



# Local Compressed Video Stream Learning for Generic Event Boundary Detection

Libo Zhang<sup>1,2,3</sup> · Xin Gu<sup>2</sup> · Congcong Li<sup>2</sup> · Tiejian Luo<sup>2</sup> · Heng Fan<sup>4</sup>

Received: 5 January 2023 / Accepted: 25 September 2023

© The Author(s), under exclusive licence to Springer Science+Business Media, LLC, part of Springer Nature 2023

## Abstract

Generic event boundary detection aims to localize the generic, taxonomy-free event boundaries that segment videos into chunks. Existing methods typically require video frames to be decoded before feeding into the network, which contains significant spatio-temporal redundancy and demands considerable computational power and storage space. To remedy these issues, we propose a novel compressed video representation learning method for event boundary detection that is fully end-to-end leveraging rich information in the compressed domain, *i.e.*, RGB, motion vectors, residuals, and the internal group of pictures (GOP) structure, without fully decoding the video. Specifically, we use lightweight ConvNets to extract features of the P-frames in the GOPs and spatial-channel attention module (SCAM) is designed to refine the feature representations of the P-frames based on the compressed information with bidirectional information flow. To learn a suitable representation for boundary detection, we construct the local frames bag for each candidate frame and use the long short-term memory (LSTM) module to capture temporal relationships. We then compute frame differences with group similarities in the temporal domain. This module is only applied within a local window, which is critical for event boundary detection. Finally a simple classifier is used to determine the event boundaries of video sequences based on the learned feature representation. To remedy the ambiguities of annotations and speed up the training process, we use the Gaussian kernel to preprocess the ground-truth event boundaries. Extensive experiments conducted on the Kinetics-GEBD and TAPOS datasets demonstrate that the proposed method achieves considerable improvements compared to previous end-to-end approach while running at the same speed. The code is available at <https://github.com/GX77/LCVSL>.

**Keywords** Generic event boundary detection (GEBD) · Spatial-channel attention module (SCAM) · Group similarity · Local frames bag

## 1 Introduction

In recent years, video has become an integral part of human life, significantly impacting various aspects of our daily rou-

tines and activities. By 2023, the video content will make up 80% of all consumer internet traffic.<sup>1</sup> When perceiving video contents, people will naturally and spontaneously segment events, breaking down longer events into a series of shorter temporal units (Shou et al., 2021). However, this mechanism is tough for machine learning, although it is so natural to the human brain. To this end, Generic Event Boundary Detection

---

Communicated by Dima Damen.

---

Libo Zhang, Xin Gu have equal contribution to this work.

---

✉ Heng Fan  
heng.fan@unt.edu  
Libo Zhang  
libo@iscas.ac.cn  
Xin Gu  
guxin21@mails.ucas.ac.cn  
Congcong Li  
licongcong18@mails.ucas.edu.cn  
Tiejian Luo  
tjluo@ucas.ac.cn

---

<sup>1</sup> Institute of Software Chinese Academy of Sciences, Beijing, China  
<sup>2</sup> University of Chinese Academy of Sciences, Beijing, China  
<sup>3</sup> Nanjing Institute of Software Technology, Nanjing, China  
<sup>4</sup> Department of Computer Science and Engineering, University of North Texas, Denton, TX, USA

