






What's in the Flow? Exploiting Temporal Motion Cues for Unsupervised Generic Event Boundary Detection

Sourabh Vasant Gothe , Vibhav Agarwal , Sourav Ghosh , Jayesh Rajkumar Vachhani ,
Pranay Kashyap, Barath Raj Kandur Raja 

Samsung R&D Institute Bangalore, India

{sourab.gothe, vibhav.a, sourav.ghosh, jay.vachhani, kashyap.p, barathraj.kr} @samsung.com

Abstract

Generic Event Boundary Detection (GEBD) task aims to recognize generic, taxonomy-free boundaries that segment a video into meaningful events. Current methods typically involve a neural model trained on a large volume of data, demanding substantial computational power and storage space. We explore two pivotal questions pertaining to GEBD: Can non-parametric algorithms outperform unsupervised neural methods? Does motion information alone suffice for high performance? This inquiry drives us to algorithmically harness motion cues for identifying generic event boundaries in videos. In this work, we propose FlowGEBD, a non-parametric, unsupervised technique for GEBD. Our approach entails two algorithms utilizing optical flow: (i) Pixel Tracking and (ii) Flow Normalization. By conducting thorough experimentation on the challenging Kinetics-GEBD and TAPOS datasets, our results establish FlowGEBD as the new state-of-the-art (SOTA) among unsupervised methods. FlowGEBD exceeds the neural models on the Kinetics-GEBD dataset by obtaining an F1@0.05 score of 0.713 with an absolute gain of 31.7% compared to the unsupervised baseline and achieves an average F1 score of 0.623 on the TAPOS validation dataset.

1. Introduction

In 2023, video has accounted for 82.5% of all web traffic, making it the most popular form of content on the internet. Increased video consumption has made video understanding a critical task in computer vision, comprising video classification [8, 23, 35], object segmentation [39, 40], action localization [3, 37, 41], and captioning [38, 42], among others. However, the memory requirements of the models, the feasibility of real-time inference, and domain generalization are typical constraints on these solutions.

Current state-of-the-art video models [9, 17, 21, 31, 33] have been mainly focused on building upon a limited set

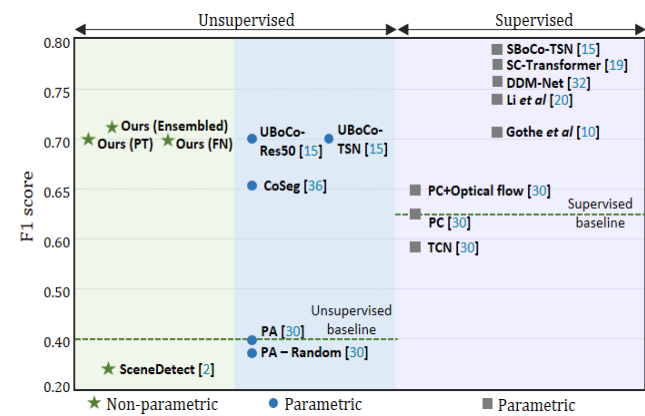


Figure 1. F1@0.05 scores of different methods on the Kinetics-GEBD validation dataset. Our method FlowGEBD achieves state-of-the-art results among unsupervised methods compared to non-parametric [2] and parametric [15, 30, 36] benchmarks.

of predefined action classes and usually process short clips to generate global video-level predictions. With the growth of video content, the number of classes is expanding, and the predefined target classes cannot encompass them all. Recently, the GEBD [30] task was introduced with the objective of studying the long-form video understanding problem through the lens of human perception [34]. GEBD aims to locate class-agnostic event boundaries in a video, regardless of its category. It considers the following high-level causes of event boundaries: changes in subject, action, shot, environment, and object of interaction. The outcome of GEBD has many potential applications: video summarization, video editing, short video segment sharing (Fig. 2), and enhancing video classification and other downstream tasks [38].

There are three main approaches to the GEBD problem: supervised, unsupervised, and self-supervised. Several supervised methods [10, 12, 14, 19, 20, 32] propose to decode the self-similarity representation produced by the frame features for boundary detection. The literature investigates a wide variety of approaches, including efficient represen-

tation learning [10, 14], transformer decoders [19], attention masks [32], and the use of compressed domain features [20]. In unsupervised approaches, CoSeg [36] investigates cognitively inspired parametric methods, and UBoCo [15] yields the best results by employing a novel contrastive loss. In self-supervised approaches, TeG [26] investigates learning temporal granularity in video representations. In contrast, Rai *et al.* [27] employ a differentiable motion feature learning module to detect spatial and temporal differences for GEBD. These methods incorporate explicit motion features into their network structures and earn a high F1 score as they are deep neural networks (DNN) guided by the ground truth. However, we approach GEBD from a different perspective, investigating a two-fold question: (a) Can non-parametric algorithms outperform unsupervised parametric methods? (b) Is motion information alone sufficient to achieve high performance? We seek answers to these questions by exploiting the motion information to detect the generic event boundaries in a video algorithmically.

In this paper, we present two unsupervised, non-parametric approaches to solve GEBD. (i) Pixel Tracking (PT) method that relies on sparse optical flow in temporal dimension to identify the boundary, and (ii) Flow normalization (FN) method that traces the max temporal dense flow to detect the event boundaries. The ensemble of both achieves an F1 score of 0.713 on Kinetics-GEBD and an F1 score of 0.375 on the TAPOS.

As shown in Fig. 1, our method achieves 31.7% absolute gain compared to the unsupervised baseline method and outperforms the supervised baseline [30] by 8.8% on the Kinetics-GEBD dataset. In summary, our main contributions are as follows:

- We propose FlowGEBD, a non-parametric (algorithmic), unsupervised method for generic event boundary detection.
- We design two algorithms by leveraging motion information: (i) Pixel Tracking (PT) and (ii) Flow-Normalization (FN) using optical flow estimation in framewise and patchwise mode to solve the GEBD task.
- We conduct extensive ablations, time complexity analysis, and sensitivity analysis to demonstrate the robustness of the proposed method.
- Our results establish FlowGEBD as the new state-of-the-art among unsupervised methods on the challenging Kinetics-GEBD and TAPOS datasets.

