

HKUST SPD - INSTITUTIONAL REPOSITORY

Title	Event-based Motion Segmentation by Cascaded Two-Level Multi-Model Fitting
Authors	Lu, Xiuyuan; Zhou, Yi; Shen, Shaojie
Source	IEEE/RSJ International Conference on Intelligent Robots and Systems, IROS 2021, / IEEE. Piscataway, NJ : IEEE, 2021, p. 4445-4452
Version	Accepted Version
DOI	10.1109/IROS51168.2021.9636307
Publisher	IEEE
Copyright	© 2021 IEEE. Personal use of this material is permitted. Permission from IEEE must be obtained for all other uses, in any current or future media, including reprinting/republishing this material for advertising or promotional purposes, creating new collective works, for resale or redistribution to servers or lists, or reuse of any copyrighted component of this work in other works.

This version is available at HKUST SPD - Institutional Repository (<https://repository.ust.hk>)

If it is the author's pre-published version, changes introduced as a result of publishing processes such as copy-editing and formatting may not be reflected in this document. For a definitive version of this work, please refer to the published version.

Event-based Motion Segmentation by Cascaded Two-Level Multi-Model Fitting

Xiuyuan Lu, Yi Zhou*, and Shaojie Shen

Abstract—Among prerequisites for a synthetic agent to interact with dynamic scenes, the ability to identify independently moving objects is specifically important. From an application perspective, nevertheless, standard cameras may deteriorate remarkably under aggressive motion and challenging illumination conditions. In contrast, event-based cameras, as a category of novel biologically inspired sensors, deliver advantages to deal with these challenges. Its rapid response and asynchronous nature enables it to capture visual stimuli at exactly the same rate of the scene dynamics. In this paper, we present a cascaded two-level multi-model fitting method for identifying independently moving objects (i.e., the motion segmentation problem) with a monocular event camera. The first level leverages tracking of event features and solves the feature clustering problem under a progressive multi-model fitting scheme. Initialized with the resulting motion model instances, the second level further addresses the event clustering problem using a spatio-temporal graph-cut method. This combination leads to efficient and accurate event-wise motion segmentation that cannot be achieved by any of them alone. Experiments demonstrate the effectiveness and versatility of our method in real-world scenes with different motion patterns and an unknown number of independently moving objects.

I. INTRODUCTION

Modern robotic applications (e.g., autonomous driving, AR/VR, UAVs, etc) require synthetic agents to perceive dynamic scenes in order to interact with the environment effectively. As one of the challenges cast by dynamic scenes, the multi-motion segmentation problem aims at detecting and further modeling the motions of independently moving objects (IMOs) while estimating the camera's ego-motion. The problem's complexity originates from its exploratory nature: multiple hypotheses (models) are to be considered and validated against the data to select the optimal fits. Moreover, from a sensing perspective, challenging visual effects (e.g., motion blur and over/under-exposure) make the problem even harder for solutions using standard cameras.

Event-based cameras, for example, the DVS [1], are biologically-inspired visual sensors that mimic the transient pathway of the human visual system. They acquire visual information in a completely different way from standard cameras, and consequently, produce a stream of asynchronous per-pixel intensity changes (the so-called "events") instead of intensity images. This novel principle of operation brings advantages over standard frame-based cameras to address the

All authors are with the Dept. of ECE, Hong Kong University of Science and Technology, Hong Kong, China. Email: {xluj, eeyzhou, eeshaojie}@ust.hk. This work was supported by the HKPFS and HKUST PDF matching fund 2020. *Corresponding author.

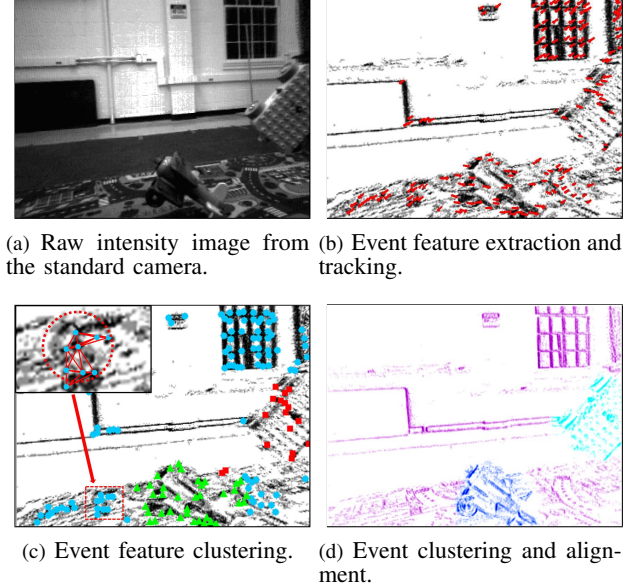


Fig. 1: The proposed cascaded two-level multi-model fitting scheme for event clustering. The raw intensity image (a) is for visualization only. Event features are extracted on the uncompensated image of warped events and continuously tracked on the image plane (b). At level one (c), event features are clustered into groups that undergo different rigid motions. At level two (d), events are clustered by associating to the refined motion models and are aligned along corresponding point trajectories.

challenging scenarios in robotic vision, such as high-speed and/or high dynamic range (HDR) stereo depth estimation [2, 3], camera tracking [4, 5] and Simultaneous Localization and Mapping [6]–[8].

In this paper, we look into the problem of event-based motion segmentation, which aims at clustering events occurred within a short time interval into groups, such that each of them complies with a unique rigid motion (background motion or an IMO's motion). A novel method is developed that jointly clusters the events (data-model association) and estimates the parameters of their coherent motion (model fitting). The contribution of the work is summarized as:

- A cascaded two-level multi-model fitting scheme that solves the event-based motion segmentation problem by two steps: event feature clustering (level one) and event clustering (level two).




Fig. 2: Overview of the proposed system. The input to the system is a stream of events occurred within a short time interval. The proposed system consists of two multi-model fitting modules cascaded sequentially. Using as input the event features’ tracking, level one clusters features (labeled by \mathcal{L}_f) into groups that comply with different rigid motions \mathcal{M} . Initialized by level one, events are clustered (labeled by \mathcal{L}_e) at level two by associating with refined motion models \mathcal{M}^* .

