pose sequences reaches 19.82 mm. As stated in Supancic *et al.* [67], human annotators could achieve an accuracy about 20 mm for hand, and can easily

distinguish two poses if the error exceeds 50 mm. Therefore, our method can already achieve sufficient hand pose performance for practical gesture interaction applications, as shown in Appendix A in the supplemental material. The performance of our method is slightly worse (but comparable) than MobRecon [17]. This is reasonable due to the richer texture features of RGB images under normal scenes.

**Fast Motion Scenes.** In fast motion sequences, EvHandPose outperforms EventHands with lower 2D-MPJPE. This is because in training procedure EvHandPose learns the optical flow generated by hand movement on the one hand, and the dynamic time

TABLE 2: Quantitative experiments results of different methods and ablation study results.

Purposes	Methods	Training data	Metrics	Normal		Fast motion	Strong light		Flash	
				Fixed	Random	Random	Fixed	Random	Fixed	Random
Main results	EvHandPose	EvRealHands (events)	3D/2D-MPJPE PA-MPJPE	19.82/4.39 <b>10.19</b>	29.19/6.70 <b>12.72</b>	-8.18 -	<b>28.98/6.80</b> 11.67	<b>33.59/7.64</b> <b>14.59</b>	32.80/7.55 12.85	50.41/10.97 16.92
	EventHands [10]	EvRealHands (events)	3D/2D-MPJPE PA-MPJPE	38.51/10.57 14.18	55.37/14.86 17.42	-/15.04 -	47.97/13.49 17.41	57.43/15.32 17.78	54.30/15.07 17.18	70.86/18.63 21.48
	Jalees <i>et al.</i> [45]	EvRealHands (events)	3D/2D-MPJPE PA-MPJPE	77.50/17.42 27.56	76.15/35.70 17.14	-/33.85 -	73.34/13.13 28.47	80.18/21.46 19.19	75.88/32.09 24.21	89.58/31.92 19.76
	MobRecon [17]	EvRealHands (RGB)	3D/2D-MPJPE PA-MPJPE	<b>19.28/3.82</b> 10.90	<b>24.53/5.07</b> 12.85	-/13.75 -	40.54/11.94 19.51	59.84/16.32 25.38	<b>30.30/7.00</b> 13.22	<b>44.24/10.26</b> 17.66
	SingleHDR [66] + MobRecon [17]	EvRealHands (RGB)	3D/2D-MPJPE PA-MPJPE	24.80/5.04 12.54	30.66/6.24 14.02	-/19.24 -	63.96/22.26 25.38	73.35/22.91 27.61	35.10/8.98 14.90	49.55/12.42 18.52
Ablation: Event representation	EvHandPose (w/o flow)	EvRealHands (events)	3D-MPJPE PA-MPJPE	25.19 10.70	37.36 14.77	11.03 -	40.30 13.15	42.81 17.21	39.90 12.90	60.17 19.78
	ECI [36]	EvRealHands (events)	3D-MPJPE PA-MPJPE	23.13 10.56	33.76 14.03	9.81 -	32.16 12.39	39.80 16.27	36.29 13.01	49.54 17.81
	Voxel [38]	EvRealHands (events)	3D-MPJPE PA-MPJPE	26.41 11.31	35.48 13.99	9.67 -	35.96 13.47	40.14 16.47	33.69 <b>12.32</b>	47.04 18.34
	TORÉ [39]	EvRealHands (events)	3D-MPJPE PA-MPJPE	23.78 10.83	33.57 13.46	9.39 -	32.55 <b>11.43</b>	38.54 15.47	35.90 12.45	51.47 18.18
	Time surface [37]	EvRealHands (events)	3D-MPJPE PA-MPJPE	31.01 12.25	40.32 15.13	9.94 -	38.28 13.14	44.64 16.87	40.65 14.66	58.89 20.31
	EST [40]	EvRealHands (events)	3D-MPJPE PA-MPJPE	24.01 11.14	36.13 14.84	11.83 -	30.44 12.71	39.37 16.68	35.11 12.63	55.61 19.47
	Matrix-LSTM [41]	EvRealHands (events)	3D-MPJPE PA-MPJPE	58.36 27.12	67.01 26.73	34.88 -	59.71 28.52	67.61 26.30	65.00 25.06	78.22 30.37
	Ablation: Weakly-supervision framework	EvHandPose (w/o hand-edge)	EvRealHands (events)	3D-MPJPE PA-MPJPE	20.44 10.42	29.34 12.83	9.00 -	29.13 11.52	34.58 15.07	33.52 12.49
EvHandPose (w/o CM)		EvRealHands (events)	3D-MPJPE PA-MPJPE	20.85 10.43	29.62 12.73	8.77 -	29.90 12.08	34.89 14.88	36.42 13.21	54.40 <b>16.69</b>
EvHandPose (w/o smooth)		EvRealHands (events)	3D-MPJPE PA-MPJPE	20.48 10.24	29.59 12.87	8.76 -	29.16 12.23	33.77 15.17	34.77 12.75	52.57 17.28
EvHandPose (super)		EvRealHands (events)	3D-MPJPE PA-MPJPE	22.76 10.33	32.91 13.54	9.05 -	30.16 11.68	36.26 15.17	37.26 12.68	54.91 17.64
Ablation: Real-synthetic domain gap	EvHandPose	EventHands [10] (events)	3D-MPJPE PA-MPJPE	70.13 22.10	77.19 24.02	32.46 -	79.07 23.15	82.35 23.11	92.20 21.75	99.47 29.31
	EvHandPose	InterHand2.6M [57] (synth. events)	3D-MPJPE PA-MPJPE	51.01 24.02	65.96 18.18	21.13 -	56.88 26.00	68.57 18.35	61.29 21.50	80.23 20.61
	NGA [34]	EvRealHands (events & RGB)	3D-MPJPE PA-MPJPE	39.34 12.68	58.63 14.93	20.60 -	37.56 15.19	47.65 18.15	56.33 16.24	62.34 20.18
	EvTransfer [35]	EvRealHands (events & RGB)	3D-MPJPE PA-MPJPE	32.65 12.90	37.64 15.46	22.42 -	41.65 14.78	43.44 17.86	46.33 15.64	58.90 19.62

window make EvHandPose more adaptable to sub-segments of different intervals. Compared to MobRecon [17], EvHandPose has a 2D-MPJPE reduction of 40.5%. The main reasons for worse performance of MobRecon can be summarized as: Fast motion will lead to severe hand motion blur effect of RGB camera, and it is often infeasible for MobRecon [17] or human annotators to extract accurate hand region and keypoints from blurred images. However, the asynchronous imaging nature of the event camera allows it to capture microsecond motion information and EvHandPose adopts motion representations to record the hand movement precisely as shown in Figure 7 and Figure 8. The edge representation of the hand is still clear in event frames by setting the event time window to very short intervals (*e.g.*, 5 ms). The stable edge and motion representations enable EvHandPose to achieve accurate high fps hand pose estimation, although there are no fast motion data in training.

**Strong Light Scenes.** As shown in Table 2, EvHandPose predicts more stable and accurate hand poses than MobRecon [17] on strong light sequences with MPJPE 11.6 ~ 26.2 mm lower. As shown in Figure 7, under strong light, hand overexposure leads to loss of texture information, and it is challenging to estimate hand pose from overexposed images in RGB-based methods. However,

EvHandPose can still extract edge and hand flow feature from the event stream under strong light, thus achieving robust 3D hand pose estimation. As shown in Figure 7, EvHandPose involves hand motion information in Event-to-Pose module, which makes it more robust and accurate than EventHands [10]. EvHandPose has a 4.4 ~ 9.2 mm increase in MPJPE in strong light scenes compared to normal scenes. One reason is that the distribution of the events generated under strong light is different from that under normal scenes, even though they both record the edge motion. The other reason is that background and shadow of the hand under strong light will also produce a large number of events, making it hard to predict accurate hand poses. The HDR-preprocessing method does not perform better than MobRecon [17] under strong light scenes. While analyzing the HDR-processed images, we find that the images with overexposure in our dataset are difficult to recover using HDR techniques. Furthermore, HDR techniques tend to distort the edges and texture distribution of the hands, which interferes in hand pose estimation.