2. Related Work

2.1. Generic Event Boundary Detection

Generic event boundary detection (GEBD) [30] aims to localize the moments where humans naturally perceive

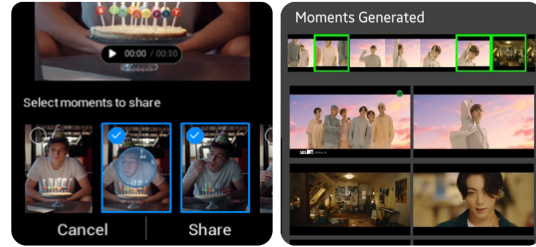


Figure 2. FlowGEBD enables applications on smartphones, like short video segment sharing, summarization, editing by identifying generic video moments

taxonomy-free event boundaries that break a longer video into shorter temporal segments. Previous methods [12, 30, 32] formulate the GEBD task as binary classification, which predicts the boundary label of each frame by considering the temporal context information. However, it could be more efficient because the redundant computation is conducted while generating the representations of consecutive frames. Kang *et al.* [14] proposed to use the temporal self-similarity matrix (TSM) as intermediate representation and used contrastive learning as an auxiliary to learn better from the TSM results. Li *et al.* [20] proposed solving GEBD using the compressed video features and achieved $4.5\times$ faster-running speed than the baseline method [30] on GPU. Recently, Gothe *et al.* [10] developed the most miniature model to solve the GEBD task with the lowest inference time on GPU. SC-Transformer [19] introduced a structured partition of sequences (SPoS) mechanism to learn structured context using a transformer-based architecture for GEBD. To enrich motion information, optical flow is introduced as a new modality in [13]. TeG [26] proposed a generic self-supervised model for learning persistent and more fine-grained features and uses a 3D-ResNet-50 encoder as its backbone. However, all these methods require substantial memory and computational resources along with the labeled data.

Regarding unsupervised GEBD approaches, PySceneDetect [2] is a Python library that detects shot changes by considering pixel changes in the HSV colorspace. However, generic event boundaries consist of various boundary causes like the change of action, subject, and environment, implying that only a tiny portion of event boundaries can be detected with this approach. PredictAbility (PA) [30] computationally assesses the predictability score over time and then locates the event boundaries by detecting the local minima of the predictability sequence. CoSeg [36] devises a transformer-based frame feature reconstruction scheme and adopts ResNet-18 [11] as the backbone. UBoCo [15] proposes an unsupervised/supervised method using the TSM as the video representation. UBoCo’s unsupervised framework for GEBD combines Recursive TSM Parsing (RTP) and the Boundary Contrastive (BoCo) loss. However these models belong to a high memory regime.

2.2. Learning motion and visual correspondences

Motion plays a crucial role in video understanding, and many SOTA models [13, 18, 22, 27, 32] incorporate motion information by using optical flows. Lucas and Kanade’s image registration method [24], also known as gradient-based optical flow, enables motion estimation possible with high-speed computation. Pyramidal Lucas and Kanade [1], Gunnar Farneback [5–7] are other well-known methods for motion estimation.

DDM-Net [32] applies progressive attention to multilevel dense difference maps (DDM) to characterize motion patterns and jointly aggregate motion and appearance cues in a supervised setting. MotionSqueeze (MS) [18] introduces an end-to-end trainable, model-agnostic and lightweight module to extract motion features on the fly for video understanding. However, it requires training via backpropagation and integration with pre-existing video architectures. Rai *et al.* [27] presents a self-supervised model for GEBD by reformulating training objectives at frame-level and clip-level to learn effective video representations using the MS [18] module. However, these are parametric methods that require training on large datasets. To the best of our knowledge, there is no unsupervised and non-parametric (algorithmic) solution with high performance in generic event boundary detection.

3. Proposed Methodology

GEBD takes a video as input and returns a set of boundary timestamps. Mathematically, it maps an ordered sequence of L frames, $\langle f_1, f_2, \dots, f_L \rangle$ (that may also be represented as $\vec{F} \in \mathbf{F}$), to a set of timestamps $\{b_1, b_2, \dots, b_M\}$ ($= \mathcal{B} \in \mathbf{B}$), that denote the event boundaries. It then naturally follows that $M \leq (L - 1)$. For all practical purposes, $M \ll L$ and $\forall b_i \in \mathcal{B}, \exists j$, such that timestamp b_i corresponds to a unique frame f_j . Thus, we formulate the GEBD task as:

$$\mathcal{T} : \mathbf{F} \rightarrow \mathbf{B} \quad (1)$$

Here, we describe our approach, FlowGEBD, that solves this task using pixel tracking, flow normalization, and their ensemble (with temporal refinement) as shown in Fig. 3.

3.1. FlowGEBD with Pixel Tracking (PT)

In this section, we present a method that leverages sparse optical flow to determine event boundaries by monitoring the flow of a subset of pixels.

3.1.1 Framewise mode

We process a video frame-by-frame, considering each frame as a unit. Each frame f of width w and height h comprises a 2-dimensional matrix of pixels, $p_{u,v}$, where $u, v \in \mathbb{Z}^+$ (positive integers), $u \in [1, w]$, and $v \in [1, h]$. We only consider the luminance information of pixels. Hence, $p_{u,v}$ can be represented as a real number ($p_{u,v} \in \mathbb{R}$), $0 \leq p_{u,v} \leq 1$.

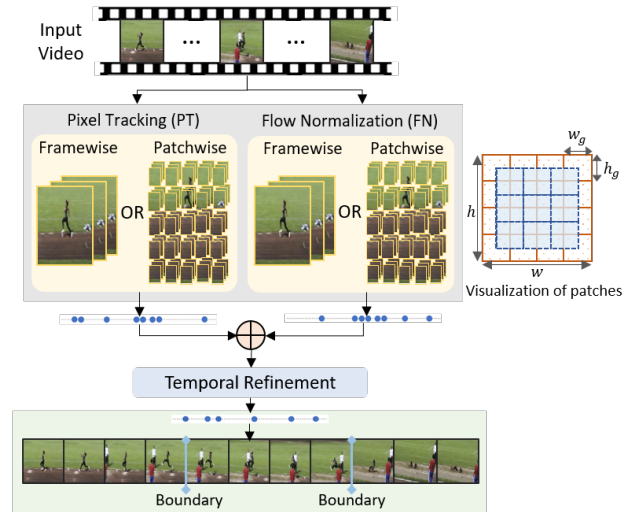


Figure 3. FlowGEBD accepts a video as input and predicts a set of event boundaries, \mathcal{B} . Visual representation of patches with $n_w = n_h = 4$ (right). \square : Base patches, \square : Centroidal

The apparent motion of pixel $p_{u,v}$ between two consecutive frames caused by the movement of an object or camera is measured by optical flow. For each frame f_i with a subsequent frame f_{i+1} , the optical flow Φ_i can be denoted as a 2-dimensional matrix of displacement vectors [31]. Each element in the displacement vector $\vec{d}_{u,v}$ denotes the horizontal and vertical motion of $p_{u,v}$ between frames f_i and f_{i+1} .

Method. The key intuition is that an event boundary can be determined by monitoring the optical flow of a subset of pixels. This underlying assumption is supported by Shou *et al.* [30] who consider change in brightness, rapid camera movements, etc. as definitive indicators of event boundaries. So, for the first frame f_1 in the sequence \vec{F} , we use uniform random sampling or Shi-Tomasi corner detection algorithm [29] to identify a set $\mathcal{P}_{\text{base}}$ comprised of key features (pixels $p_{u,v}$). Then, for every subsequent frame f_i , we compute the sparse flow for these pixels using the iterative Lucas-Kanade method [24]. We consider a pixel as also an element of the current key pixel set $\mathcal{P}_{\text{current}} (= \mathcal{P}_i)$, if and only if it has non-zero displacement $\vec{d}_{u,v}$ from the previous frame.