- An efficient hypothesis proposing unit (level one) developed by integrating a robust model fitting method into a progressive multi-model fitting pipeline. The result of level one – a compact set of motion instances – is used to initialize the process in level two, which leads to efficient event-wise labeling.
- An extensive evaluation on available datasets showing state-of-the-art performance.

In the rest of the paper, we first review related work (Section II), and then explain our method (Section III and IV). Finally, the proposed approach is evaluated in Section V.

II. RELATED WORK

The literature review focuses on works that do not require prior knowledge, including the shape of IMOs, and the correlation between the tracked geometric primitives and the event camera’s motion. Early prior-free pipelines (e.g., [9, 10]) follow a greedy strategy, which first recover the dominant background motion and deal with the remaining events induced by the IMOs subsequently.

Using the result of [10] as initialization, [11] presented the first method that solved the model estimation (fitting) and event-motion association (labeling) sub-problems jointly. The segmentation problem was formulated in an expectation-maximization (EM) scheme, which iteratively switched between updating the event-motion association and refining the motion models. Both the E-step and M-step leveraged the idea of motion compensation [12]. The method, however, required to know the number of IMOs as a prior knowledge.

More recently, a hierarchical clustering method was presented in [13] and further improved in [14]. The method first clustered event features by applying the K-means method. Then the resulting redundant clusters were passed to an iterative method, which consisted of clusters merging and motion parameters refinement. The method does not require to know the number of IMOs in advance (by setting K to a big number). However, the clusters merging criteria was somehow designed under a hybrid metric of image contrast and centroid distance. Thus, clusters undergoing identical motions could not be merged if their centroids were far away from each other. Finally, the event-motion association (event-wise labeling) was made according to if an event was within the convex hull of the features in a cluster. Inaccurate event labeling was witnessed near the overlapping areas of background and IMOs.

To address the flaws of [11] [13], the method in [15] adapted motion compensation [12] to a graph-based energy formulation. The idea of negative representation [8] is borrowed so that the original maximization problem is converted to a minimization one. The number of clusters could converge to the real number of IMOs because of the applied constraint on the length of description. However, this method failed to detect IMOs with a small size because of the bad inlier-outlier ratio in the initialization. In addition to the above methods, we have to mention a supervised end-to-end learning-based pipeline in [16], which jointly estimated optical flow, 3D motion and object segmentation. It provided a state-of-the-art dataset for evaluation of IMOs segmentation, though it was not close to any other approaches in the literature review.

Our method draws on the advantages of existing methods and overcomes their shortcomings with a new design. Inspired by [15], we also formulate the event-based segmentation as a geometric multi-model fitting problem. Unlike [15], we propose a cascaded two-level multi-model fitting scheme. In the first level, we bring in event feature tracking [13, 14, 17] for efficient hypothesis proposal. Instead of using K-means [13, 14], we cluster the features with a progressive multi-model fitting method. The model fitting and feature-motion association sub-problems are solved jointly, thus not requiring incremental cluster merging locally [13, 14]. In the second level, we leverage the spatio-temporal graph cut method [15] to solve the event clustering problem.

III. PROBLEM STATEMENT AND PRELIMINARIES

The goal is to cluster events occurred within a short time interval into several groups such that each complies with a unique rigid motion. Each motion model parametrizes the 2D point trajectories of events on the image plane induced by an IMO. Once the events are clustered, they can be warped according to the estimated motions, thus producing an image of warped events with the highest contrast (Fig. 1 (d)).

In this section, we first review the principles of event cameras and clarify the input of our algorithm. Second, we discuss the motion model fitting methods using as input the feature correspondences and the raw events, respectively. Finally, the event-based motion segmentation problem is identified as a geometric multi-model fitting one, followed by an overview of the proposed cascaded two-level multi-model fitting scheme.

A. Event Stream and Features

Unlike standard cameras which capture full frames at a fixed frame rate, event cameras have independent pixels that sense brightness changes asynchronously. An event is triggered if the variation of the sensed intensity (in log scale) $I(\mathbf{x}, t)$, at a pixel location $\mathbf{x} = (u, v)^T$ and time instance t , exceeds a nominal threshold C_{th} :

$$|\log I(\mathbf{x}, t) - \log I(\mathbf{x}, t - \Delta t)| > C_{th}, \quad (1)$$

where Δt is the time since the last event was triggered at this pixel. Each event e , denoted by a tuple $\langle u, v, t, p \rangle$, records spatio-temporal information that includes u, v the pixel coordinate of the event, t the timestamp at which the event occurs, and $p \in \{-1, +1\}$ the polarity indicating the sign of the brightness change.

We denote the input of our algorithm, a stream of events occurred within a short time interval $[t - \delta t, t]$, as $\mathcal{E}_{[t - \delta t, t]} = \{e_k\}_{k=1}^{N_e}$. Besides, we also use the event feature correspondences as input. We extract event features at time instance $t - \delta t$ and t , respectively, and establish feature correspondences by tracking the 2D motion of these features (detailed in Section IV-A). The set of feature correspondences is denoted by $\mathcal{F} = \{\mathbf{f}_i\}_{i=1}^{N_f}$, where $\mathbf{f}_i = \{\mathbf{x}_i^{t - \delta t} \leftrightarrow \mathbf{x}_i^t\}$ refers to a pair of feature correspondence.

B. Motion Model Fitting

Here we clarify the ways to fit a motion model with feature correspondences and raw events, respectively. At level one of our method, motion models can be estimated in a closed form given feature correspondences. Using the four-parameter model $\mathbf{m} := \{m_u, m_v, m_s, m_\theta\}^T$ [9] as an example, the point trajectory on the image plane is parametrized as:

$$\mathbf{x}'_1 = \mathbf{x}_1 + \left[\begin{pmatrix} m_u \\ m_v \end{pmatrix} + (m_s + 1)\mathbf{R}_{m_\theta}\mathbf{x}_1 - \mathbf{x}_1 \right] \delta t, \quad (2)$$

where $\mathbf{x}_1 \leftrightarrow \mathbf{x}'_1$ denotes a pair of corresponding features, $\mathbf{R}_{m_\theta} := \begin{pmatrix} \cos m_\theta & -\sin m_\theta \\ \sin m_\theta & \cos m_\theta \end{pmatrix}$ the in-plane rotation, and δt the time interval during which the features are tracked. The minimal case of Eq. 2 requires two pairs of feature correspondences, which can be solved by the direct linear transform (DLT) method. Other models, such as 2D-flow model [18], and 3D-Rotation model [19], can be worked out similarly. For a single-instance case in presence of outliers, a robust solution typically uses a minimal-set solver inside a RANSAC framework [20] and subsequently refines the estimate via a nonlinear optimization with all inliers.