**Flash Scenes.** As shown in Table 2, EvHandPose outperforms EventHands [10] and Jalees *et al.*'s method [45] with MPJPE over 20 mm lower. Because EvHandPose can capture multiple neighboring hand motions from event streams and utilize the

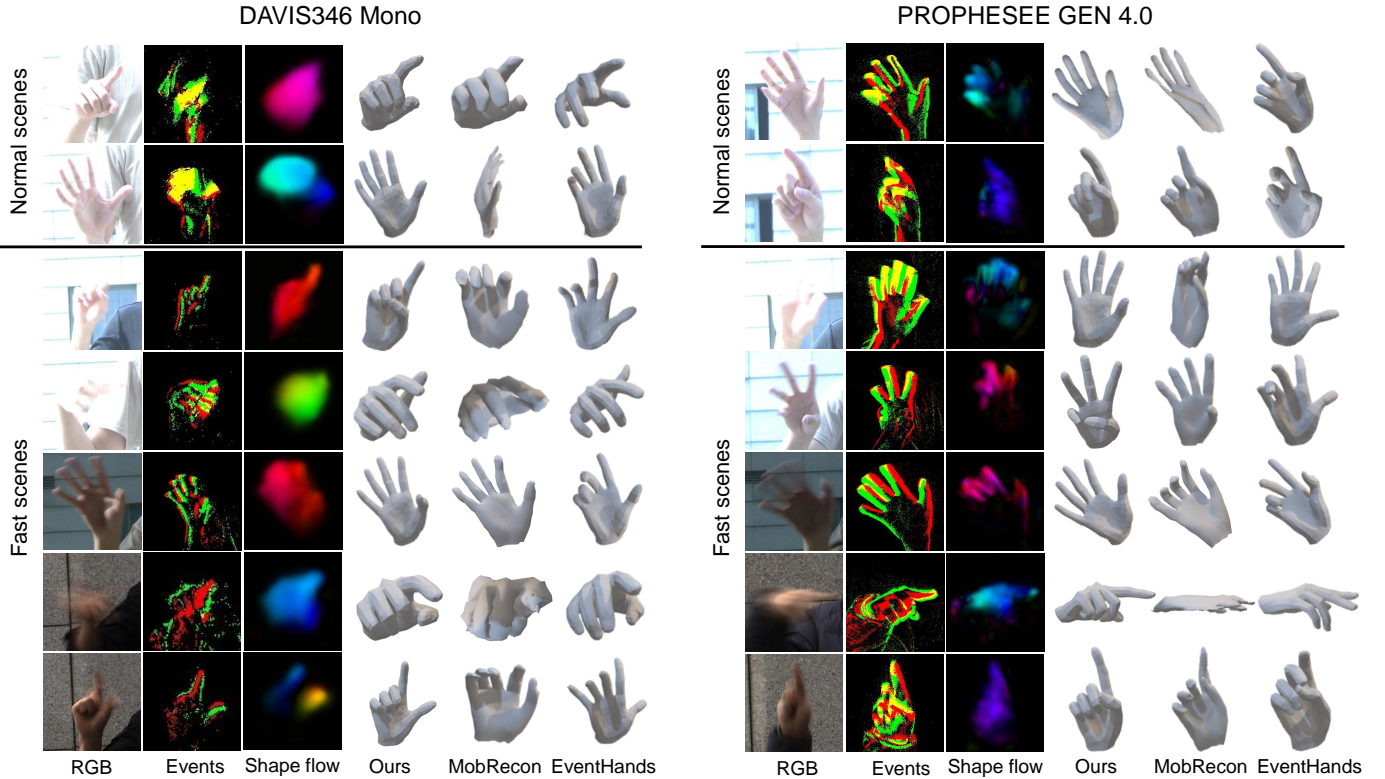


Fig. 8: Qualitative results for evaluating the generalization ability to outdoor scenes and another type of event camera. Left part is for the results on the event camera DAVIS346 Mono, and the right for PROPHESSEE GEN 4.0. Columns from left to right: RGB images, events, predicted shape flows, poses from EvHandPose (ours), poses from MobRecon [17], poses from EventHands [10]. When facing strong illumination or fast hand motion issues in outdoor scenes, EvHandPose outperforms RGB-based method MobRecon [17] and previous event-based method EventHands [10] by a large margin.

temporal information to alleviate the interference of noisy events from the background illumination change. However, the MPJPE of EvHandPose is slightly lower than MobRecon [17] by 2.5 ~ 6.2 mm. This is due to some inherent limitations from the DAVIS346 camera we use, which shows a large amount of events erupting in a short time for the scene with a broad dynamic range. This phenomena causes too few effective event observations in a short period of time, which has little useful information for pose estimation. This might be improved by proposing effective temporal filter for hand pose or using more advanced event cameras.

**Generalization to Outdoor Scenes.** In order to compare the generalization ability to outdoor scenes, we train EvHandPose and the baselines in indoor scenes and evaluate them in outdoor scenes. As shown in Figure 8, EvHandPose predicts more stable hand poses from event streams in outdoor scenes compared with the event-based method [10], and MobRecon [17] which faces overexposure and motion blur issues. The strong generalization performance in outdoor scenes benefits from the effectiveness of our motion representation to retain hand motions and the weakly-supervision framework to learn diverse hand poses from event streams.

**Generalization to Another Type of Event Camera.** In order to demonstrate the generalization ability of EvHandPose to the another type of event camera, we also qualitatively evaluate those methods on outdoor PROPHESSEE sequences, which are trained on indoor DAVIS346 sequences. As shown in Figure 8 on the sequences from the outdoor PROPHESSEE camera, EvHandPose significantly outperforms MobRecon [17] than those in DAVIS346 outdoor

TABLE 3: Comparison of computational cost, latency, and robustness of different methods.

Method	Computation ↓		Latency ↓		MPJPE ↓	
	Params (M)	FLOPs (G)	Imaging (ms)	Inference (ms)	Strong light (mm)	Fast motion (pixel)
MobRecon [17]	<b>10.73</b>	<b>0.42</b>	15	5.15	53.01	13.75
Jalees <i>et al.</i> [45]	-	-	<1	212.26	77.67	33.85
EventHands [10]	20.33	6.47	<1	3.34	53.96	15.04
EvHandPose	47.36	3.44	<1	<b>1.55</b>	<b>31.91</b>	<b>8.18</b>

scenes, mainly due to the higher pixel resolution (1280×720 pixels) thus richer events information of PROPHESSEE. Besides, our motion representation also enables the capability of cross-camera generalization.

**Computational Cost, Latency, and Accuracy.** We conduct a further comparison about computational cost, latency, and accuracy among these methods in Table 3. Computational cost is measured by model parameters (Params) and floating point operations per inference (FLOPs). Latency is mainly composed of data imaging process time (Imaging) and model inference time (Inference). Since Jalees *et al.* [45]’s method is an optimization-based methods, the number of iteration steps strongly influences the computational cost and latency. We try our best to find the results with best accuracy and least iteration steps (1000 iterations in their official codes). As shown in Table 3, EvHandPose can achieve 40% lower MPJPE in strong light and fast motion scenes with the lowest latency, thus achieving robust estimation with high temporal resolution.

This derives from low latency and robustness of the asynchronous imaging mechanism of event cameras compared with frame-based cameras. Also, the simple yet effective motion representations of EvHandPose contribute to the superior performance, too. Additionally, our backbone is based on the vanilla ResNet [48], and we anticipate future modification on the model structure similar to MobRecon [17] to reduce computational cost and latency.

### 6.3 Ablation Study

We conduct ablation studies on each contributed module of our method. Specially, we evaluate the effectiveness of the predicted hand flow as event representations for hand pose estimation, the effect of weakly-supervision framework for sparse annotation issue, and the effect of EvRealHands to bridge the real-synthetic domain gap in event-based 3D hand pose estimation.

**Effect of Hand Flow as Event Representation.** To demonstrate the effectiveness of our hand motion representation, we compare our representation with other widely adopted event representations, including ECI [36], time surface [37], voxel [38], TORE [39], EST [40], and Matrix-LSTM [41]. For fair comparison, these methods share the same ResNet34 [48] backbone with our method. Quantitative results in Table 2 show that our hand flow representation outperforms them with MPJPE 4 ~ 28 mm lower and PA-MPJPE 1 ~ 17 mm lower. Removing hand flow representation from EvHandPose (denoted as “w/o flow” in Table 2) results in a 4 ~ 11 mm increase in MPJPE. As qualitative results shown in Figure 7, EvHandPose predicts hand flow from events in its Event-to-Pose module. On the one hand, shape flow can alleviate motion ambiguity issue because it contains the temporal information of hand movements as shown in the first row of the results under strong light in Figure 7. On the other hand, hand shape flow can separate the hand from the background, thus improving the accuracy of estimation.

#### Effect of Weakly-supervision Framework for Sparse Annotation.

We evaluate the influence of the weakly-supervision framework and each loss item on the performance of EvHandPose. In Table 2, EvHandPose, EvHandPose (w/o hand-edge), EvHandPose (w/o CM), EvHandPose (w/o smooth), and EvHandPose (super) respectively mean the methods with complete loss, without hand-edge loss, without CM loss, without smooth loss, and supervised method without any weakly-supervision loss. As shown in Table 2, compared with the EvHandPose under supervision, our weakly-supervision framework lead to a 1.2 ~ 4.5 mm decrease in 3D-MPJPE in each scene. According to the experiments, all losses contribute to the performance on the normal, fast motion, and flash scenes. This improvement derives from the effective usage of dense event streams under sparse 3D annotation.

#### Effect of EvRealHands for Real-Synthetic Domain Gap.

To demonstrate how we solve the real-synthetic domain gap in event-based hand pose estimation, we train EvHandPose on EventHands [10] or synthetic data from the existing RGB-based dataset Interhand2.6M [57] and evaluate the models on our real sequences. The synthetic data from Interhand2.6M [57] are generated by the event simulator v2e [33] with single hand RGB sequences of 7 subjects. The results in Table 2 show that the issues in previous synthetic dataset [10] (unnatural hand poses and limitations of event simulators) result in a 20 ~ 50 mm increase in MPJPE for EvHandPose. Although synthetic Interhand2.6M are composed of natural hand poses, there is still a 25 ~ 30 mm MPJPE gap. We select NGA [34] and EvTransfer [35] as representative domain

adaptation methods and use the same ResNet34 [48] backbone for comparison. Even though we train them using our EvRealHands dataset, NGA [34] and EvTransfer [35] still underperform our method by 12 ~ 23 mm MPJPE lower. This performance gap shows that collecting a large-scale real-world event-based hand dataset is beneficial for event-based 3D hand pose estimation.