In each of these frame-by-frame iterations, whenever the ratio of elements in $\mathcal{P}_{\text{current}}$ to $\mathcal{P}_{\text{base}}$ falls below a predefined threshold θ_1 , we infer to have encountered an event boundary and record the current frame index. In such a scenario, we resample new key pixels $\mathcal{P}_{\text{base}}$ from the current frame. If no such event boundary is encountered in an iteration, we maintain $\mathcal{P}_{\text{base}}$ as a constant reference until a boundary is identified. To determine the final set of boundaries, we apply a temporal boundary refinement algorithm (Section 3.3). At any time, this algorithm depends only on the past frame to deduce an event boundary, thereby exhibiting the causal property. The overall approach is detailed in Algorithm 1.

Algorithm 1: FlowGEBD using Pixel Tracking
(Framewise mode)

Data: Video of resolution $w \times h$ as a sequence of L frames, $\vec{F} = \langle f_1, f_2, \dots, f_L \rangle$

Result: Event Boundaries, $\mathcal{B} = \{b_1, b_2, \dots, b_M\}$

$\mathcal{B}_{\text{temp}} \leftarrow \{\}$

$\mathcal{P}_1 \leftarrow \mathcal{P}_{\text{base}} \leftarrow \text{samplePixels}(f_1)$

// $O(1)$ for uniform random

for $f_i \in (\vec{F} - f_1)$ **do**

$\Phi_{i-1} \leftarrow \text{sparseFlow}(f_{i-1}, f_i, \mathcal{P}_{i-1})$

// $O(wh |\mathcal{P}_{\text{base}}|)$

$\mathcal{P}_i \leftarrow \bigcup_{\substack{|\vec{d}_{u,v}| \neq 0 \\ \vec{d}_{u,v} \in \Phi_{i-1}}} \{p_{u,v}\}$

// Non-zero flow, $O(wh)$

if $\frac{|\mathcal{P}_i|}{|\mathcal{P}_{\text{base}}|} < \theta_1$ **then**

$\mathcal{B}_{\text{temp}} \leftarrow \mathcal{B}_{\text{temp}} \cup \{i\}$

$\mathcal{P}_i \leftarrow \mathcal{P}_{\text{base}} \leftarrow \text{samplePixels}(f_i)$

end

end

// for-loop $\implies O(Lwh |\mathcal{P}_{\text{base}}|)$

$\mathcal{B} \leftarrow \text{refine}(\mathcal{B}_{\text{temp}})$

// Refer to Algorithm 3

return \mathcal{B}

samplePixels(\cdot): Uniform random or Shi-Tomasi corner detection
sparseFlow(\cdot): Sparse optical flow using iterative Lucas-Kanade method
 θ_1 : A constant threshold

Can we improve further? The framewise approach monitors the key pixels in a frame. However, for a video of a moderately large field of view, it is often the case that an event boundary may be denoted by the change in actions of certain subjects in the video, even if the background remains static. Furthermore, the main subjects of the frame are typically positioned along the grid lines and at the intersections to make it more aesthetically pleasing following the Rule of thirds [4]. For such cases, the event boundary can be determined more accurately if we decompose a frame into a grid of patches instead of having the entire frame as a single unit. Patchwise processing provides advantages like capturing subject change in small areas and detecting action change from one patch to another. Thus, in the next section, we propose an approach that processes each frame as a composition of multiple patches.

3.1.2 Patchwise mode

A patch g_f , derived from frame f , consists of a subset of the frame pixels. More specifically, patch $g_f(u, v, w_g, h_g)$ consists of all pixels, $p_{i,j} \in f$, where $i, j \in \mathbb{Z}^+$, $i \in [u, u + w_g)$ and $j \in [v, v + h_g)$. We denote the set of all such patches in frame f as \mathbf{G}_f .

We define two categories of patches where patches of the same category do not overlap each other. The first among these are “*base patches*”, which distribute the pixels equally and independently along the width and height of the frame. We refer to the other category as “*centroidal patches*”, as each of their edges joins the centroids of adjacent base patches. Centroidal patches help capture events that span across the intersections of base patches. Fig. 3 (right) depicts the arrangement of patches, where the total number of patches (each of width w_g and height h_g) is given by:

$$\mathcal{N}_g = \underbrace{n_w \times n_h}_{\text{Base patches}} + \underbrace{(n_w - 1) \times (n_h - 1)}_{\text{Centroidal patches}} \quad (2)$$

Where n_w and n_h represent the cardinality of base patches along the frame width and height, respectively.

Method. In this mode, we independently process the entire frame sequence \mathcal{N}_g times, once for each patch, considering only the corresponding patches. During this, we skip the temporal boundary refinement stage mentioned in Algorithm 1. We then take a union of the \mathcal{N}_g predicted boundary sets and apply boundary refinement to derive \mathcal{B} . It may be noted that framewise mode is a specialized case of patchwise, where $n_w = n_h = 1 \implies \mathcal{N}_g = 1$.

3.2. FlowGEBD with Optical Flow Normalization

We explore another approach that leverages dense optical flow and determines event boundaries by observing the normalized optical flow.

3.2.1 Framewise mode

As we have observed in FlowGEBD with pixel tracking (section 3.1.2), the framewise mode is a specialized case of patchwise. Hence, we discuss the Optical Flow Normalization method in the more generalized patchwise mode, and the same can be adapted for framewise by using $n_w = n_h = 1$.

3.2.2 Patchwise mode

Method. For every frame f_i , we identify the set of \mathcal{N}_g patches, denoted by \mathbf{G}_{f_i} , for a fixed n_w and n_h . Then, for every consecutive frame f_{i-1} and f_i , we compute \mathcal{N}_g dense optical flows (one between each patch pair, $g_{f_{i-1}}$ and g_{f_i}). For this, we use the Gunnar Farneback algorithm [7]. Then, we use the maximum flow displacement corresponding to all patch pixels as the “*flow of the patch*” or its “*PatchFlow*”. After processing L frames and their \mathcal{N}_g patches, we accumulate the PatchFlow for each patch across temporal dimension and normalize the displacement values. We hypothesize that considerable displacement of PatchFlow in the temporal dimension constitutes a change in action or event. For any patch, if the normalized value for frame index i exceeds a constant threshold θ_2 , we deem to have encountered an event boundary and add the corresponding frame index i to our

Algorithm 2: FlowGEBD using Optical Flow Normalization (Patchwise mode)