In the second level, the resulting motion instances by the first level will be refined with all associated events. We apply the *motion compensation* [12] method for model fitting with raw events. The idea of *motion compensation* refers to the process of warping all involved events with an associated motion model to a reference time instance. During the process, events are aligned incrementally along the point trajectories on the image plane, and the best fit is found when the contour strength (contrast) of the image of warped events (IWE) is maximized. The contour strength can be assessed by

a variety of dispersion metrics. Here, we utilize the variance loss because of its advantages in terms of accuracy and computation complexity over other alternatives [21]. The motion compensation method can be summarized as follows. First, all involved events \mathcal{E} are geometrically transformed according to a warping function \mathbf{W} ,

$$e_k \doteq (\mathbf{x}_k, t_k) \mapsto e'_k \doteq (\mathbf{x}'_k, t_{ref}), \quad (3)$$

resulting in a number of warped events $\mathcal{E}' = \{e'_k\}_{k=1}^{N_e}$ at a reference time t_{ref} . The warping function \mathbf{W} is parametrized by the motion model \mathbf{m} . It determines the image-plane trajectories along which the events are aligned. Second, the warped events \mathcal{E}' are aggregated into an IWE,

$$I(\mathbf{x}; \mathbf{m}) \doteq \sum_{k=1}^{N_e} \delta(\mathbf{x} - \mathbf{x}'_k(\mathbf{m})), \quad (4)$$

where each pixel \mathbf{x} counts the number of warped events that fall within it. To establish a differentiable objective in the contrast maximization, the Dirac function $\delta(\cdot)$ is typically replaced by a Gaussian $\mathcal{N}(\mathbf{x}; \mathbf{0}, \epsilon^2 \text{Id})$ of $\epsilon = 1$ pixel width. Finally, the variance of the IWE quantifies the goodness of fitting: $\mathbf{m}^* = \arg \max_{\mathbf{m}} \sigma^2(I(\mathbf{x}; \mathbf{m}))$. Our formulation supports any type of parametric motion model.

C. Segmentation via Geometric Multi-Model Fitting

Segmentation refers to the problem of clustering data according to some proximity in data, such as identical color and texture pattern (photometric proximity) or being consistent with a common geometric model (geometric proximity). Event-based motion segmentation lies in the latter case because the absolute intensity is not available in raw data. The problem can be posed as an energy-based multi-model fitting problem, which jointly solves data-model association and model fitting sub-problems. The data-model association sub-problem, also known as labeling, is a discrete optimization problem. To solve this sub-problem, a graph structure is typically established on the data, which introduces additional consideration on spatial regularity, and more importantly, enables to apply modern graph-cut solvers. The overall energy has a unified form $E(L, \mathcal{M}) = E_{data}(L, \mathcal{M}) + E_{reg}(L, \mathcal{M})$, which is a function of two sets of variables: L the labeling function that associates a model to a data, and \mathcal{M} the models that represent the commonality shared within each cluster.

We follow this formulation in the sense of energy-based multi-model fitting with a graph structure. Furthermore, we propose a cascaded two-level multi-model fitting scheme. In each level, the motion segmentation problem is formulated as a joint optimization problem over the data-model (features-motion in the first level, and events-motion in the second level) association and the motion models' parameters. The data-model association is represented by the labeling function $L(d^1) : \mathbb{D} \rightarrow \mathcal{L} = \{1, \dots, N\}$, which associates a data d to a label $l \in \mathcal{L}$ indicating which IMO the data belongs to. The motion patterns of IMOs are represented by a set

¹ d refers to a data which could be either an pair of corresponding features $\mathbf{f} \in \Omega \times \Omega$ or an event $e \in \Omega \times \mathcal{T}$.

of motion models $\mathcal{M} = \{\mathbf{m}_1, \dots, \mathbf{m}_N\}$. Each motion model encodes the coherent motion that a cluster of data complies with. The overall system is illustrated in Fig. 2. We detail the multi-model fitting method at each level in Section. IV.

IV. METHODOLOGY

In this section, we detail the proposed multi-model fitting method at each level. First, we discuss the approach for feature clustering at level one, which incorporates robust motion estimation into a progressive multi-model fitting scheme. Using the resulting motion instances as an initial model pool, we show in the second level that the event clustering problem can be efficiently addressed with a spatio-temporal graph cut method.

A. Feature Clustering by Progressive Multi-Model Fitting

The multi-model fitting method, at level one, jointly solves the motion parameters estimation and feature-motion association sub-problems, using the feature correspondences as input. We extract Shi-Tomasi corners [22] on raw IWEs (without motion compensation²) at time $t - \delta t$ and t , respectively, and obtain feature correspondences using Lucas-Kanade optical flow method [23] (Fig. 1(b)). To cluster event features into groups that comply with different motion patterns (Fig. 1(c)), we perform motion model fitting using feature correspondences with RANSAC and integrate it into the Progressive-X scheme [24]. In this scheme, motion model hypotheses are proposed one-by-one and added to a set of active instances maintained by the PEARL optimization [25], which looks for the optimal configuration over all instances and feature correspondences. The result of this level consists of the labeling of features \mathcal{L}_f and associated motion instances \mathcal{M} . We detail the key components in this scheme as follows.

Hypothesis Proposal. To propose sufficient motion model hypotheses, a proposal engine is needed to increasingly generate yet unseen hypotheses. A variant of RANSAC algorithm – Graph-Cut (GC) RANSAC [26] is applied for this purpose. The GC-RANSAC method consists of two key steps: sampling and inliers/outliers clustering.