### 6.4 Limitations

Although EvHandPose shows promising hand pose results even in several challenging scenarios, we find that there are still several cases in which our method fails to predict correct hand poses. Typical cases showing large hand pose errors can be summarized as follows:

- Motion ambiguity: Although motion information and recurrent model are adopted to alleviate the motion ambiguity issue, it is not completely removed. It is still challenging to predict hand pose from event streams if a hand is kept stationary for a long time span;
- New pose: We observe that MPJPE of our method is 5 ~ 15 mm smaller in the fixed poses than in the random poses. One possible reason is that the distribution of random poses has greater diversity, leading to a slightly performance drop in random pose sequences. This effect can be reduced by including more data with diverse hand poses;

## 7 CONCLUSION

We propose a novel neural network for event-based monocular 3D hand pose estimation on sparse-labeled event streams. In order to deal with the asynchronous data format and motion ambiguity issue, we adopt edge and hand flow representations to utilize the spatial-temporal information of event streams effectively. To tackle the sparse annotation challenge, we design contrast maximization and edge constraints in a weakly-supervision framework. And we construct the first large-scale real-world dataset for event-based hand pose estimation to address the domain gap between the synthetic and real data. Experiments on our dataset demonstrate that our method outperforms the recent event-based hand pose method [10] in all cases of EvRealHands testing dataset and RGB-based method [17] in challenging scenarios, and can generalize well to outdoor scenes and a different type of event camera. Although our model has not been trained on fast motion sequences, it can achieve stable hand pose estimation at 120 fps. Our dataset and method provides a new benchmark and baseline of robust event-based 3D hand pose estimation, and they might inspire related researches in the future, such as hand gesture recognition directly from event streams. We also hope to improve EvRealHands by capturing hands using more advanced event cameras (such as [68] and [69]) with higher resolution and less noise.

## ACKNOWLEDGMENTS

This work is supported by National Science and Technology Major Project (Grant No. 2021ZD0109803), Beijing Natural Science Foundation (Grant No. L232028 and L233024), and National Natural Science Foundation of China (Grand No. 62136001, 62088102), and Beijing Hospitals Authority Clinical Medicine Development of Special Funding Support No. ZLRK202330. The authors would like to thank Yiran Zhang, Qida Hao, and Bingxuan Wang for their help to collect and annotate the EvRealHands dataset, Jian Cheng for his help in the downstream hand gesture recognition task, Xinyu Zhou for his help to collect outdoor datasets, and Yixin Yang for her help in the HDR-preprocessing method.

## REFERENCES

- [1] V. Lepetit, "Recent advances in 3d object and hand pose estimation," *arXiv preprint arXiv:2006.05927*, 2020.
- [2] X. Deng, D. Zuo, Y. Zhang, Z. Cui, J. Cheng, P. Tan, L. Chang, M. Pollefeys, S. Fanello, and H. Wang, "Recurrent 3d hand pose estimation using cascaded pose-guided 3d alignments," *IEEE Transactions on Pattern Analysis and Machine Intelligence (PAMI)*, vol. 45, no. 1, pp. 932–945, 2022.
- [3] P. Lichtsteiner, C. Posch, and T. Delbrück, "A 128×128 120 dB 15  $\mu$ s latency asynchronous temporal contrast vision sensor," *IEEE J. Solid State Circuits*, pp. 566–576, 2008.
- [4] A. Mitrokhin, C. Fermüller, C. Parameshwara, and Y. Aloimonos, "Event-based moving object detection and tracking," in *2018 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 2018.
- [5] D. Gehrig, H. Rebecq, G. Gallego, and D. Scaramuzza, "Asynchronous, photometric feature tracking using events and frames," in *Proc. of European Conference on Computer Vision (ECCV)*, 2018.
- [6] H. Rebecq, G. Gallego, E. Mueggler, and D. Scaramuzza, "EMVS: Event-based multi-view stereo—3D reconstruction with an event camera in real-time," *International Journal of Computer Vision (IJCV)*, vol. 126, no. 12, pp. 1394–1414, 2018.
- [7] J. Hagenaaers, F. Paredes-Vallés, and G. De Croon, "Self-supervised learning of event-based optical flow with spiking neural networks," *Advances in Neural Information Processing Systems (NeurIPS)*, 2021.
- [8] H. Rebecq, T. Horstschäfer, G. Gallego, and D. Scaramuzza, "Evo: A geometric approach to event-based 6-Dof parallel tracking and mapping in real time," *IEEE Robotics and Automation Letters (IRAL)*, 2016.
- [9] Y. Zhou, G. Gallego, and S. Shen, "Event-based stereo visual odometry," *IEEE Transactions on Robotics (TRO)*, vol. 37, no. 5, pp. 1433–1450, 2021.
- [10] V. Rudnev, V. Golyanik, J. Wang, H.-P. Seidel, F. Mueller, M. Elgharib, and C. Theobalt, "EventHands: Real-time neural 3D hand pose estimation from an event stream," in *Proc. of International Conference on Computer Vision (ICCV)*, 2021.
- [11] H. Rebecq, R. Ranftl, V. Koltun, and D. Scaramuzza, "High speed and high dynamic range video with an event camera," *IEEE Transactions on Pattern Analysis and Machine Intelligence (PAMI)*, vol. 43, no. 6, pp. 1964–1980, 2019.
- [12] M. Mostafavi, L. Wang, and K.-J. Yoon, "Learning to reconstruct HDR images from events, with applications to depth and flow prediction," *International Journal of Computer Vision (IJCV)*, vol. 129, no. 4, pp. 900–920, 2021.
- [13] C. Zimmermann and T. Brox, "Learning to estimate 3D hand pose from single RGB images," in *Proc. of International Conference on Computer Vision (ICCV)*, 2017.
- [14] U. Iqbal, P. Molchanov, T. B. J. Gall, and J. Kautz, "Hand pose estimation via latent 2.5 D heatmap regression," in *Proc. of European Conference on Computer Vision (ECCV)*, 2018.
- [15] L. Ge, Z. Ren, Y. Li, Z. Xue, Y. Wang, J. Cai, and J. Yuan, "3D hand shape and pose estimation from a single RGB image," in *Proc. of Conference on Computer Vision and Pattern Recognition (CVPR)*, 2019.
- [16] X. Zhang, Q. Li, H. Mo, W. Zhang, and W. Zheng, "End-to-end hand mesh recovery from a monocular RGB image," in *Proc. of International Conference on Computer Vision (ICCV)*, 2019.
- [17] X. Chen, Y. Liu, Y. Dong, X. Zhang, C. Ma, Y. Xiong, Y. Zhang, and X. Guo, "MobRecon: Mobile-friendly hand mesh reconstruction from monocular image," *Proc. of Conference on Computer Vision and Pattern Recognition (CVPR)*, 2022.
- [18] I. Oikonomidis, M. I. Lourakis, and A. A. Argyros, "Evolutionary quasi-random search for hand articulations tracking," in *Proc. of Conference on Computer Vision and Pattern Recognition (CVPR)*, 2014.
- [19] D. Tzionas, L. Ballan, A. Srikantha, P. Aponte, M. Pollefeys, and J. Gall, "Capturing hands in action using discriminative salient points and physics simulation," *International Journal of Computer Vision (IJCV)*, vol. 118, no. 2, pp. 172–193, 2016.
- [20] I. Oikonomidis, N. Kyriazis, and A. A. Argyros, "Tracking the articulated motion of two strongly interacting hands," in *Proc. of Conference on Computer Vision and Pattern Recognition (CVPR)*, 2012.
- [21] L. Ballan, A. Taneja, J. Gall, L. V. Gool, and M. Pollefeys, "Motion capture of hands in action using discriminative salient points," in *Proc. of European Conference on Computer Vision (ECCV)*, 2012.
- [22] N. Kyriazis and A. Argyros, "Scalable 3D tracking of multiple interacting objects," in *Proc. of Conference on Computer Vision and Pattern Recognition (CVPR)*, 2014.
- [23] J. Romero, D. Tzionas, and M. J. Black, "Embodied Hands: Modeling and capturing hands and bodies together," *ACM Transactions on Graphics (Proc. of ACM SIGGRAPH)*, vol. 36, no. 6, 2017.
- [24] Y. Cai, L. Ge, J. Cai, and J. Yuan, "Weakly-supervised 3D hand pose estimation from monocular RGB images," in *Proc. of European Conference on Computer Vision (ECCV)*, 2018.
- [25] A. Spurr, U. Iqbal, P. Molchanov, O. Hilliges, and J. Kautz, "Weakly supervised 3D hand pose estimation via biomechanical constraints," in *Proc. of European Conference on Computer Vision (ECCV)*, 2020.
- [26] X. Deng, Y. Zhu, Y. Zhang, Z. Cui, P. Tan, W. Qu, C. Ma, and H. Wang, "Weakly supervised learning for single depth-based hand shape recovery," *IEEE Transactions on Image Processing (TIP)*, vol. 30, pp. 532–545, 2020.
- [27] D. Kulon, R. A. Guler, I. Kokkinos, M. M. Bronstein, and S. Zafeiriou, "Weakly-supervised mesh-convolutional hand reconstruction in the wild," in *Proc. of Conference on Computer Vision and Pattern Recognition (CVPR)*, 2020.
- [28] Y. Chen, Z. Tu, D. Kang, L. Bao, Y. Zhang, X. Zhe, R. Chen, and J. Yuan, "Model-based 3D hand reconstruction via self-supervised learning," in *Proc. of Conference on Computer Vision and Pattern Recognition (CVPR)*, 2021.
- [29] C. Wan, T. Probst, L. V. Gool, and A. Yao, "Self-supervised 3D hand pose estimation through training by fitting," in *Proc. of Conference on Computer Vision and Pattern Recognition (CVPR)*, 2019.
- [30] A. Boukhayma, R. d. Bem, and P. H. Torr, "3D hand shape and pose from images in the wild," in *Proc. of Conference on Computer Vision and Pattern Recognition (CVPR)*, 2019.
- [31] N. Neverova, J. Thewlis, R. A. Guler, I. Kokkinos, and A. Vedaldi, "Slim densepose: Thrifty learning from sparse annotations and motion cues," in *Proc. of Conference on Computer Vision and Pattern Recognition (CVPR)*, 2019.
- [32] Y. Hasson, B. Tekin, F. Bogo, I. Laptev, M. Pollefeys, and C. Schmid, "Leveraging photometric consistency over time for sparsely supervised hand-object reconstruction," in *Proc. of Conference on Computer Vision and Pattern Recognition (CVPR)*, 2020.
- [33] D. Gehrig, M. Gehrig, J. Hidalgo-Carrió, and D. Scaramuzza, "Video to events: Recycling video datasets for event cameras," in *Proc. of Conference on Computer Vision and Pattern Recognition (CVPR)*, 2020.
- [34] Y. Hu, T. Delbrück, and S. Liu, "Learning to exploit multiple vision modalities by using grafted networks," in *Proc. of European Conference on Computer Vision (ECCV)*, 2020.
- [35] N. Messikommer, D. Gehrig, M. Gehrig, and D. Scaramuzza, "Bridging the gap between events and frames through unsupervised domain adaptation," *IEEE Robotics and Automation Letters (IRAL)*, 2022.
- [36] A. I. Maqueda, A. Loquercio, G. Gallego, N. García, and D. Scaramuzza, "Event-based vision meets deep learning on steering prediction for self-driving cars," in *Proc. of Conference on Computer Vision and Pattern Recognition (CVPR)*, 2018.
- [37] X. Lagorce, G. Orchard, F. Galluppi, B. E. Shi, and R. B. Benosman, "Hots: a hierarchy of event-based time-surfaces for pattern recognition," *IEEE Transactions on Pattern Analysis and Machine Intelligence (PAMI)*, vol. 39, no. 7, pp. 1346–1359, 2016.
- [38] A. Z. Zhu, L. Yuan, K. Chaney, and K. Daniilidis, "Unsupervised event-based learning of optical flow, depth, and egomotion," in *Proc. of Conference on Computer Vision and Pattern Recognition (CVPR)*, 2019.
- [39] R. W. Baldwin, R. Liu, M. Almatrafi, V. K. Asari, and K. Hirakawa, "Time-ordered recent event (TORE) volumes for event cameras," *IEEE Transactions on Pattern Analysis and Machine Intelligence (PAMI)*, 2023.
- [40] D. Gehrig, A. Loquercio, K. G. Derpanis, and D. Scaramuzza, "End-to-end learning of representations for asynchronous event-based data," in *Proc. of International Conference on Computer Vision (ICCV)*, 2019.
- [41] M. Cannici, M. Ciccone, A. Romanoni, and M. Matteucci, "A differentiable recurrent surface for asynchronous event-based data," in *Proc. of European Conference on Computer Vision (ECCV)*, A. Vedaldi, H. Bischof, T. Brox, and J. Frahm, Eds., 2020.
- [42] A. Z. Zhu, L. Yuan, K. Chaney, and K. Daniilidis, "EV-FlowNet: Self-supervised optical flow estimation for event-based cameras," 2018.
- [43] L. Xu, W. Xu, V. Golyanik, M. Habermann, L. Fang, and C. Theobalt, "EventCap: Monocular 3D capture of high-speed human motions using an event camera," in *Proc. of Conference on Computer Vision and Pattern Recognition (CVPR)*, 2020.
- [44] S. Zou, C. Guo, X. Zuo, S. Wang, P. Wang, X. Hu, S. Chen, M. Gong, and L. Cheng, "EventHPE: Event-based 3D human pose and shape estimation," in *Proc. of International Conference on Computer Vision (ICCV)*, 2021.
- [45] J. Nehvi, V. Golyanik, F. Mueller, H.-P. Seidel, M. Elgharib, and C. Theobalt, "Differentiable event stream simulator for non-rigid 3D tracking," in *Proc. of Conference on Computer Vision and Pattern Recognition (CVPR)*, 2021.