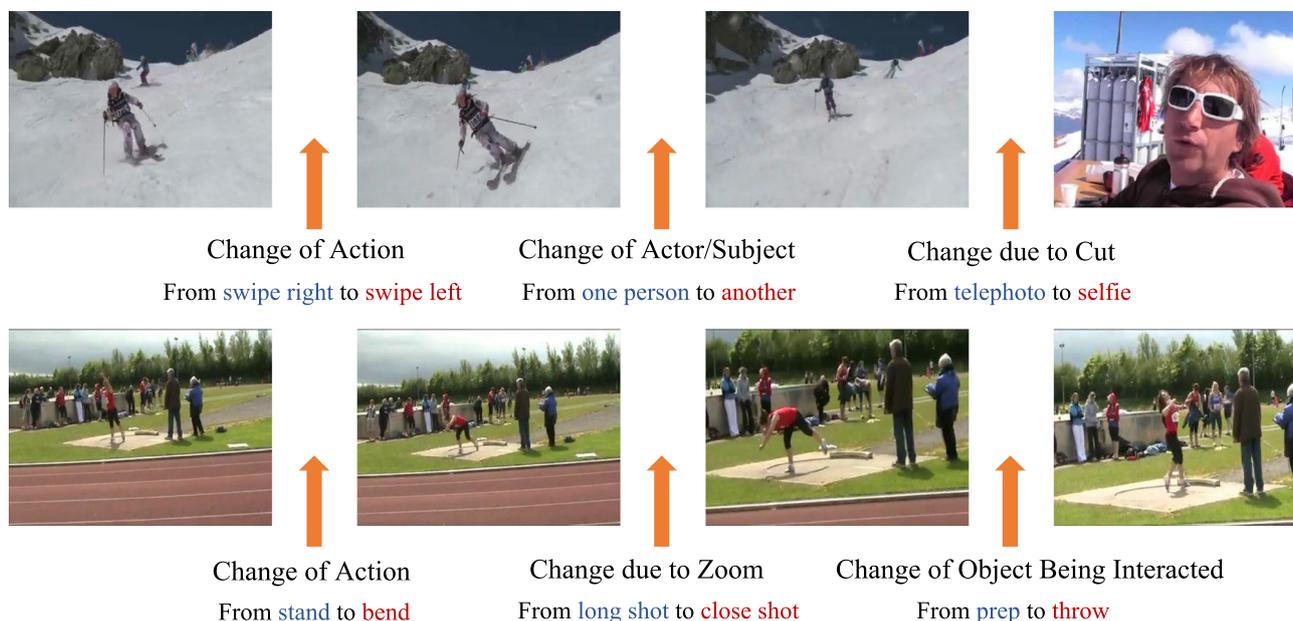
<sup>1</sup> <https://www.meltycone.com/blog/video-marketing-statistics-for-2023>.

(Shou et al., 2021) (GEBD) is proposed to allow machines to develop such an ability.

GEBD aims to localize the moments in which humans naturally perceive event boundaries. The high-level causes of event boundaries in the GEBD task are the following: 1) *change in spatial domain*: significant changes in the color or brightness of the environment. 2) *change in the temporal domain*: an old action ends or a new action starts. Notably, these causes can happen simultaneously or are intermingled together, which lead to complicated event boundary variations, as shown in Fig. 1. To solve the GEBD task, we can simply regard it as a video representation learning problem following the main methods. Currently, the two-stream networks (Simonyan & Zisserman, 2014; Feichtenhofer et al., 2016, 2019) and 3D convolutional networks (Taylor et al., 2010; Ji et al., 2013; Tran et al., 2015; Varol et al., 2018) are two popular network architectures in the video understanding field. The two-stream networks usually incorporate two different modalities of information to learn complementary representation, for example, decoded RGB video frames and optical flow. 3D convolutional network is another choice to model temporal information using the spatio-temporal filters. Transformers were successfully applied in computer vision (Dosovitskiy et al., 2021a), the new trend in video understanding is using the Transformers, including (Dosovitskiy et al., 2021b; Arnab et al., 2021; Liu et al., 2022; Fan et al., 2021; Zhang et al., 2021), which achieve competitive results. Despite their success, these methods are not optimal for the

GEBD task since consecutive decoded RGB frames contain high temporal redundancy and are not practical for real-time applications.

Recently, another alternative for video understanding is learning directly from compressed domain. Several methods (Zhang et al., 2016; Li et al., 2020; Wu et al., 2018; Shou et al., 2019; Wang et al., 2019; Yu et al., 2021; Huang et al., 2021) have demonstrated the advantages of directly taking compressed information in video stream as input for video understanding. These methods usually run in two orders of magnitude faster than the methods using optical flow while achieving competitive results (Shou et al., 2019). This tremendous improvement in speed comes from the using of motion vectors and residuals, which designed for storage and transmission of videos and almost compute-free. The rich information in motion vectors can be regarded as an alternative to the compute-intensive optical flow. To better utilize motion vectors and residuals, different methods have been developed for efficient and effect compressed video representation learning. Specifically, CoViAR (Wu et al., 2018) first converts motion vectors and residuals into 2D representations like images and then directly feeds them into 2D CNNs for action recognition. This method lacks interactions between I-frames and P-frames and thus achieves inferior results. DMC-Net (Shou et al., 2019) improves the CoViAR (Wu et al., 2018) method by reconstructing the optical flow based on motion vectors and residuals and a discriminator is applied to guide the reconstruction. How-