1) Sampling. A minimal set \mathcal{S} is sampled, with which a motion model hypothesis is fitted using the linear solver (III-B). Applicable sampler choices include NAPSAC [27] and PROSAC [28]. We utilize NAPSAC because of its local sampling characteristics.

2) Inliers/outliers clustering. Inliers are separated from outliers through a binary classification process. To this end, a neighbourhood graph is established, with each node representing a feature correspondence. Actually, the topology of the neighbourhood graph is simply determined by the feature locations (nodes) on one IWE and a naive neighbour searching scheme determining the neighbourhood (edges), and thus, leading to a 2D graph (see zoom-in area of Fig. 1(c)). The binary classification task is formulated as a

standard Markov Random Field (MRF) problem, consisting of a data term and a smoothness term. The data term is designed in a way that rewards a data point close to a model if it is labeled as an inlier, while a modified *Potts* model is applied as the smoothness term to deal with extreme cases in which the data contain significantly more outliers close to the desired model than inliers [26]. The model hypothesis is then refitted with all inliers using a nonlinear optimizer.

In addition to these two steps, we need a quality function to verify a hypothesis. The quality function's return reflects the support from data. The design of the qualifying criteria is different from ordinary RANSAC algorithms specified for a single-instance fitting problem. In a single-instance case, the quality function measures the goodness of fit w.r.t a unique hypothesis. The simplest way is to count inliers. On the contrary, in a multi-instance fitting problem, the quality function is supposed to count data that are inliers only to the hypothesis h , while being outliers to other existing instances h_{\cup} (called compound model instance in [24]). In our problem, a data point \mathbf{f} (a pair of corresponding features) is regarded as an inlier to a motion model hypothesis if the geometric error $e_{\mathbf{f}, \mathbf{m}_h} = \|\mathbf{x}_i^t - W(\mathbf{x}_i^{t-\delta t}; \mathbf{m}_h)\|$ is smaller than ζ (a predefined threshold), and the minimal geometric error from other model instances $e_{\mathbf{f}, \mathbf{m}_{h_{\cup}}} = \min_{j \in \{h_{\cup}\}} \|\mathbf{x}_i^t - W(\mathbf{x}_i^{t-\delta t}; \mathbf{m}_j)\|$ is greater than ζ . In practice, the truncated quality function of MSAC [29] has shown its superiority to inlier counting of RANSAC in terms of accuracy and sensitivity to a predefined threshold. Thus, we define our quality function Q as

$$Q(\mathbf{m}_h) = |\mathcal{F}| - \sum_{\mathbf{f} \in \mathcal{F}} \min \left(1, \max \left(\frac{e_{\mathbf{f}, \mathbf{m}_h}^2}{\gamma(\zeta)^2}, 1 - \frac{e_{\mathbf{f}, \mathbf{m}_{h_{\cup}}}^2}{\gamma(\zeta)^2} \right) \right), \quad (5)$$

where $|\cdot|$ denotes the cardinality of a set, and $\gamma(\zeta) = \frac{3}{2}\zeta$. As a consequence, a data point leads to a small score if it is consistent with both the newly proposed hypothesis and the existing compound instances or if it is inconsistent with the newly proposed hypothesis.

The above two steps are carried out iteratively until the probability of finding a hypothesis with better support is lower than a confidence threshold μ . As a consequence, the required number of iteration is calculated as $K = \frac{\log(1-\mu)}{\log(1-\eta^{|\mathcal{S}|})}$, where $\eta = |\mathcal{I}|/|\mathcal{F}|$ denotes the inlier ratio. The resulting hypothesis returned by the GC-RANSAC loop is added to the pool of active model instances, which participates in the following multi-instance optimization.

Multi-Instance Optimization. The goal of this step, on one hand, is to find some of the data a better fit through iterative relabeling and re-fitting. On the other hand, the newly added hypothesis is further validated to see if it is redundant under the joint concern over model consistency and spatial coherence. If not redundant, it will be maintained as an active instance in the model pool. The formulation of this task is similar to the inliers/outliers classification mentioned above. However, a more complex multi-instance fitting problem is to be solved. To this

²We observed no apparent difference in performance whether or not a global motion compensation is performed as pre-processing [13, 14].

Algorithm 1: Feature Clustering with Progressive Multi-Model Fitting

input : Feature correspondence set \mathcal{F} , threshold μ
output: Feature labeling \mathcal{L}_f , model instance \mathcal{M}
Initialize $\mathcal{L}_f = \emptyset$, $\mathcal{M} = \emptyset$, $\omega^* = 0$, $\bar{I} = |\mathcal{S}| + 2$
 $\mathcal{G} \leftarrow$ Neighbourhood graph (\mathcal{F})
while $\bar{I} > |\mathcal{S}| + 1$ **do** \triangleright Prog-X Loop [24]
 for $k = 1$ **to** K **do** \triangleright GC-RANSAC Loop [26]
 $\mathcal{S} \leftarrow$ Sampled by NAPSAC [27]
 $\mathbf{m}_h \leftarrow$ Linear model fitting (\mathcal{S})
 $\omega \leftarrow$ Compute support of \mathbf{m}_h (Eq. 5)
 while $\omega > \omega^*$ **do**
 $\omega^* \leftarrow \omega$, $\mathbf{m}_h^* \leftarrow \mathbf{m}_h$, $\mathcal{I}^* \leftarrow \mathcal{I}$
 $\mathcal{I} \leftarrow$ Graph-Cut ($\mathbf{m}_h, \mathcal{G}$)
 $\mathbf{m}_h \leftarrow$ Nonlinear model fitting (\mathcal{I})
 $\omega \leftarrow$ Compute support of \mathbf{m}_h (Eq. 5)
 end
 end
 $\mathcal{M} \leftarrow \mathcal{M} \oplus \mathbf{m}_h^*$
 $\{\mathcal{L}_f, \mathcal{M}\} \leftarrow$ PEARL(\mathcal{M}, \mathcal{F}) [25]
 $\bar{I} \leftarrow$ Compute the upper limit of $|\mathcal{I}^*|$ (Eq. 6)
end

end, the aforementioned neighbourhood graph is re-used. The PEARL [25] optimization method is used to solve this problem. Since the newly proposed hypotheses come one-by-one, the pool of instances is kept small, and thus, the computational complexity is not high. Consequently, the pool of model instances is augmented, and the data are relabeled by refined model instances.