Data: Video of resolution $w \times h$ as a sequence of L frames, $\vec{F} = \langle f_1, f_2, \dots, f_L \rangle$; Parameters: Patch width w_g , and height h_g

Result: Event Boundaries, $\mathcal{B} = \{b_1, b_2, \dots, b_M\}$

```

G $f_1$   $\leftarrow$  patches( $f_1, w_g, h_g$ ) // Gets  $\mathcal{N}_g$  patches (Fig. 3)
for  $f_i \in (\vec{F} - f_1)$  do
  G $f_i$   $\leftarrow$  patches( $f_i, w_g, h_g$ )
   $\Phi_{i-1}^g \leftarrow \{\}$  // Placeholder for all optical flows in frame  $f_i$ 
  for  $(g_{f_{i-1}}, g_{f_i}) \in (\mathbf{G}_{f_{i-1}}, \mathbf{G}_{f_i} |_{\text{corresponding patches}})$  do
     $\varphi \leftarrow$  denseFlow( $g_{f_{i-1}}, g_{f_i}$ ) //  $O(w_g h_g)$ 
     $\Phi_{i-1}^g \leftarrow \Phi_{i-1}^g \cup \{\max(\varphi)\}$ 
  end
end // inner-for-loop  $\implies O(\mathcal{N}_g w_g h_g)$ 
end // outer-for-loop  $\implies O(L \mathcal{N}_g w_g h_g)$ 
// We now have PatchFlows =  $\{\Phi_1^g, \Phi_2^g, \dots, \Phi_{L-1}^g\}, \dots$ 
// ... where  $i^{\text{th}}$  element = set  $\Phi_i^g$  of patch optical flows for  $f_i$ 
 $\mathcal{B}_{\text{temp}} \leftarrow \{\}$ 
for  $(\Phi_1^g, \Phi_2^g, \dots, \Phi_{L-1}^g) \in \text{PatchFlows} |_{\text{corresponding patches}}$  do
   $\vec{\Phi} \leftarrow (\Phi_1^g, \Phi_2^g, \dots, \Phi_{L-1}^g)$ 
   $\hat{\Phi} \leftarrow \frac{\vec{\Phi}}{\|\vec{\Phi}\|_2}$  // L2-norm
   $\mathcal{B}_{\text{temp}} \leftarrow \mathcal{B}_{\text{temp}} \cup \{\arg_{b_i} \hat{\Phi}_i > \theta_2 (\hat{\Phi})\}$ 
end // for-loop  $\implies O(L \mathcal{N}_g)$ 
 $\mathcal{B} \leftarrow$  refine( $\mathcal{B}_{\text{temp}}$ ) // Refer to Algorithm 3
return  $\mathcal{B}$ 

```

denseFlow(\cdot): Dense optical flow using the Gunnar Farneback's algorithm
 θ_2 : A constant threshold

working set of event boundaries $\mathcal{B}_{\text{temp}}$. Finally, we apply temporal refinement on $\mathcal{B}_{\text{temp}}$ to compute the refined event boundary set \mathcal{B} . This approach for patchwise processing of dense optical flow is detailed in Algorithm 2.

3.3. FlowGEBD with ensembling of Pixel Tracking and Flow Normalization

Pixel tracking helps determine event boundaries based on the sparse optical flow of a few key pixels. On the other hand, flow normalization aggregates the dense optical flow of all pixels and offers a lossless method to determine a set of event boundaries using ‘‘PatchFlow’’. To obtain an ensemble of both approaches, we independently take the event boundaries from both without performing the temporal refinement stage. Instead, we take a union of the predicted sets from these two approaches and perform temporal refinement over the union. **Temporal refinement.** We analyze the elements of a set of predicted boundary timestamps along the corresponding

Algorithm 3: Temporal Refinement of Boundaries

Data: Event Boundaries, $\mathcal{B} = \{b_1, b_2, \dots, b_M\}$

Result: Refined Event Boundaries, $\tilde{\mathcal{B}} \subset \mathcal{B}$

```

 $\tilde{\mathcal{B}}_{\text{temp}} \leftarrow \{\}$ ;  $\vec{\mathcal{B}} \leftarrow$  sorted( $\mathcal{B}$ ) //  $O(M \log M)$ 
 $\vec{K} \leftarrow \langle \rangle$  // Placeholder for a cluster of boundary elements
for  $b'_i \in \vec{\mathcal{B}}$  do
  if  $(\vec{K} \neq \langle \rangle) \wedge (\forall k \in \vec{K} \mid (|b'_i - k| \geq \theta_3))$ 
    then
       $\tilde{\mathcal{B}}_{\text{temp}} \leftarrow \tilde{\mathcal{B}}_{\text{temp}} \cup \{\text{median}(\vec{K})\}$ 
       $\vec{K} \leftarrow \langle \rangle$ 
    end
     $\vec{K} \leftarrow \vec{K} \cup \langle b'_i \rangle$ 
  end
 $\tilde{\mathcal{B}}_{\text{temp}} \leftarrow \tilde{\mathcal{B}}_{\text{temp}} \cup \{\text{median}(\vec{K})\}$  // Flush  $\vec{K}$ 
return  $\tilde{\mathcal{B}}_{\text{temp}}$  as  $\tilde{\mathcal{B}}$ 

```

θ_3 : A constant threshold

temporal dimension to identify ‘‘rare boundaries’’ and ‘‘popular boundaries’’. Rare or isolated boundaries are those where the event changes in a single patch and typically with no neighboring timestamps, i.e., for a rare boundary b_r , there exists no other boundary within the duration $b_r \pm \theta_3$. On the other hand, popular boundaries are dense clusters where multiple boundaries have been determined within the temporal vicinity.

We may interpret contiguous popular boundaries as belonging to one cluster and each rare boundary as a standalone single-element cluster. Then, we select one representative element for each cluster by identifying its median boundary element and consider only such elements for the final set of refined event boundaries. The generic Algorithm 3 identifies such clusters and determines an optimal boundary in each of them.

4. Experiments

In this section, we conduct multiple experiments and evaluate both algorithms, followed by the ensembled method.