- [46] J. P. Lewis, M. Cordner, and N. Fong, "Pose space deformation: a unified approach to shape interpolation and skeleton-driven deformation," in *ACM Transactions on Graphics (Proc. of ACM SIGGRAPH)*, 2000, pp. 165–172.
- [47] X. Shi, Z. Chen, H. Wang, D.-Y. Yeung, W.-K. Wong, and W.-c. Woo, "Convolutional LSTM network: A machine learning approach for precipitation nowcasting," *Advances in Neural Information Processing Systems (NeurIPS)*, vol. 28, 2015.
- [48] K. He, X. Zhang, S. Ren, and J. Sun, "Deep residual learning for image recognition," in *Proc. of Conference on Computer Vision and Pattern Recognition (CVPR)*, 2016.
- [49] K. Shoemake, "Animating rotation with quaternion curves," in *ACM Transactions on Graphics (Proc. of ACM SIGGRAPH)*, 1985.
- [50] G. Gallego, H. Rebecq, and D. Scaramuzza, "A unifying contrast maximization framework for event cameras, with applications to motion, depth, and optical flow estimation," in *Proc. of Conference on Computer Vision and Pattern Recognition (CVPR)*, 2018.
- [51] B. D. Lucas and T. Kanade, "An iterative image registration technique with an application to stereo vision," in *International Joint Conference on Artificial Intelligence (IJCAI)*, 1981.
- [52] G. Gallego, M. Gehrig, and D. Scaramuzza, "Focus is all you need: Loss functions for event-based vision," in *Proc. of Conference on Computer Vision and Pattern Recognition (CVPR)*, 2019.
- [53] J. Tompson, M. Stein, Y. LeCun, and K. Perlin, "Real-time continuous pose recovery of human hands using convolutional networks," *ACM Transactions on Graphics*, 2014.
- [54] J. Zhang, J. Jiao, M. Chen, L. Qu, X. Xu, and Q. Yang, "3d hand pose tracking and estimation using stereo matching," *CoRR*, 2016. [Online]. Available: <http://arxiv.org/abs/1610.07214>
- [55] S. Yuan, Q. Ye, B. Stenger, S. Jain, and T. Kim, "Bighand2.2m benchmark: Hand pose dataset and state of the art analysis," in *Proc. of Conference on Computer Vision and Pattern Recognition (CVPR)*, 2017.
- [56] C. Zimmermann, D. Ceylan, J. Yang, B. Russell, M. Argus, and T. Brox, "Freihand: A dataset for markerless capture of hand pose and shape from single RGB images," in *Proc. of International Conference on Computer Vision (ICCV)*, 2019.
- [57] G. Moon, S.-I. Yu, H. Wen, T. Shiratori, and K. M. Lee, "Interhand2. 6M: A dataset and baseline for 3D interacting hand pose estimation from a single RGB image," in *Proc. of European Conference on Computer Vision (ECCV)*, 2020.
- [58] S. Brahmabhatt, C. Tang, C. D. Twigg, C. C. Kemp, and J. Hays, "ContactPose: A dataset of grasps with object contact and hand pose," in *Proc. of European Conference on Computer Vision (ECCV)*, 2020.
- [59] T. Simon, H. Joo, I. Matthews, and Y. Sheikh, "Hand keypoint detection in single images using multiview bootstrapping," in *Proc. of Conference on Computer Vision and Pattern Recognition (CVPR)*, 2017.
- [60] S. Han, B. Liu, R. Cabezas, C. D. Twigg, P. Zhang, J. Petkau, T.-H. Yu, C.-J. Tai, M. Akbay, Z. Wang *et al.*, "MEgATrack: monochrome egocentric articulated hand-tracking for virtual reality," *ACM Transactions on Graphics (Proc. of ACM SIGGRAPH)*, vol. 39, no. 4, pp. 87–1, 2020.
- [61] J. Heikkilä and O. Silvén, "A four-step camera calibration procedure with implicit image correction," in *Proc. of Conference on Computer Vision and Pattern Recognition (CVPR)*, 1997.
- [62] P. Duan, Z. W. Wang, X. Zhou, Y. Ma, and B. Shi, "Eventzoom: Learning to denoise and super resolve neuromorphic events," in *Proc. of Conference on Computer Vision and Pattern Recognition (CVPR)*, 2021.
- [63] Q. De Smedt, H. Wannous, J.-P. Vandeboorde, J. Guerry, B. L. Saux, and D. Filliat, "3D hand gesture recognition using a depth and skeletal dataset: Shrec'17 track," in *Proceedings of the Workshop on 3D Object Retrieval*, 2017.
- [64] F. Zhang, V. Bazarevsky, A. Vakunov, A. Tkachenka, G. Sung, C.-L. Chang, and M. Grundmann, "Mediapipe hands: On-device real-time hand tracking," *arXiv preprint arXiv:2006.10214*, 2020.
- [65] J. C. Gower, "Generalized procrustes analysis," *Psychometrika*, 1975.
- [66] Y.-L. Liu, W.-S. Lai, Y.-S. Chen, Y.-L. Kao, M.-H. Yang, Y.-Y. Chuang, and J.-B. Huang, "Single-image HDR reconstruction by learning to reverse the camera pipeline," in *Proc. of Conference on Computer Vision and Pattern Recognition (CVPR)*, 2020, pp. 1651–1660.
- [67] J. S. Supancic III, G. Rogez, Y. Yang, J. Shotton, and D. Ramanan, "Depth-based hand pose estimation: methods, data, and challenges," *Proc. of International Conference on Computer Vision (ICCV)*, 2015.
- [68] S. Chen and M. Guo, "Live demonstration: CeleX-V: A 1M pixel multi-mode event-based sensor," in *Proc. of Conference on Computer Vision and Pattern Recognition (CVPR) Workshops*, 2019, pp. 1682–1683.
- [69] T. Finateu, A. Niwa, D. Matolin, K. Tsuchimoto, A. Mascheroni, E. Reynaud, P. Mostafalu, F. T. Brady, L. Chotard, F. LeGoff, H. Takahashi, H. Wakabayashi, Y. Oike, and C. Posch, "A 1280×720 back-illuminated