Termination Criterion The progressive multi-model fitting pipeline, summarized in Alg. 1, runs iteratively until the termination criterion is met. According to [24], the overall progress is terminated when $\bar{I} < |\mathcal{S}| + 1$, where \bar{I} denotes the upper bound of the number of inliers to a not yet found instance. The inliers in this context refer to data that are independent on the compound model. The upper bound \bar{I} with confidence μ in the n_{th} iteration is defined as

$$\bar{I}(\mathcal{F}, h_{\cup}, |\mathcal{S}|, n, \mu) = (|\mathcal{F}| - |h_{\cup}|) \sqrt[n]{1 - \sqrt[n]{1 - \mu}}, \quad (6)$$

where, with a slight abuse of notation, $|h_{\cup}|$ denotes the number of inliers to the compound model. This criterion guarantees that, upon termination, the probability of an existing unseen instance with at least $|\mathcal{S}| + 1$ inliers is smaller than $1 - \mu$.

B. Event Clustering via Spatio-Temporal Graph Cut

The goal of this step is to achieve event-wise labeling such that each event is associated with a specific motion model. We borrow the idea of event-based motion segmentation in [15], which also formulates the segmentation problem as an energy-based multi-model fitting problem. In contrast to the 2D neighbourhood graph applied in IV-A, a 3-D

fully connected graph is established by exploiting the spatio-temporal characteristics of events. The proposed objective is a hybrid function of discrete labeling L and continuous model parameters \mathcal{M} , defined as

$$E(L, \mathcal{M}) \doteq E_D(L, \mathcal{M}) + E_P(L) + E_M(L), \quad (7)$$

which simultaneously considers the data-model consistency (data term E_D) and spatio-temporal coherence across the data (Potts model E_P). The data term considers the goodness of fit by adapting the idea of motion compensation, while the smoothness term enhances spatially coherent labeling. Additionally, an *Minimum Description Length* (MDL) term (E_M) is added in order to make the model pool as compact as possible. This constraint would reduce model redundancy and thus prompt the remaining models to correspond exactly to IMOs in the scene. Due to space limitations, the definition on each term is omitted. Readers can refer to [15] for more details. A block-coordinate descent strategy is used to solve Eq. 7, optimizing L and \mathcal{M} in an alternating manner. The original initialization method of [15] proposes hypotheses by fitting models to events within a number of hierarchically divided spatio-temporal volumes. This strategy, though free of optical flow computation, generates many similar (redundant) hypotheses together with false-positive ones, and thus, is relatively slow. Besides, it typically fails to detect small IMOs because of the bad inlier ratio. Thanks to the feature clustering in the first level, the resulting model instances can be used as the initial pool of models. With this compact and reasonably accurate pool of motion instances, the inference of event-model association can be done efficiently. We evaluate our algorithm in the following section.

V. EXPERIMENTS

In this section, we evaluate the proposed algorithm. First, we introduce the datasets and evaluation metrics used (Section V-A). Second, we provide quantitative and qualitative evaluation results on these datasets (Section V-B and V-C). Then we discuss the parameters used and the computational performance (Section V-D). Finally, we conclude with the limitations of the proposed method (Section V-E).

A. Datasets and Evaluation Metrics

To extensively evaluate the algorithm, we use almost all publicly available datasets [9, 11, 16, 30]. Except for *DistSuf* [30] using a static event camera, all the others were captured with a hand-held moving camera, in either indoor or outdoor environments. The number of IMOs was up to

TABLE I: Summary of datasets' characteristics

Dataset	Camera	Camera status	#IMO	Env.	HDR	Non-rigid
EED [9]	DAVIS240	Moving	1-3	Indoor	Yes	No
EV-IMO [16]	DAVIS346	Moving	1-3	Indoor	No	No
DistSuf [30]	DAVIS240	Static	1-3	In/Outdoor	No	No
EMSMC [11]	DAVIS346	Moving	1-3	In/Outdoor	Yes	Yes
Our Data	DAVIS346	Moving	1-3	Indoor	No	Yes

TABLE II: Quantitative evaluation on EED dataset [9] using detection rate of IMOs (in %).

EED sequence	EED [9]	EMSMC [11]	0-MMS [14]	EMSGC [15]	Ours
what is background	89.21	100.0	-*	100.0	100.0
occlusion	90.83	92.31	-	100.0	100.0
fast drone	92.78	96.30	-	96.30	96.30
light variation	84.52	80.51	-	93.51	92.21
multiple objects	87.32	96.77	-	0	95.67
Average	88.93	93.18	94.2	77.96	96.84

* Numbers are not available for each individual sequence in [14].

TABLE III: Quantitative evaluation on EV-IMO dataset [16] using the IoU metric (in%)

Method	EV-IMO [16]	MOMS [13]	0-MMS [14]	EMSGC [15]	Ours
IoU	77.00	74.82	80.37	76.81	80.73

three, and challenging cases including HDR scenes and non-rigid motions were also considered. Details of these datasets are summarized in Table. I. Two standard metrics are used in the quantitative evaluation, including *detection rate* and *Intersection over Union (IoU)*. Details about the metrics can be found in [9, 15].

B. Quantitative Evaluation

EED Dataset [9]. We evaluate our algorithm on all the five sequences using the *detection rate* metric and compared against state-of-the-art approaches [9, 11, 14, 15]. As shown in Table II, our algorithm overall outperforms the other methods. Note that our method gave less accurate result than EMSGC [15] on the *light variation* sequence. The problem was caused by false-positive feature correspondences used as input in the first level. These features were extracted from noisy events induced by a strobe light flashing in the background. Such a bad information-noise ratio generated many false-positive motion model hypotheses, leading to inaccurate segmentation results in the second level. Nevertheless, EMSGC [15] failed to detect small-size IMOs (*Multiple objects* sequence) because of its initialization that struggled in the case of bad inlier-outlier ratio, whereas our method can handle this due to the proposed feature tracking.