4.1. Dataset

Kinetics-GEBD. Our approach is evaluated primarily on the challenging Kinetics-GEBD [30], a benchmark dataset for locating the boundaries of generic events in the video. It consists of 54,691 videos of 10 seconds each that span a broad spectrum of video domains in the wild and is open-vocabulary, taxonomy-free. The ratio of the train, validation, and test sets in Kinetics-GEBD is equal, with each set including roughly 18,000 videos chosen from Kinetics-400 [16]. FlowGEBD is an algorithmic unsupervised method, so we do not require train data. We evaluate our methods on the val-

Supervision	Rel.Dis. threshold	0.05	0.1	0.15	0.2	0.25	0.30	0.35	0.4	0.45	0.5	Avg
Supervised	PC [30] (Baseline)	0.625	0.758	0.804	0.829	0.844	0.853	0.859	0.864	0.867	0.870	0.817
	PC + Optical Flow [20]	0.646	0.776	0.818	0.842	0.856	0.864	0.868	0.874	0.877	0.879	0.830
	Gothe <i>et al.</i> [10]	0.712	-	-	-	-	-	-	-	-	-	-
	SBoCo-Res50 [15]	0.732	-	-	-	-	-	-	-	-	-	0.866
	DDM-Net [32]	0.764	0.843	0.866	0.880	0.887	0.892	0.895	0.898	0.900	0.902	0.873
	Li <i>et al.</i> [20]	0.743	0.830	0.857	0.872	0.880	0.886	0.890	0.893	0.896	0.898	0.865
	SC-Transformer [19]	0.777	0.849	0.873	0.886	0.895	0.900	0.904	0.907	0.909	0.911	0.881
	SBoCo-TSN [15]	0.787	-	-	-	-	-	-	-	-	-	-
Unsupervised	PA - Random [30] [†]	0.336	0.435	0.484	0.512	0.529	0.541	0.548	0.554	0.558	0.561	0.506
	PA [30] [†]	0.396	0.488	0.520	0.534	0.544	0.550	0.555	0.558	0.561	0.564	0.527
	CoSeg [36] [†]	0.656	0.758	0.783	0.794	0.799	0.803	0.804	0.806	0.807	0.809	0.782
	UBoCo-Res50 [15] [†]	0.703	0.839	0.862	0.885	0.889	0.893	0.894	0.898	0.900	0.902	0.866
	UBoCo-TSN [15] [†]	0.702	0.846	0.862	0.879	0.888	0.889	0.895	0.897	0.904	0.905	0.866
	SceneDetect [2] [*]	0.275	0.300	0.312	0.319	0.324	0.327	0.330	0.332	0.334	0.335	0.318
	Ours (PT 1) [*]	0.702	0.819	0.844	0.855	0.860	0.863	0.866	0.867	0.869	0.870	0.841
	Ours (FN 2) [*]	0.691	0.826	0.860	0.877	0.885	0.889	0.892	0.894	0.896	0.897	0.861
	Ours (Ensembled) [*]	0.713	0.828	0.850	0.858	0.862	0.864	0.866	0.867	0.868	0.869	0.845

Table 1. F1 results on Kinetics-GEGB validation set with different Rel.Dis. thresholds. FlowGEGB achieves the best F1@0.05 scores for unsupervised setting (31.7% absolute gain over unsupervised baseline, PA [30]). † : parametric (neural) Methods * : Non-parametric

Supervision	Rel.Dis. threshold	0.05	0.1	0.15	0.2	0.25	0.30	0.35	0.4	0.45	0.5	Avg
Supervised	PC [30]	0.522	0.595	0.628	0.647	0.660	0.666	0.672	0.676	0.680	0.684	0.643
	DDM-Net [32]	0.604	0.681	0.715	0.735	0.747	0.753	0.757	0.760	0.763	0.767	0.728
	SC-Transformer [19]	0.618	0.694	0.728	0.749	0.761	0.767	0.771	0.774	0.777	0.780	0.742
Unsupervised	PA - Random [30] [†]	0.158	0.233	0.273	0.310	0.331	0.347	0.357	0.369	0.376	0.384	0.314
	PA [30] [†]	0.360	0.459	0.507	0.543	0.567	0.579	0.592	0.601	0.609	0.615	0.543
	SceneDetect [2] [*]	0.035	0.045	0.047	0.051	0.053	0.054	0.055	0.056	0.057	0.058	0.051
	Ours (PT 1) [*]	0.355	0.489	0.562	0.619	0.655	0.677	0.693	0.703	0.714	0.721	0.619
	Ours (FN 2) [*]	0.346	0.487	0.562	0.619	0.658	0.678	0.695	0.706	0.715	0.722	0.618
	Ours (Ensembled) [*]	0.375	0.502	0.569	0.624	0.658	0.677	0.695	0.703	0.711	0.717	0.623

Table 2. F1 results on TAPOS validation set with different Rel.Dis. thresholds. The ensembled method achieves the best F1 score compared to other unsupervised methods. † : parametric (neural) Methods * : Non-parametric Methods

validation set, as the annotations of the test sets are not public.

TAPOS. In addition to Kinetics-GEGB, we experiment on the TAPOS dataset [28] containing Olympics sports videos with 21 actions. The dataset authors manually defined how to break each action into sub-actions during annotation. Following [30], we re-purpose TAPOS for the GEGB task by performing boundaries localization between sub-actions in each action instance. TAPOS contains 1790 instances for validation, and we evaluate on the same.

Implementation and Evaluation. We run all our experiments on Intel(R) Xeon(R) Silver 4210 CPU @ 2.20GHz equipped machine. We sample the video at 4 FPS and resize it to 160×160 as preprocessing. As described in [30], we use F1 at 0.05 Relative Distance (Rel.Dis.) as our primary evaluation metric. The predicted boundary is deemed accurate for a certain Rel.Dis. threshold if the difference between the predicted and ground truth timestamps is smaller than the threshold. We report F1 scores of different thresholds to range from 0.05 to 0.5 with a gap of 0.05.

4.2. Main Results

Kinetics-GEGB. Table 1 illustrates the results of our methods on the Kinetics-GEGB validation set along with unsupervised and supervised benchmarks. PT achieves a higher

F1@0.05 of 0.702 than FN, with a strong recall of 0.91 (average) and a 0.77 (average) precision. Intuitively, PT is able to detect action change across patches along with subject change as the cardinality of the patch increases. FN obtains a high Avg. F1 value through a balanced precision and recall of 0.81 and 0.90, respectively. The FlowGEGB (Ensembled) outperforms all previous unsupervised methods with the highest F1@0.05 of 0.713 using a refinement approach (Algorithm 3) that combines PT and FN. Compared to unsupervised baseline PA [30], FlowGEGB obtains a significant gain of 31.7% in F1@0.05 and exceeds DNN-based unsupervised methods [15, 36], demonstrating the effectiveness of our proposed algorithms. Additionally, compared to the supervised baselines PC [30] and PC + Optical Flow [20], our method achieves 8.8% and 6.7% absolute improvement, respectively.

TAPOS. We also conduct experiments on the TAPOS [28]; the results are summarized in Table 2. The dataset is not inherently well-suited for GEGB as it comprises a pre-defined set of 21 action classes. Hence, we separate sub-action instances from each action video and treat them as a single video for GEGB. Shou *et al.* [30] have shown that the GEGB model trained on TAPOS underperforms on the Kinetics-GEGB dataset due to a change in boundary semantics. However, our algorithm is robust enough to be




Figure 4. *Pixel Tracking*: Visual representation of 3×3 patchwise pixel tracking along temporal dimension ($\theta_1 = 0.4$)




Figure 5. *Flow Normalization*: Visual representation of normalized 3×3 patchwise max flow along temporal dimension ($\theta_2 = 0.25$)

applied directly to the TAPOS dataset. Compared to the unsupervised benchmark PA [30], our method obtains Avg. F1 score of 0.623, gaining 8% absolute improvement. We found no alternative SOTA unsupervised methods for GEBD on the TAPOS dataset to compare our results directly.