stacked temporal contrast event-based vision sensor with 4.86 $\mu$ m pixels, 1.066GEPS readout, programmable event-rate controller and compressive data-formatting pipeline," in *IEEE International Solid-State Circuits Conference*, 2020, pp. 112–114.



**Jianping Jiang** received his B.Eng. degree from Tsinghua University in 2020 and M.Sci. degree from the School of Computer Science, Peking University. His research interests span mixed reality, 3D avatar.



**Jiahe Li** received his B.S. degree in School of Artificial Intelligence, Beijing Normal University. He is currently a master student at the Institute of Software, Chinese Academy. His research interests include computer vision, 3D reconstruction, human motion tracking and synthesis.



**Baowen Zhang** received his B.Eng. degree from Southeast University, in 2020. He is currently a master graduate student at the Institute of Software, Chinese Academy of Sciences. His main research interest is computer vision.



synthesis, and natural user interfaces.

**Xiaoming Deng** is currently a Professor with the Institute of Software, Chinese Academy of Sciences (CAS). He received the BS and MS degrees from Wuhan University, and the PhD degree from the Institute of Automation, CAS. He has been a Research Fellow at the National University of Singapore, and a Postdoctoral Fellow at the Institute of Computing Technology, CAS, respectively. His main research topics are in computer vision, and specifically related to 3D reconstruction, human motion tracking and



Boxin Shi received the BE degree from the Beijing University of Posts and Telecommunications, the ME degree from Peking University, and the PhD degree from the University of Tokyo, in 2007, 2010, and 2013. He is currently a Boya Young Fellow Associate Professor (with tenure) and Research Professor at Peking University, where he leads the Camera Intelligence Lab. Before joining PKU, he did research with MIT Media Lab, Singapore University of Technology and Design, Nanyang Technological University, National Institute of Advanced Industrial Science and Technology, from 2013 to 2017. His papers were awarded as Best Paper Runner-Up at ICCP 2015 and selected as Best Papers from ICCV 2015 for IJCV Special Issue. He is an associate editor of TPAMI/IJCV and an area chair of CVPR/ICCV/ECCV. He is a senior member of IEEE.

# Supplemental Material: EvHandPose: Event-based 3D Hand Pose Estimation with Sparse Supervision

Jianping Jiang<sup>†</sup>, Jiahe Li<sup>†</sup>, Baowen Zhang,  
Xiaoming Deng<sup>‡</sup>, *Member, IEEE*, Boxin Shi<sup>‡</sup>, *Senior Member, IEEE*

## APPENDIX A ADDITIONAL EXPERIMENTAL RESULTS

### A.1 Downstream Application: Gesture Recognition

To demonstrate the potential of event-based hand pose estimation on downstream tasks, we perform skeleton-based hand gesture recognition [1] on fixed hand pose sequences (15 types similar to [2]) in EvRealHands. We train STA-GCN [1], a state-of-the-art hand gesture recognition model built on graph convolution network, on ground truth poses and evaluate it on predicted hand poses from EvHandPose and other baselines. The training and testing splitting is the same as EvHandPose. As results shown in Table S1, results from EvHandPose can achieve 20% accuracy over those from EventHands [3] and Jalees *et al.* [4] in strong light and flash scenes. Using predicted hand poses from EvHandPose can improve the gesture recognition accuracy in strong light scenes by 13% compared with the RGB-based method MobRecon [5].

### A.2 Implementation Details

Since we use temporal information through recurrent model, we need to divide an event stream into  $N$  sub-segments. In our experiment,  $N$  is 12, and we utilize the outputs of the last three sub-segments for supervision. For supervision, the interval of each sub-segment is 66.66 ms, which is the same with the period of annotations. For weakly-supervision, we only let second last sub-segment have 3D annotation, so that we can dynamically choose the interval of the sub-segment and the model can learn more hand information with different intervals. For the sub-segment with annotation, we use its ground truth hand pose instead of predicted hand pose in Pose-to-IWE module. The interval is randomly selected from 10 ms to 100 ms for weakly-supervision. Certainly, all the sub-segments have the same duration in a sequence. For

TABLE S1: Accuracy (%) of hand gesture recognition.

Input results	Normal	Strong light	Flash
EventHands [3]	68.3	60.0	33.3
Jalees <i>et al.</i> [4]	75.0	70.0	53.3
MobRecon [5]	<b>88.3</b>	76.7	<b>76.7</b>
EvHandPose	83.3	<b>90.0</b>	73.3

evaluation, we set the interval of each sub-segment to 5 ms for fast motion sequences and 66.66 ms for other sequences.

In order to make the network effectively extract the feature of the hand, we calculate a bounding box for each sub-segment and then crop out the events inside this bounding box. Obtaining a bounding box for a event sub-segment is different from the RGB image due to the different temporal distributions. Given 3D joints at some timestamp, we project them onto the image plane to obtain the 2D joints and cover the 2D joints with an exact rectangle. We set the square with the center of the rectangle and 1.5 times the length of the longest side of the rectangle as bounding box. To compute a bounding box for a event sub-segment, we first approximately estimate the 3D joints of the start time and end time of this sub-segment with quadratic interpolation. Then we calculate the bounding boxes of the 3D joints and apply a bounding box that covers the two bounding boxes exactly as the final bounding box. For fast motion sequences, our machine annotated 3D joints are not very precise. We apply annotated 2D joints interpolation for generating bounding boxes.

We apply data augmentation including scale, rotation, and translation on events and resize the cropped event frames to  $128 \times 128$ . We use ResNet34 [6] as the backbone of feature encoders, and apply Adam [7] as the optimizer for all our training procedures. The learning rate of FlowNet and EvHandPose in supervision is 0.0005 and for EvHandPose in weakly-supervision is 0.0001. We train EvHandPose on two TITAN RTX 4090 GPUs with batch size 32 for 44 K iterations.

### A.3 Training and Evaluation Data

We collect sequences from 7 out of 10 subjects as training data, 1 as validation data, and 2 as evaluation data. The training data included about 4 minutes of strong light sequences, 2 minutes

- J. Jiang, B. Shi are with National Key Laboratory for Multimedia Information Processing, National Engineering Research Center of Visual Technology, and AI Innovation Center, School of Computer Science, Peking University, Beijing 100871, China.
- J. Li, B. Zhang, X. Deng are with Beijing Key Laboratory of Human Computer Interactions, Institute of Software, Chinese Academy of Sciences, Beijing 100190, China, and University of Chinese Academy of Sciences, Beijing, China.
- <sup>†</sup> Equal contribution. <sup>‡</sup> Corresponding authors: shiboxin@pku.edu.cn, xiaoming@iscas.ac.cn