EV-IMO Dataset [16]. We also evaluate our algorithm using the *IoU* metric when ground-truth dense segmentation masks are provided. As can be seen from the quantitative results in Table III, our algorithm outperforms the other state-of-the-art solutions. The numbers for baseline methods are still taken from the corresponding publications. Qualitative results are given in Fig 4, where we respectively show the labelled features, the labelled events (IWEs), and the dense segmentation masks (convex hulls) overlapped on corresponding grayscale images. The results demonstrate that objects can be successfully separated from the background as long as they undergo motions independent of the camera’s ego-motion. Note that the feature labeling result from level one is sometimes not perfect. We observed redundant motion model instances, and some weak features from the same IMO that may be associated to different motion models. Both of

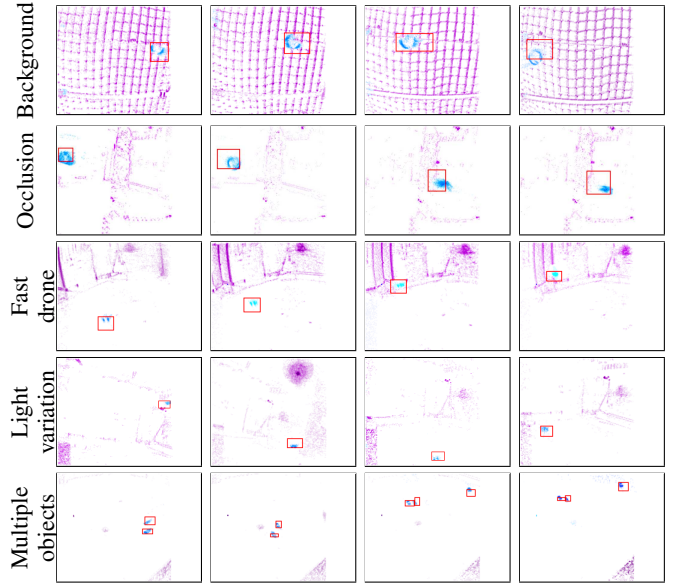


Fig. 3: Segmentation result on EED dataset [9]. The red bounding boxes show the location of the IMOs. Note that all ground truth bounding boxes are manually annotated on the DAVIS grayscale images that are not temporally aligned with event data. Hence, offsets are witnessed for fast moving objects.

these two issues can be resolved in level two thanks to the applied regularization terms (*Potts* model and MDL term), as long as true-positive motion model proposals are obtained from level one.

C. Qualitative Evaluation

In addition to the above quantitative results, we also provide extensive qualitative evaluation on a set of real-world data, including DistSurf dataset [30], EMSMC dataset [11], and our data. Exemplary results are shown in Fig. 5.

The DistSurf dataset consists of two sequences: cars and hands. As shown in the first two columns of Fig. 5, our algorithm successfully separates an IMO from the other, even when they overlap. Results of the two selected sequences from EMSMC dataset are illustrated in the third and fourth column of Fig. 5. The *skateboarder* sequence was captured in an outdoor HDR scene with the event camera looking towards the sun while a pedestrian and a skateboarder pass by. This result demonstrates that event-based solutions are qualified in HDR scenarios. Also, the compact segmentation is due to the MDL term applied in level two, which approximates complex non-rigid motions with a small number of rigid ones. The *fan & coin* sequence captures a rotating fan (rotation speed around 1800 rpm) and a free-falling coin. Our algorithm supports different types of motion models, and thus allows us to model and separate the rotating fan blades and the falling coin. Our data captured indoor scenes with multiple people passing by. It is more complicated compared to the aforementioned ones because multiple non-rigid motions and occlusions coexist. Our algorithm performs well on these sequences.

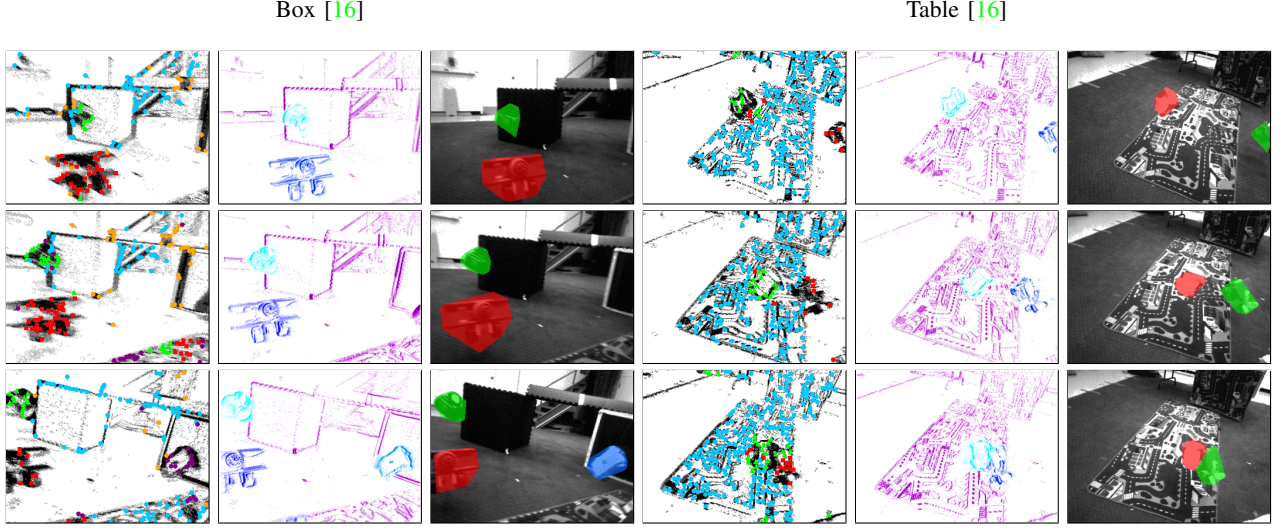


Fig. 4: Segmentation results on the EV-IMO dataset [16]. Column 1 and 4 show the feature labeling result of two sequences: *Box* and *Table*. Features clustered to the same group are labeled with an identical color. Column 2 and 5 show the labeled IWEs. Note that we use a different color scheme for the visualization of IWEs than that for feature labeling. Column 3 and 6 show the corresponding dense segmentation masks on the corresponding grayscale images from DAVIS [31]. The grayscale images are for visualization only.