Tuning Thresholds θ_1 and θ_2 serve as thresholds that govern the behavior of the PT and FN algorithms, respectively. Additionally, θ_3 is a threshold regulating neighboring boundaries for clustering in the refinement process.

As per the hypothesis, in PT, a notable decrease in pixel count during temporal tracking indicates pivotal event changes, such as changes in subjects or environment. Our empirical findings indicate that marking an event boundary with a drop exceeding 60% ($\theta_1 = 0.4$) yields better performance. In Fig. 4, it can be observed that the event boundaries are perfectly aligned with the trough. Likewise, in FN, our empirical observations indicate that marking a frame as an event boundary is effective when it contributes to over 25% ($\theta_2 = 0.25$) of the overall normalized motion, signifying event changes in the video. Fig. 5 illustrates the alignment of peaks with event boundaries, visually validating our approach. The third threshold θ_3 indicates the distance between two boundaries to consider them as belonging to the same cluster during refinement. Specifically, it is the Euclidean distance to split observations into clusters. We set $\theta_3 = 0.50$, i.e., twice the unit timestamp (time per frame) considered in our algorithm ($4 \text{ FPS} \implies 1 \text{ unit} = 1/4 = 0.25$).

Method	Sampling		Spatial Processing			F1@0.05
	Random	Corners	Frame-wise	Base Patch	Centroidal	
PT-1	✓		✓			0.492
		✓		✓		0.533
	✓	✓		✓		0.659
				✓	✓	0.678
FN-2	✓	✓		✓		0.678
			✓	✓		0.702
	NA	NA		✓	✓	0.486
Ensembled	✓			✓	✓	0.713

Table 3. Effect of Sampling and Spatial Processing on Pixel Tracking (PT) and Flow-Normalized (FN) Algorithms

$\mathcal{N}_g \leftarrow (n_w = n_h)$	F1@0.05		
	PT-1	FN-2	Ensembled
$n_w = 3$	0.679	0.652	0.709
$n_w = 4$	0.696	0.694	0.710
$n_w = 5$	0.702	0.691	0.713

Table 4. F1 score of the proposed method with respect to patch size on GEBD-Kinetics validation set

All the experiments reported in Tables 1 and 2 are conducted in patchwise mode with $n_w = n_h = 5$, $\theta_1 = 0.4$, $\theta_2 = 0.25$, $\theta_3 = 0.5$. However, it is demonstrated in Section 4.4 that FlowGEBD is robust and insensitive to these thresholds. Qualitative results of FlowGEBD are presented in supplementary.

4.3. Ablation Studies

Effect of Sampling. Table 3 illustrates the result of two sampling techniques. In random sampling, we uniformly sample the fixed fraction of pixels from each patch. The corner detection [29] looks for a significant change in pixel intensity in all directions. This sometimes results in the sampling of fewer corner pixels. Thus, we observe that random sampling of pixels gives better F1 scores for PT. It may be noted that the sampling method does not apply to the FN since it computes dense optical flow.

Effect of Spatial Granularity. As detailed in Section 3.1.1 and 3.1.2, we conclude from Table 3 that computing patchwise offers higher performance than processing the entire frame as a single unit. Further, by introducing Centroidal patches, we can capture event change at the intersection of the base patches, leading to a noticeable increase in F1.

Effect of Patch Size. Besides the spatial granularity, patch size is essential to predict the accurate event boundaries, as illustrated in Fig. 3. We capture the effect of varying n_w in Table 4. A higher \mathcal{N}_g (indicative of a smaller patch size) results in more candidate boundary sets, reducing the likelihood of missing uncommon boundaries. Moreover, it helps effectively trace events in tiny regions. We have determined the ideal value for our approach to be $n_w = n_h = 5$. Processing beyond a specific patch size can introduce noisy boundaries to the candidate pool, lowering the F1.

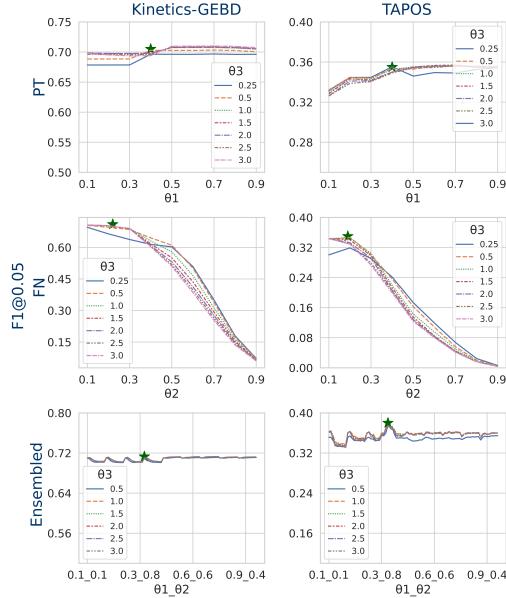


Figure 6. Sensitivity analysis of thresholds $\theta_1, \theta_2, \theta_3$. \star marks the best F1@0.05 score.

4.4. Sensitivity Analysis of thresholds

We conduct an extensive ablation study of thresholds and analyze the impact on the performance. We sample θ_1, θ_2 uniformly between 0.1 to 0.9, and for θ_3 , we vary it from 0.5 to 3.0 in steps of 0.5.

Fig. 6 shows the analysis of sensitivity of the thresholds on Kinetics-GEBD and TAPOS datasets in patchwise ($n_w = 5$) mode. Our findings reveal that PT demonstrates robust performance across a wide range of threshold values (θ_1 and θ_3), consistently exhibiting the same trend for both the Kinetics and TAPOS datasets. In FN, as we increase the normalized flow threshold θ_2 , the number of detected boundaries will reduce gradually. The same effect is observed in Fig. 6, where the gradual change in performance indicates relative stability and generalization on both datasets.

The ensembled method shows the harmonious collaboration of PT and FN, attaining optimal F1 scores across all combinations ($9 \times 9 \times 6$). The mean standard deviations of F1@0.05 for PT, FN, and Ensembled on Kinetics-GEBD are 0.005, 0.02, and 0.0006, respectively, while on TAPOS, they are 0.002, 0.01, and 0.01. These findings highlight the robustness of FlowGEBD and its insensitivity to thresholds.