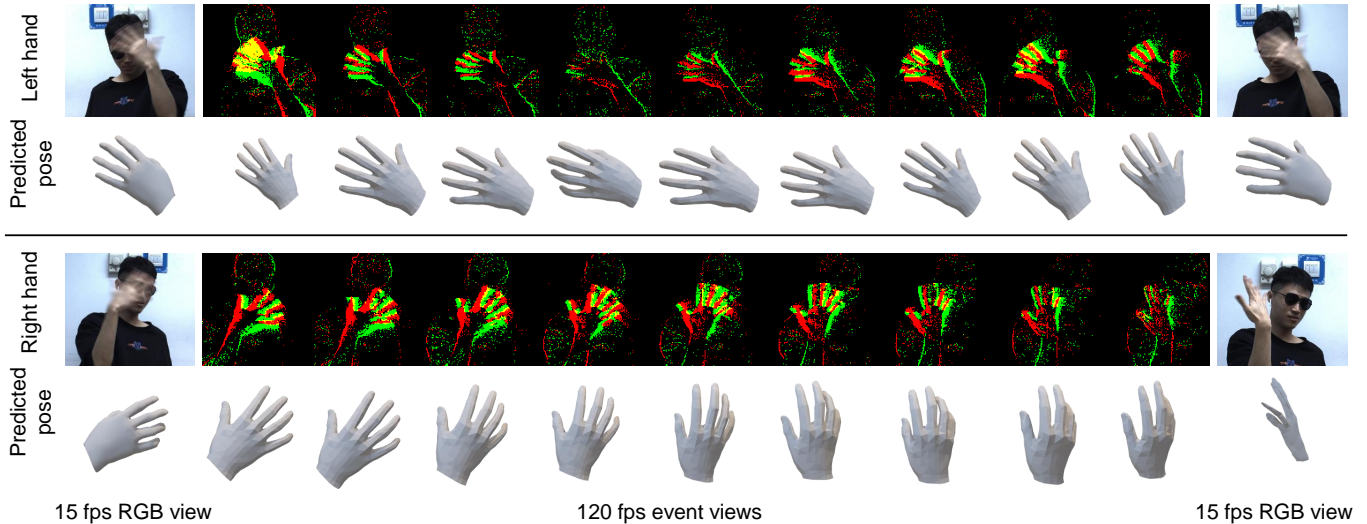


Fig. S1: Illustration of 120 fps event-based hand pose estimation. The two columns on the left and right end are neighboring RGB frames at 15 fps and their predicted hand poses by MobRecon [5]. The nine middle columns are the corresponding event frames and predicted poses by EvHandPose at 120 fps. EvHandPose achieves robust pose estimation compared with MobRecon [5] at much higher fps.

of flash sequences, and no fast motion data. We evaluate the method in indoor (8 sequences under normal scenes, 4 sequences of strong light, 4 sequences of flash, and 2 sequences of fast motion), DAVIS346 outdoor scenes (3 sequences under challenging outdoor illumination, 3 fast motion sequences), and PROPHESSEE outdoor scenes. In the sequences mentioned above, the left and right hands are equally divided. In our training process, we flip the left hand data into the right hand for increasing the amount of data.

#### A.4 Additional Main Results

Figure S2 shows the 3D PCK curves and AUC of EvHandPose and baselines under various scenes. EvHandPose outperforms EventHands [3] and the method of Jalees *et al.* [4] in all the scenes with AUC 0.16 higher. Besides, EvHandPose outperforms RGB-based method MobRecon [5] by AUC 0.11  $\sim$  0.22 higher under strong light scenes, and achieves comparable results under normal and flash scenes. Figure S3 shows additional quantitative results in strong light scenes between EvHandPose and MobRecon [5]. It can be concluded that EvHandPose can predict more stable hand poses under strong light scenes than MobRecon [5], which derives from the effective hand motion representations to retain the hand movements under challenging illumination.

#### A.5 Comparison of Flow Representations

Several existing works encode the meshes into the flow representations. EventHPE [8] projects body flow with vertex movements to the image plane with bilinear interpolation, denoted as linear flow, to encode the body shape. Hasson *et al.* [9] perform linear interpolation on vertex motion and differentially render the hand motion onto the image plane, denoted as vertex flow. We perform qualitative and quantitative comparison among linear flow, vertex flow, and our mesh flow to show the superiority of mesh flow for event-based hand pose estimation.

**Qualitative Comparison of Flow Representations** Hand flow representation should approximate the hand motion as precisely as possible, *i.e.*, the events triggered at the same hand position

should gather together to form an IWE with high contrast after being warped by the hand flow. In Figure S4, we show three kinds of hand flow representations and their corresponding IWEs on three sequences. As shown in the visualization of hand flows, linear flow leads to holes in flow representation, because the vertices of triangles on the mesh are not dense enough to cover the image plane. Vertex flow can not approximate the rotation motion of fingers, and the events triggered from the fingers are not warped efficiently. Mesh flow performs interpolation on the hand parameters to approximate the finger motion and renders the hand motion differentially, thus leading to IWEs with higher contrast.

**Quantitative Comparison of Flow Representations** The quality of hand flow representation is essential for the contrast of IWE, thus influencing the effect of contrast maximization constraint in our self-supervision framework. We evaluate the performance of EvHandPose under the CM loss by different hand flow representations. As shown in Table S2, our mesh flow leads to 1.5  $\sim$  3 mm MPJPE decrease compared with linear flow and vertex flow on EvRealHands. The main reason is that mesh flow can describe the hand movement precisely and impose strong priors for the constraints in self-supervision.

#### A.6 Comparison of Hand-edge Loss

Hand-edge loss is designed to enforce the alignment of the projection of predicted hand mesh and the edge of IWE, while it is challenging to find reliable correspondences between events and hand mesh vertices. EventCap [10] obtains the closest event for each vertex by the criterion based on the spatial and temporal distance in its event-based pose refinement step (denoted as vertex-to-event). However, it suffers from two issues: 1) Events are not aligned at the same time and 2) fingers will gather together in its optimization loop.

We compare the effect of the hand-edge losses using Pose-to-IWE module (denoted as IWE-to-vertex) and vertex-to-event strategies by applying them as constraints in hand model fitting problem. We select 65 annotated event timestamps in normal random sequences as testing samples. Given an annotated event



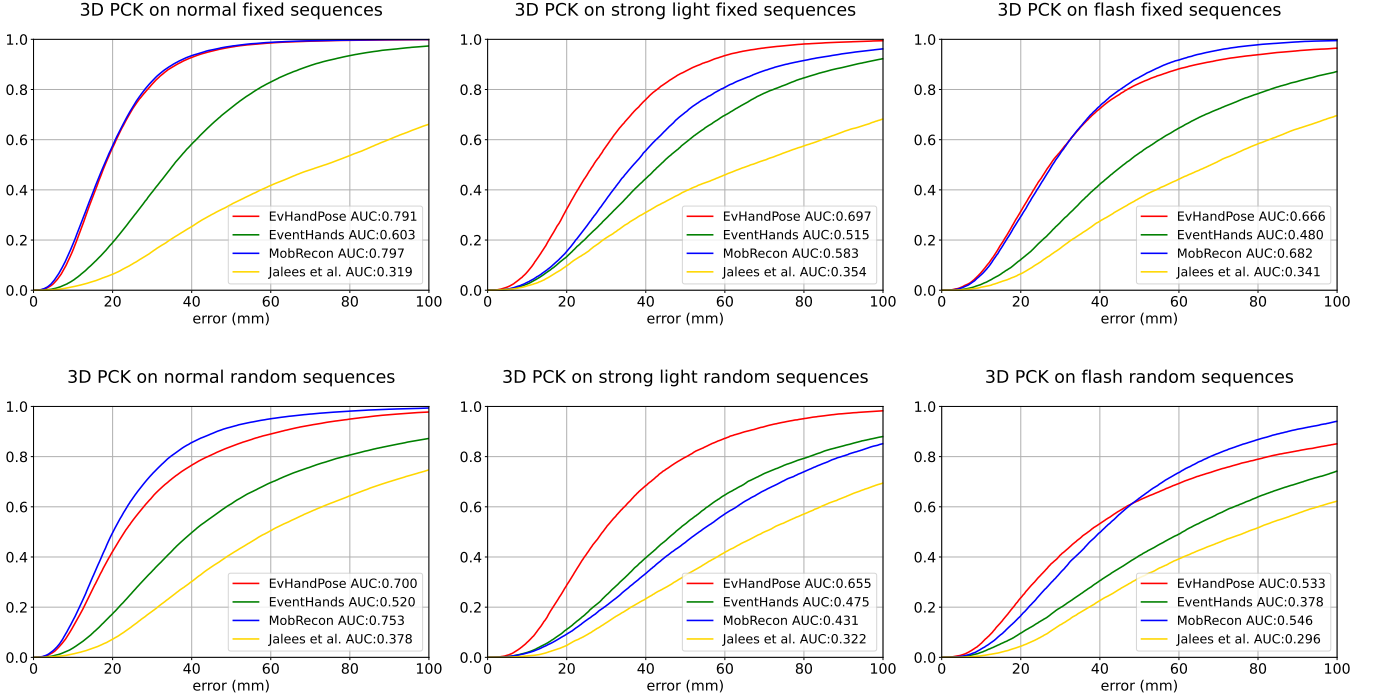


Fig. S2: 3D PCK curves of different methods on different scenes.




Fig. S3: Errors on two strong light sequences. Under strong light, EvHandPose has lower prediction error than MobRecon [5].

sub-segment, we add Gaussian noise with variance  $\sigma$  on its ground truth hand parameters as its initial hand parameters:

$$\varphi_{\text{init}} = \varphi_{\text{gt}} + N(0, \sigma), \quad (\text{S1})$$

where  $\varphi_{\text{init}}$  is initial hand parameters,  $\varphi_{\text{gt}}$  is ground truth hand parameters, and  $N(\cdot, \cdot)$  is the Gaussian noise.