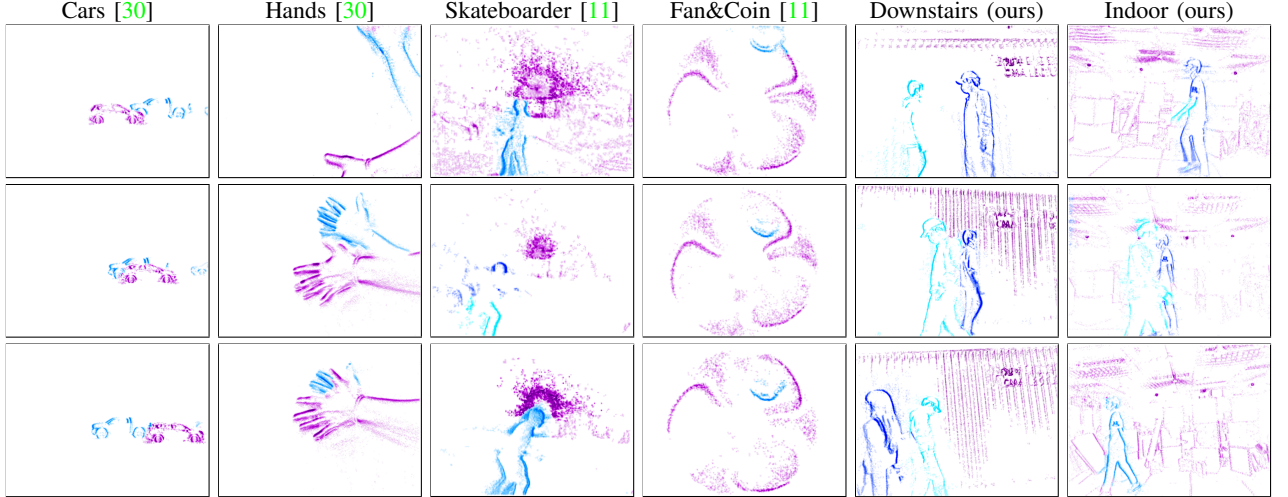


Fig. 5: Qualitative results on a set of real-world data.

D. Parameters and Computational Performance

The segmentation is performed on events occurred within a constant time interval, and we process events in a fashion of sliding window. The width of the temporal window δt is set by 15–35 ms according to the scene dynamics and textures. The confidence threshold μ is set to 0.95. The inlier-outlier (geometric error in Section IV-A) threshold is set to 1.5 pixel. The parameters used at level two are taken from [15].

We ran our algorithm using a single thread on an Intel i7-8700K CPU. The average computation time of each module in Fig. 2 is listed in Table IV. It is worth mentioning that the feature clustering algorithm at level one performs thousands of iterations to propose hypotheses extensively. To speed up, we reuse historical motion model to bootstrap the

estimation at next time, and thus, only data corresponding to new appearing IMOs need to be repeatedly sampled and validated. Numbers in parentheses refer to the time when reusing historical models. With the resulting compact pool of models, the spatio-temporal graph cut method at level two can be finished efficiently. The overall computation time for each independent segmentation takes 0.36 second approximately, which is 20 times faster than [15].

TABLE IV: Computation time of our algorithm

Module	Feature Detection and Tracking	Feature Clustering (Level One)	Event Clustering (Level Two)	Overall
Avg. Time (ms)	3	110 (15)	250	363 (268)

E. Limitations

Two limitations of the proposed method are witnessed during our experiment. First, our method cannot distinguish two IMOs under similar motion. For example, two cars moving from right to left in the *Cars* sequence [30] are likely to be recognized as one IMO no matter how far they are from each other. This issue can be resolved by introducing semantic cues. Second, the width of the temporal window has to be selected according to the characteristics of each data. This is due to the fact that the applied linear motion model is held only within a short time interval, thus a narrow window is preferred. On the other hand, the applied motion-compensation scheme needs enough textures to evaluate the sharpness, thus requiring a relatively wide window. The parameter δt is set by making a trade-off between the two considerations. A better way to resolve this limitation is to apply nonlinear motion models, such as a Bézier spline [32], which is able to describe complex motion within a relatively longer temporal window. We leave these to future work.

VI. CONCLUSION

We present a novel approach for event-based motion segmentation in this paper. Our method is a cascaded two-level multi-model fitting scheme, which addresses the segmentation problem in a two-step fashion. In the first level, we leverage event feature tracking and solve the feature clustering problem via a progressive multi-model fitting pipeline. In the second level, the event clustering problem is formulated as inferring a 3-D MRF and solved using a spatio-temporal graph-cut method. The sequentially cascaded scheme leads to a more accurate and efficient event-based motion segmentation, which cannot be achieved by any of the components alone. The extensive evaluation demonstrates the versatility of the proposed method.