4.5. Time Complexity of FlowGEBD

We theoretically assess the time complexity of FlowGEBD. From Algorithms 1, 2, 3, we observe that the time complexities (in patchwise mode) of pixel tracking, flow normalization, and temporal refinement are $O(LN_g w_g h_g |\mathcal{P}_{\text{base}}|)$, $O(LN_g w_g h_g)$, and $O(M \log M) (\equiv O(L \log L))$, respectively. So, the overall time complexity is given by

	Method	Params	Latency (ms)	F1@0.05
On GPU	PC [30]	23.5	46.4	0.625
	Gothé <i>et al.</i> [10]	6.79	1.2	0.712
	Li <i>et al.</i> [20]	≥ 34.6	4.7	0.743
	PA [30]	23.5	46.4	0.396
	UBoCo-Res50 [15]	≥ 23.5	≥ 46.4	0.703
	UBoCo-TSN [15]	≥ 90	≥ 90.2	0.702
On CPU	Gothé <i>et al.</i> [10]	6.79	84.35	0.712
	SceneDetect [2]	NA	34.91	0.275
	Ours (PT 1)	NA	2.26	0.702
	Ours (FN 2)	NA	6.42	0.691
	Ours (Ensembled)	NA	6.5	0.713

Table 5. Comparison of Latency with other methods and their F1@0.05 on Kinetics-GEBD validation set.

$O(LN_g w_g h_g |\mathcal{P}_{\text{base}}| + LN_g w_g h_g + L \log L)$. We can further simplify this to $O(LN_g w_g h_g |\mathcal{P}_{\text{base}}|)$. Since in pixel tracking, we use a sparse set of key pixels (i.e. $|\mathcal{P}_{\text{base}}| \ll w_g h_g$), $|\mathcal{P}_{\text{base}}|$ is a fraction of w and h . Additionally, $N_g w_g h_g \propto wh$. So, we infer that the latency of FlowGEBD is directly proportional to wh and L . Please consult the supplementary materials for the analysis of the inference time on sample videos.

Comparison of latency. Table 5 presents the latency per frame across different methods. Most of these methods employ ResNet-50 as their backbone, resulting in an average inference time of at least 46.4 ms per frame at a resolution of 160×160 on a GPU [10]. In contrast, PT and FN exhibit considerably lower inference time, consuming 2.26 ms and 6.42 ms, respectively. The ensembled approach takes 6.5 ms on average without compromising the F1 score. The reported inference time is measured on a Samsung Galaxy S21 Ultra device with 12 GB RAM. Furthermore, the estimation of optical flow can be accelerated on GPU by utilizing NVIDIA Optical Flow SDK [25].

5. Conclusion and Discussion

We introduce FlowGEBD, a non-parametric, unsupervised approach for generic event boundary detection. FlowGEBD comprises two independent algorithms, (i) Pixel Tracking and (ii) Flow Normalization, which can be deployed framewise or patchwise. FlowGEBD achieves the state-of-the-art results (Tables 1 and 2) on the Kinetics-GEBD and TAPOS at a strict relative distance (F1@0.05). This demonstrates that the motion information acquired from an optical flow alone is sufficient and obviates the need for complex neural models to achieve high performance. We performed an extensive ablation study and threshold sensitivity analysis to demonstrate the robustness of the proposed method.

However, since FlowGEBD does not incorporate spatial semantics (high-level DNN features), it is more suitable for GEBD rather than specific action/event localization. The same effect is observed in the evaluation of TAPOS. In future work, we will explore the bi-directional processing of each frame to improve the performance.

References

- [1] Jean-Yves Bouguet et al. Pyramidal implementation of the affine lucas kanade feature tracker description of the algorithm. *Intel corporation*, 5(1-10):4, 2001. 3
- [2] Brandon Castellano. Pyscenedetect: Intelligent scene cut detection and video splitting tool, 2018. 1, 2, 6, 8
- [3] Yu-Wei Chao, Sudheendra Vijayanarasimhan, Bryan Seybold, David A Ross, Jia Deng, and Rahul Sukthankar. Rethinking the faster r-cnn architecture for temporal action localization. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 1130–1139, 2018. 1
- [4] Sagnik Dhar, Vicente Ordonez, and Tamara L Berg. High level describable attributes for predicting aesthetics and interestingness. In *CVPR 2011*, pages 1657–1664. IEEE, 2011. 4
- [5] Gunnar Farneback. Fast and accurate motion estimation using orientation tensors and parametric motion models. In *Proceedings 15th International Conference on Pattern Recognition. ICPR-2000*, volume 1, pages 135–139. IEEE, 2000. 3
- [6] Gunnar Farneback. *Polynomial expansion for orientation and motion estimation*. PhD thesis, Linköping University Electronic Press, 2002. 3
- [7] Gunnar Farneback. Two-frame motion estimation based on polynomial expansion. In *Image Analysis: 13th Scandinavian Conference, SCIA 2003 Halmstad, Sweden, June 29–July 2, 2003 Proceedings 13*, pages 363–370. Springer, 2003. 3, 4
- [8] Christoph Feichtenhofer. X3d: Expanding architectures for efficient video recognition. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 203–213, 2020. 1
- [9] Christoph Feichtenhofer, Haoqi Fan, Jitendra Malik, and Kaiming He. Slowfast networks for video recognition. In *Proceedings of the IEEE/CVF international conference on computer vision*, pages 6202–6211, 2019. 1
- [10] Sourabh Vasant Gothe, Jayesh Rajkumar Vachhani, Rishabh Khurana, and Pranay Kashyap. Self-similarity is all you need for fast and light-weight generic event boundary detection. In *ICASSP 2023-2023 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 1–5. IEEE, 2023. 1, 2, 6, 8
- [11] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 770–778, 2016. 2
- [12] Dexiang Hong, Congcong Li, Longyin Wen, Xinyao Wang, and Libo Zhang. Generic event boundary detection challenge at cvpr 2021 technical report: Cascaded temporal attention network (castanet). *arXiv preprint arXiv:2107.00239*, 2021. 1, 2
- [13] Dexiang Hong, Xiaoqi Ma, Xinyao Wang, Congcong Li, Yufei Wang, and Longyin Wen. Sc-transformer++: Structured context transformer for generic event boundary detection, 2022. 2, 3
- [14] Hyolim Kang, Jinwoo Kim, Kyungmin Kim, Taehyun Kim, and Seon Joo Kim. Winning the cvpr’2021 kinetics-gebdc challenge: Contrastive learning approach. *arXiv preprint arXiv:2106.11549*, 2021. 1, 2
- [15] Hyolim Kang, Jinwoo Kim, Taehyun Kim, and Seon Joo Kim. Uboco: Unsupervised boundary contrastive learning for generic event boundary detection. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 20073–20082, 2022. 1, 2, 6, 8
- [16] Will Kay, Joao Carreira, Karen Simonyan, Brian Zhang, Chloe Hillier, Sudheendra Vijayanarasimhan, Fabio Viola, Tim Green, Trevor Back, Paul Natsev, et al. The kinetics human action video dataset. *arXiv preprint arXiv:1705.06950*, 2017. 5
- [17] Dan Kondratyuk, Liangzhe Yuan, Yandong Li, Li Zhang, Mingxing Tan, Matthew Brown, and Boqing Gong. Movinets: Mobile video networks for efficient video recognition. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 16020–16030, 2021. 1
- [18] Heeseung Kwon, Manjin Kim, Suha Kwak, and Minsu Cho. Motionsqueeze: Neural motion feature learning for video understanding. In *Computer Vision—ECCV 2020: 16th European Conference, Glasgow, UK, August 23–28, 2020. Proceedings, Part XVI 16*, pages 345–362. Springer, 2020. 3
- [19] Congcong Li, Xinyao Wang, Dexiang Hong, Yufei Wang, Libo Zhang, Tiejian Luo, and Longyin Wen. Structured context transformer for generic event boundary detection. *arXiv preprint arXiv:2206.02985*, 2022. 1, 2, 6
- [20] Congcong Li, Xinyao Wang, Longyin Wen, Dexiang Hong, Tiejian Luo, and Libo Zhang. End-to-end compressed video representation learning for generic event boundary detection. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 13967–13976, 2022. 1, 2, 6, 8
- [21] Yanghao Li, Chao-Yuan Wu, Haoqi Fan, Karttikeya Mangalam, Bo Xiong, Jitendra Malik, and Christoph Feichtenhofer. Mvity2: Improved multiscale vision transformers for classification and detection. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 4804–4814, 2022. 1
- [22] Xingyu Liu, Joon-Young Lee, and Hailin Jin. Learning video representations from correspondence proposals. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 4273–4281, 2019. 3
- [23] Ze Liu, Jia Ning, Yue Cao, Yixuan Wei, Zheng Zhang, Stephen Lin, and Han Hu. Video swin transformer. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 3202–3211, 2022. 1
- [24] Bruce D Lucas and Takeo Kanade. An iterative image registration technique with an application to stereo vision. In *IJCAI’81: 7th international joint conference on Artificial intelligence*, volume 2, pages 674–679, 1981. 3
- [25] Abhijit Patait. An introduction to the nvidia optical flow sdk, <https://developer.nvidia.com/blog/an-introduction-to-the-nvidia-optical-flow-sdk>. *NVIDIA Developer [online]*, 13, 2019. 8
- [26] Rui Qian, Yeqing Li, Liangzhe Yuan, Boqing Gong, Ting Liu, Matthew Brown, Serge Belongie, Ming-Hsuan Yang, Hartwig Adam, and Yin Cui. Exploring temporal granular-