Then the hand model fitting problem can be formulated as solving a non-linear optimization problem:

$$\varphi_{\text{opt}} = \arg \min_{\varphi} (\lambda_{\text{edge}} \mathcal{L}_{\text{edge}} + \lambda_{\text{smooth}} \mathcal{L}_{\text{smooth}}), \quad (\text{S2})$$

TABLE S2: Quantitative experiments results of different hand flow representations.

Scenes	Poses	Metrics	EvHandPose mesh flow	EvHandPose linear flow	EvHandPose vertex flow
Normal	Fixed	MPJPE AUC	<b>20.98</b> <b>0.780</b>	24.29 0.761	23.91 0.765
	Random	MPJPE AUC	<b>31.16</b> <b>0.690</b>	34.25 0.669	34.42 0.667
Fast motion	Random	2D-MPJPE	<b>8.18</b>	11.23	11.69
Strong light	Fixed	MPJPE AUC	<b>29.62</b> <b>0.691</b>	31.38 0.683	31.04 0.684
	Random	MPJPE AUC	<b>38.45</b> <b>0.625</b>	38.77 0.620	38.47 0.622
Flash	Fixed	MPJPE AUC	<b>33.31</b> <b>0.656</b>	37.88 0.639	38.51 0.636
	Random	MPJPE AUC	<b>49.77</b> <b>0.525</b>	51.99 0.518	52.11 0.520

where  $\varphi_{\text{opt}}$  is the optimal hand parameters,  $\mathcal{L}_{\text{edge}}$ ,  $\mathcal{L}_{\text{smooth}}$  and their loss weights  $\lambda_{\text{edge}}$ ,  $\lambda_{\text{smooth}}$  are the same with those in our self-supervision framework. We use Adam [7] optimizer with learning rate 0.005 for the iterative optimization loop. For quantitative analysis, we use aligned 3D MPJPE as evaluation metric.

As shown in Table S3, our hand-edge loss can effectively reduce the MPJPE by 10% in the optimization process, while vertex-to-event hand-edge loss is less effective. Qualitative analysis in Figure S5 shows that IWE-to-vertex hand-edge loss can prevent fingers from gathering together in the optimization and obtain stable results. As shown in Figure S6, one reason is that we warp the events into the target timestamp to align the events temporally. The other reason is that finding vertex correspondence for each IWE pixel leads to reliable alignments than finding event correspondence for each vertex.

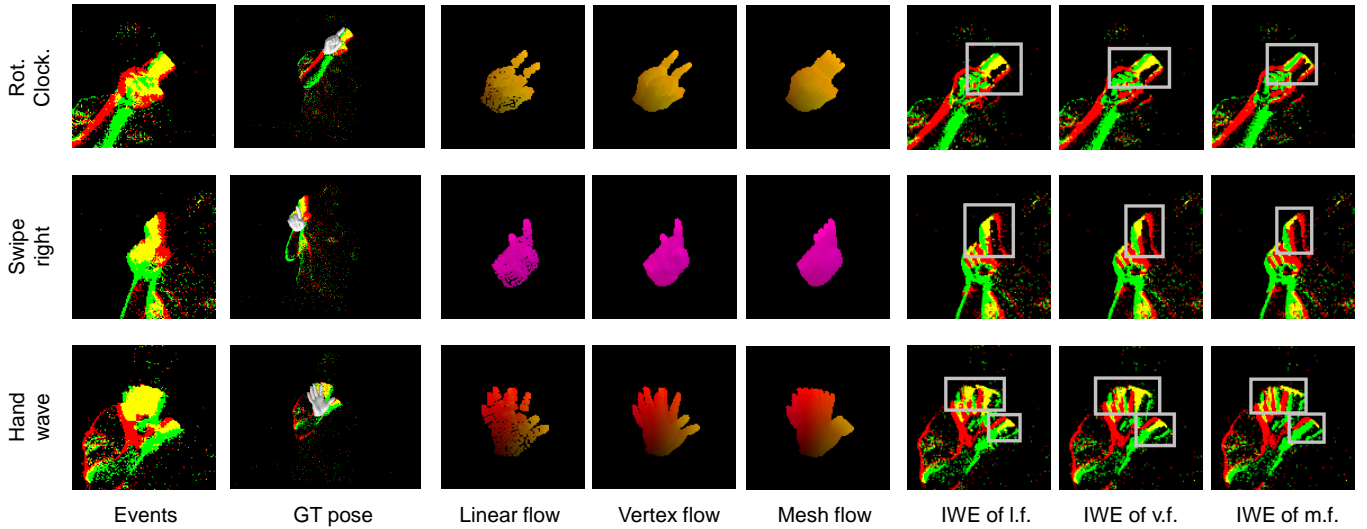


Fig. S4: Comparison among linear flow (l.f.), vertex flow (v.f.), and mesh flow (m.f.). In IWEs, red pixels represent positive events, green pixels represent negative events, and yellow pixels represent both types of events. An IWE with higher contrast will have aggregations of red and green pixels and fewer yellow pixels. As shown in the gray boxes, IWEs by our mesh flow have sharper edges than the others.

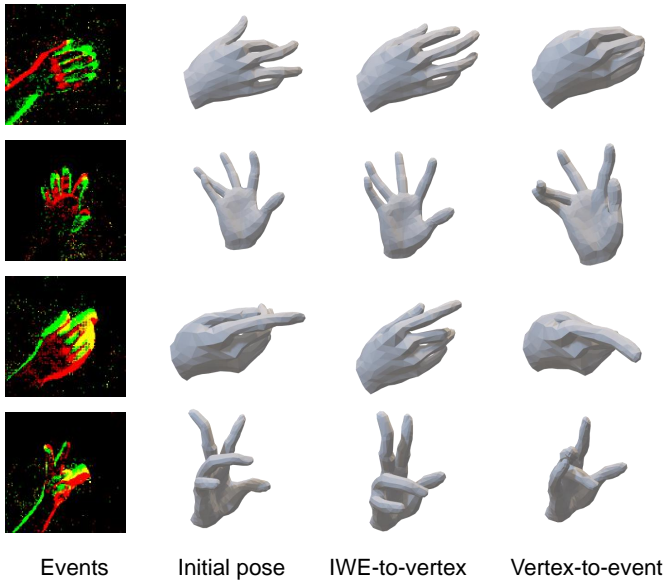


Fig. S5: Optimized hand poses with hand-edge losses using IWE-to-vertex and vertex-to-event strategies. Our hand-edge loss with IWE-to-vertex can prevent the fingers from gathering together and produce stable hand poses.

## APPENDIX B ADDITIONAL DETAILS OF DATASET EVREALHANDS

### B.1 Adequacy of Data

The amount of sequences and annotations of our dataset under various scenes is shown in Table S4. Verifying the adequacy of our dataset is important in dataset creation [11]. To achieve this, we evaluate the performance of EvHandPose on training sets with different numbers of subjects ranging from 1 to 7. As shown in Figure S7, EvHandPose achieves better performance as the number of subjects in training data increases. Since there is only one

TABLE S3: MPJPEs of the two hand-edge losses using IWE-to-vertex and Vertex-to-event strategies.  $\Delta$ MPJPE is the change of MPJPE before and after optimization.

$\sigma$	Initial MPJPE	IWE-to-vertex $\Delta$ MPJPE	Vertex-to-event $\Delta$ MPJPE
0.16	23.25	-1.95	+3.40
0.2	28.90	-2.98	+0.81
0.24	34.44	-3.40	-0.81
0.28	39.87	-3.84	-0.53
0.32	45.16	-5.35	-1.55

TABLE S4: Detailed attributes of EvRealHands. The duration of event streams, the quantity of RGB images, and their annotations in each scene are shown below. Annotations [mc] means the 3D annotations are checked or annotated manually, while Annotations [all] refer to all the machine and manual annotations. For fast motion sequences, the annotations are 2D keypoints.

Data	Normal		Strong light		Flash		Fast motion
	Fixed	Random	Fixed	Random	Fixed	Random	
Event streams (s)	850.7	2874.5	210.3	348.2	158.8	157.9	242.4
RGB images	79.7 K	267.6 K	14.8 K	28.3 K	14.9 K	13.8 K	7.8 K
Annotations [mc]	11.4 K	9.7 K	2.0 K	1.9 K	2.1 K	1.9 K	0.8 K
Annotations [all]	11.4 K	37.2 K	2.0 K	3.9 K	2.1 K	1.9 K	0.8 K

subject captured under strong light and flash in our training set, the evaluations are conducted mostly on the collected data in normal scenes. When more than 5 subjects are used for model training, the increase in subjects will no longer significantly improve the performance. EvHandPose achieves better results in strong light and fast motion scenes but very little boost in flash scenes. One explanation is that the edge and hand flow features learned by EvHandPose in normal scenes can be effectively generalized to the strong light and fast motion scenes. In flash scenes, the events generated by the background often overwhelm the events generated from the hand movement, so it makes it challenging for EvHandPose to extract effective motion clues. In future research,