REFERENCES

- [1] P. Lichtsteiner, C. Posch, and T. Delbruck, "A 128×128 120 dB 15 μ s latency asynchronous temporal contrast vision sensor," *IEEE J. Solid-State Circuits*, vol. 43, no. 2, pp. 566–576, 2008.
- [2] P. Rogister, R. Benosman, S.-H. Ieng, P. Lichtsteiner, and T. Delbruck, "Asynchronous event-based binocular stereo matching," *IEEE Trans. Neural Netw. Learn. Syst.*, vol. 23, no. 2, pp. 347–353, 2012.
- [3] Y. Zhou, G. Gallego, H. Rebecq, L. Kneip, H. Li, and D. Scaramuzza, "Semi-dense 3D reconstruction with a stereo event camera," in *Eur. Conf. Comput. Vis. (ECCV)*, 2018, pp. 242–258.
- [4] G. Gallego, J. E. A. Lund, E. Mueggler, H. Rebecq, T. Delbruck, and D. Scaramuzza, "Event-based, 6-DOF camera tracking from photometric depth maps," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 40, no. 10, pp. 2402–2412, Oct. 2018.
- [5] E. Mueggler, B. Huber, and D. Scaramuzza, "Event-based, 6-DOF pose tracking for high-speed maneuvers," in *IEEE/RSJ Int. Conf. Intell. Robot. Syst. (IROS)*, 2014, pp. 2761–2768.
- [6] H. Kim, S. Leutenegger, and A. J. Davison, "Real-time 3D reconstruction and 6-DoF tracking with an event camera," in *Eur. Conf. Comput. Vis. (ECCV)*, 2016, pp. 349–364.
- [7] 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 Robot. Autom. Lett.*, vol. 2, no. 2, pp. 593–600, 2017.
- [8] Y. Zhou, G. Gallego, and S. Shen, "Event-based stereo visual odometry," *IEEE Transactions on Robotics*, 2021.
- [9] A. Mitrokhin, C. Fermüller, C. Parameshwara, and Y. Aloimonos, "Event-based moving object detection and tracking," in *IEEE/RSJ Int. Conf. Intell. Robot. Syst. (IROS)*, 2018.
- [10] T. Stoffregen and L. Kleeman, "Simultaneous optical flow and segmentation (SOFAS) using Dynamic Vision Sensor," in *Australasian Conf. Robot. Autom. (ACRA)*, 2017.
- [11] T. Stoffregen, G. Gallego, T. Drummond, L. Kleeman, and D. Scaramuzza, "Event-based motion segmentation by motion compensation," in *Int. Conf. Comput. Vis. (ICCV)*, 2019, pp. 7243–7252.
- [12] 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 *IEEE Conf. Comput. Vis. Pattern Recog. (CVPR)*, 2018, pp. 3867–3876.
- [13] C. M. Parameshwara, N. J. Sanket, A. Gupta, C. Fermüller, and Y. Aloimonos, "MOMS with Events: Multi-object motion segmentation with monocular event cameras," *arXiv:2006.06158*, 2020.
- [14] C. M. Parameshwara, N. J. Sanket, C. D. Singh, C. Fermüller, and Y. Aloimonos, "0-mms: Zero-shot multi-motion segmentation with a monocular event camera," *arXiv:2006.06158*, 2020.
- [15] Y. Zhou, G. Gallego, X. Lu, S. Liu, and S. Shen, "Event-based motion segmentation with spatio-temporal graph cuts," *arXiv:2012.08730*, 2020.
- [16] A. Mitrokhin, C. Ye, C. Fermüller, Y. Aloimonos, and T. Delbruck, "EV-IMO: Motion segmentation dataset and learning pipeline for event cameras," in *IEEE/RSJ Int. Conf. Intell. Robot. Syst. (IROS)*, 2019.
- [17] J. P. Rodríguez-Gómez, A. G. Eguíluz, J. Martínez-de Dios, and A. Ollero, "Asynchronous event-based clustering and tracking for intrusion monitoring in uas," in *2020 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2020, pp. 8518–8524.
- [18] A. Z. Zhu, N. Atanasov, and K. Daniilidis, "Event-based feature tracking with probabilistic data association," in *IEEE Int. Conf. Robot. Autom. (ICRA)*, 2017, pp. 4465–4470.
- [19] G. Gallego and D. Scaramuzza, "Accurate angular velocity estimation with an event camera," *IEEE Robot. Autom. Lett.*, vol. 2, no. 2, pp. 632–639, 2017.
- [20] M. A. Fischler and R. C. Bolles, "Random sample consensus: a paradigm for model fitting with applications to image analysis and automated cartography," *Commun. ACM*, vol. 24, no. 6, pp. 381–395, 1981.
- [21] G. Gallego, M. Gehrig, and D. Scaramuzza, "Focus is all you need: Loss functions for event-based vision," in *IEEE Conf. Comput. Vis. Pattern Recog. (CVPR)*, 2019, pp. 12 272–12 281.
- [22] J. Shi and C. Tomasi, "Good features to track," in *IEEE Conf. Comput. Vis. Pattern Recog. (CVPR)*, 1994, pp. 593–600.
- [23] B. D. Lucas and T. Kanade, "An iterative image registration technique with an application to stereo vision," in *Int. Joint Conf. Artificial Intell. (IJCAI)*, 1981, pp. 674–679.
- [24] D. Barath and J. Matas, "Progressive-x: Efficient, anytime, multi-model fitting algorithm," in *Int. Conf. Comput. Vis. (ICCV)*, October 2019, pp. 3780–3788.
- [25] H. Isack and Y. Boykov, "Energy-based geometric multi-model fitting," *International journal of computer vision*, vol. 97, no. 2, pp. 123–147, 2012.
- [26] D. Barath and J. Matas, "Graph-cut ransac," in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2018, pp. 6733–6741.
- [27] P. H. Torr, S. J. Nasuto, and J. M. Bishop, "Napsac: High noise, high dimensional robust estimation-it's in the bag," in *British Machine Vision Conference (BMVC)*, 2002.
- [28] O. Chum and J. Matas, "Matching with prosac-progressive sample consensus," in *2005 IEEE computer society conference on computer vision and pattern recognition (CVPR'05)*, vol. 1. IEEE, 2005, pp. 220–226.
- [29] P. H. S. Torr, "Bayesian model estimation and selection for epipolar geometry and generic manifold fitting," *International Journal of Computer Vision*, vol. 50, no. 1, pp. 35–61, 2002.
- [30] M. M. Almatrafi, R. Baldwin, K. Aizawa, and K. Hirakawa, "Distance surface for event-based optical flow," *IEEE Trans. Pattern Anal. Mach. Intell.*, 2020.
- [31] C. Brandli, R. Berner, M. Yang, S.-C. Liu, and T. Delbruck, "A 240×180 130dB 3 μ s latency global shutter spatiotemporal vision sensor," *IEEE J. Solid-State Circuits*, vol. 49, no. 10, pp. 2333–2341, 2014.
- [32] H. Seok and J. Lim, "Robust feature tracking in dvs event stream using bézier mapping," in *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision*, 2020, pp. 1658–1667.