- ity in self-supervised video representation learning. *arXiv preprint arXiv:2112.04480*, 2021. **2**
- [27] Ayush K Rai, Tarun Krishna, Julia Dietlmeier, Kevin McGuinness, Alan F Smeaton, and Noel E O'Connor. Motion aware self-supervision for generic event boundary detection. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision*, pages 2728–2739, 2023. **2, 3**
- [28] Dian Shao, Yue Zhao, Bo Dai, and Dahua Lin. Intra-and inter-action understanding via temporal action parsing. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 730–739, 2020. **6**
- [29] Jianbo Shi et al. Tomasi. good features to track. In *1994 Proceedings of IEEE Conference on Computer Vision and Pattern Recognition*, pages 593–600. sn, 1994. **3, 7**
- [30] Mike Zheng Shou, Stan Weixian Lei, Weiyao Wang, Deepti Ghadiyaram, and Matt Feiszli. Generic event boundary detection: A benchmark for event segmentation. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 8075–8084, 2021. **1, 2, 3, 5, 6, 7, 8**
- [31] Karen Simonyan and Andrew Zisserman. Two-stream convolutional networks for action recognition in videos. *Advances in neural information processing systems*, 27, 2014. **1, 3**
- [32] Jiaqi Tang, Zhaoyang Liu, Chen Qian, Wayne Wu, and Limin Wang. Progressive attention on multi-level dense difference maps for generic event boundary detection. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 3355–3364, 2022. **1, 2, 3, 6**
- [33] Du Tran, Heng Wang, Lorenzo Torresani, Jamie Ray, Yann LeCun, and Manohar Paluri. A closer look at spatiotemporal convolutions for action recognition. In *Proceedings of the IEEE conference on Computer Vision and Pattern Recognition*, pages 6450–6459, 2018. **1**
- [34] Barbara Tversky and Jeffrey M Zacks. Event perception. *Oxford handbook of cognitive psychology*, 1(2):3, 2013. **1**
- [35] Limin Wang, Yuanjun Xiong, Zhe Wang, Yu Qiao, Dahua Lin, Xiaoou Tang, and Luc Van Gool. Temporal segment networks for action recognition in videos. *IEEE transactions on pattern analysis and machine intelligence*, 41(11):2740–2755, 2018. **1**
- [36] Xiao Wang, Jingen Liu, Tao Mei, and Jiebo Luo. Coseg: Cognitively inspired unsupervised generic event segmentation. *arXiv preprint arXiv:2109.15170*, 2021. **1, 2, 6**
- [37] Xiang Wang, Zhiwu Qing, Ziyuan Huang, Yutong Feng, Shiwei Zhang, Jianwen Jiang, Mingqian Tang, Changxin Gao, and Nong Sang. Proposal relation network for temporal action detection. *arXiv preprint arXiv:2106.11812*, 2021. **1**
- [38] Yuxuan Wang, Difei Gao, Licheng Yu, Weixian Lei, Matt Feiszli, and Mike Zheng Shou. Geb+: A benchmark for generic event boundary captioning, grounding and retrieval. In *Computer Vision–ECCV 2022: 17th European Conference, Tel Aviv, Israel, October 23–27, 2022, Proceedings, Part XXXV*, pages 709–725. Springer, 2022. **1**
- [39] Xiaohao Xu, Jinglu Wang, Xiang Ming, and Yan Lu. Towards robust video object segmentation with adaptive object calibration. In *Proceedings of the 30th ACM International Conference on Multimedia*, pages 2709–2718, 2022. **1**
- [40] Yufei Xu, Qiming Zhang, Jing Zhang, and Dacheng Tao. Vitae: Vision transformer advanced by exploring intrinsic inductive bias. *Advances in Neural Information Processing Systems*, 34:28522–28535, 2021. **1**
- [41] Chen-Lin Zhang, Jianxin Wu, and Yin Li. Actionformer: Localizing moments of actions with transformers. In *Computer Vision–ECCV 2022: 17th European Conference, Tel Aviv, Israel, October 23–27, 2022, Proceedings, Part IV*, pages 492–510. Springer, 2022. **1**
- [42] Linchao Zhu and Yi Yang. Actbert: Learning global-local video-text representations. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 8746–8755, 2020. **1**