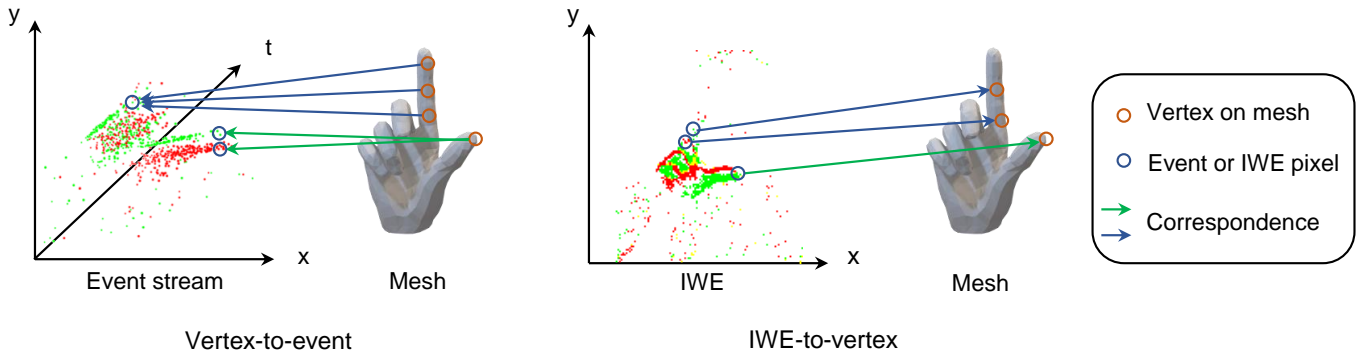


Fig. S6: Comparison of vertex-to-event and IWE-to-vertex method for hand-edge losses. In vertex-to-event method, an event may correspond to several vertices (shown in blue arrows), which will make fingers gather together in optimization. However, in IWE-to-vertex method, one IWE pixel can only match one vertex, which can avoid the issue of vertex-to-event method. In vertex-to-event method, vertex matches its corresponding event using the spatial and temporal distance, which will lead to misalignment (shown in green arrows) because temporal distance is not a good measure of the spatial distance at the target timestamp. In IWE-to-vertex method, we solve this issue by warping the events to the target timestamp.

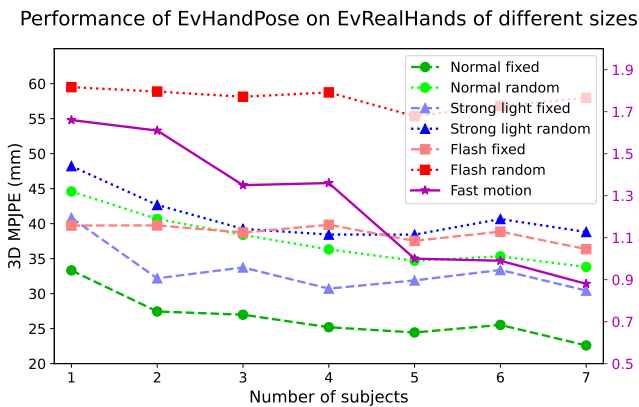


Fig. S7: Performance of EvHandPose on EvRealHands with different sizes. EvHandPose nearly reaches the best performance when there are 5 subjects in the training set.

more data under challenging scenes (strong light, fast motion, etc.) can be collected to bridge the performance gap between normal scenes and challenging scenes.

**B.2 Scanned Hand Meshes**

For shape annotation, we use an Artec Eva 3D Scanner to scan the hands of each subject, and the shapes are shown in Figure S8.




Fig. S8: Scanned hand meshes. We use 3D scanner to get precise hand mesh for shape annotation.

**B.3 Hand Gestures**

We visualize the event streams of 15 hand gestures similar to [2] in Figure S9. These gestures are *Grab, Tap, Expand, Pinch, Rotation, Rotation counter clockwise, Swipe right, Swipe left, Swipe up, Swipe down, Swipe X, Swipe +, Swipe V, Shake, Handwave*. The visualization shows that event cameras can effectively capture the hand movements from the still background. By annotating the sequences, we obtain 15 gesture classes for hand gesture recognition.

**REFERENCES**

- [1] W. Zhang, Z. Lin, J. Cheng, C. Ma, X. Deng, and H. Wang, “STA-GCN: two-stream graph convolutional network with spatial-temporal attention for hand gesture recognition,” *The Visual Computer*, 2020.
- [2] Q. De Smedt, H. Wannous, J.-P. Vandeborre, J. Guerry, B. L. Saux, and D. Filliat, “3D hand gesture recognition using a depth and skeletal dataset: Shrec’17 track,” in *Proceedings of the Workshop on 3D Object Retrieval*, 2017.
- [3] V. Rudnev, V. Golyanik, J. Wang, H.-P. Seidel, F. Mueller, M. Elgharib, and C. Theobalt, “EventHands: Real-time neural 3D hand pose estimation from an event stream,” in *Proc. of International Conference on Computer Vision (ICCV)*, 2021.
- [4] J. Nehvi, V. Golyanik, F. Mueller, H.-P. Seidel, M. Elgharib, and C. Theobalt, “Differentiable event stream simulator for non-rigid 3D tracking,” in *Proc. of Conference on Computer Vision and Pattern Recognition (CVPR)*, 2021.
- [5] X. Chen, Y. Liu, Y. Dong, X. Zhang, C. Ma, Y. Xiong, Y. Zhang, and X. Guo, “MobRecon: Mobile-friendly hand mesh reconstruction from monocular image,” *Proc. of Conference on Computer Vision and Pattern Recognition (CVPR)*, 2022.
- [6] K. He, X. Zhang, S. Ren, and J. Sun, “Deep residual learning for image recognition,” in *Proc. of Conference on Computer Vision and Pattern Recognition (CVPR)*, 2016.
- [7] D. P. Kingma and J. Ba, “Adam: A method for stochastic optimization,” in *International Conference on Learning Representations (ICLR)*, 2015.
- [8] S. Zou, C. Guo, X. Zuo, S. Wang, P. Wang, X. Hu, S. Chen, M. Gong, and L. Cheng, “EventHPE: Event-based 3D human pose and shape estimation,” in *Proc. of International Conference on Computer Vision (ICCV)*, 2021.
- [9] Y. Hasson, B. Tekin, F. Bogo, I. Laptev, M. Pollefeys, and C. Schmid, “Leveraging photometric consistency over time for sparsely supervised hand-object reconstruction,” in *Proc. of Conference on Computer Vision and Pattern Recognition (CVPR)*, 2020.
- [10] L. Xu, W. Xu, V. Golyanik, M. Habermann, L. Fang, and C. Theobalt, “EventCap: Monocular 3D capture of high-speed human motions using an event camera,” in *Proc. of Conference on Computer Vision and Pattern Recognition (CVPR)*, 2020.

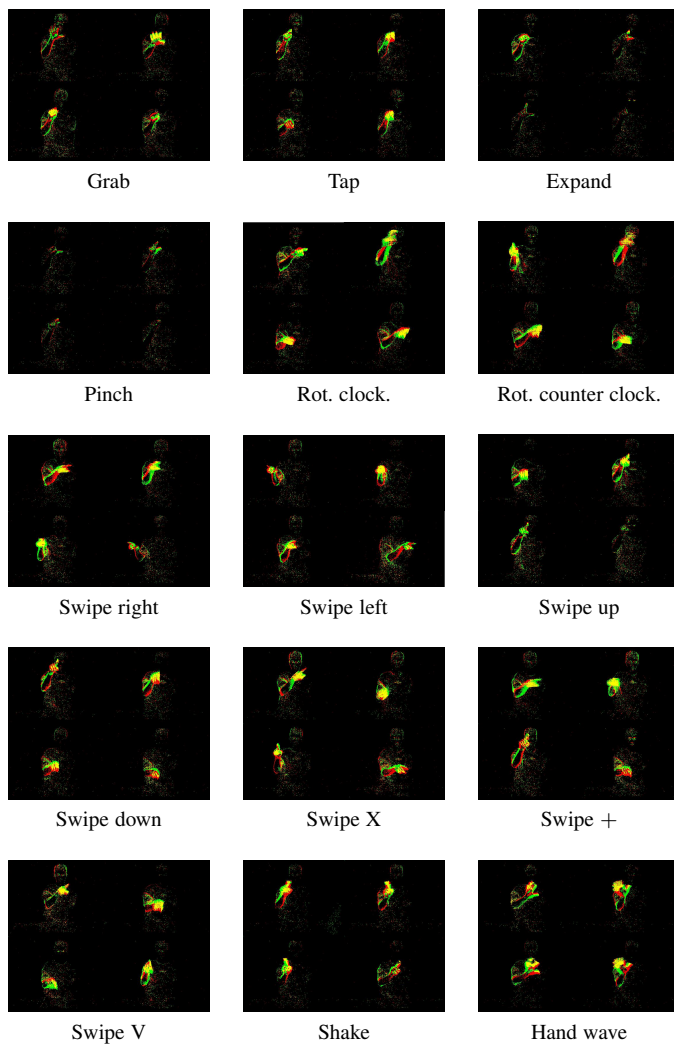


Fig. S9: Illustration of the 15 fixed pose gestures in event frames (“Rot.” and “clock.” are short for rotation and clockwise).

- [11] G. Varol, J. Romero, X. Martin, N. Mahmood, M. J. Black, I. Laptev, and C. Schmid, “Learning from synthetic humans,” in *Proc. of Conference on Computer Vision and Pattern Recognition (CVPR)*, 2017.