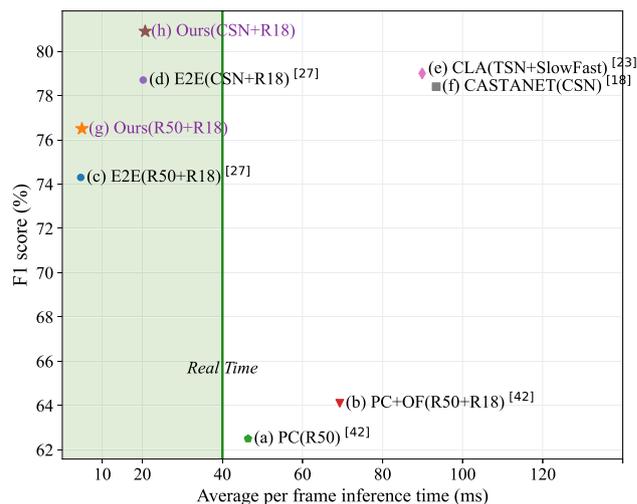
**Fig. 1** Examples of generic event boundaries. The first video is segmented at a change of action, then the change of subject into another one and the change of a different scene. The second video is segmented due

to action change, camera zoom and object interaction. The taxonomy-free nature of Kinetics-GEBD event boundaries makes it harder than existing tasks

ever, it still needs optical flow in the training stage. SIFP (Li et al., 2020) uses the slow I pathway receiving a sparse sampling I-frame clip and the fast P pathway receiving a dense sampling pseudo-optical flow clip, which eliminates the dependence on traditional optical flows calculated from raw videos. Although the aforementioned method achieves promising results, they are still far from satisfactory, which lack effective fusion strategies between different modalities, such as decoded I-frames, motion vectors, and residuals. GEBD task is more sensitive to the local temporal context, which needs a new mechanism to learn from compressed information.

In this paper, we focus on GEBD and develop a fast end-to-end method that can effectively learn from a local compressed video stream. The previous attempt (Shou et al., 2021) formulate it as a classification task by considering the context information of the candidate boundaries. However, it neglects the temporal relations between consecutive frames and operates inefficiently during feature extraction stage. Inspired by Zhang et al. (2016), Wu et al. (2018), Shou et al. (2019), Wang et al. (2019), Yu et al. (2021), Huang et al. (2021), we designed an end-to-end trained network to exploit discriminative features for GEBD in the compressed domain, *i.e.*, MPEG-4, which can save decoding cost and improve feature extraction efficiency. Specifically, most modern codecs split a video into several group of pictures (GOP), where each GOP is formed by one I-frames and  $T$  P-frames. To solve the difficulty that arose from the long chain of dependency of the P-frames, following CoViar (Wu et al., 2018), we use the backtracing technique to compute the accumulated motion vectors and residuals in linear time. In this way, the consecutive P-frames in each GOP depend only on the reference I-frame, which can be processed in parallel.

In contrast to the I-frame, it is difficult to learn the discriminative features of the P-frames. Refining the features of the reference I-frame based on the motion vectors and residuals becomes an intuitive option. Motion vectors and residuals provide information to reconstruct P-frames by referring the dependent I-frames. In addition to that, they also provide motion information that obtained from the video encoding process. To that end, we design a lightweight spatial-channel attention module to refine the features of the reference I-frame with the guidance of the motion vectors and residuals. In this way, the features of P-frames and I-frames are converted to the same feature space, which benefits the subsequent processing. After obtaining P-frame features, we split frame sequences into successive local frames bags. Each local frames bag only contains a fixed number of frames and is responsible for providing necessary context information to determine whether the central frame belongs to an event or not. Then we compute frame differences using group similarity in each local frames bag based on temporal context extracted by the long short-term memory (LSTM) module.



**Fig. 2** The inference time vs. F1 score of different methods on the Kinetics-GEBD validation dataset (Shou et al., 2021). **a** The previous method (Shou et al., 2021) PC is relatively fast with inferior results. **b** After integrating the optical flow (OF) module, the accuracy is improved with much slower running speed. **c, d** (Li et al., 2022) leverages motion vectors and residuals in the compressed domain and achieves a competitive F1 score with a faster running speed. **e, f** CLA (Kang et al., 2021) and CASTANET (Hong et al., 2021) take fully decoded RGB frames as input, which are much slower than the methods conducted in compressed domain. **g, h** Compared with **c, d**, our method obtains about 2.0% absolute improvements while keeping almost the same running speed. The green region indicates that the methods are run in real time. Best viewed in color for all figures throughout the paper

This module can predict the event boundaries of videos accurately and it actually imitates humans, *i.e.*, look back and forth around the candidate frames to determine event boundaries, by comparing the extracted features before and after the candidate frames. In addition, to remedy the ambiguities of annotations and speed up the training process, we use the Gaussian kernel to preprocess the ground-truth event boundaries instead of using the “hard labels” of boundaries. Extensive experiments conducted on the Kinetics-GEBD and TAPOS datasets to demonstrate the effectiveness of the proposed method. Specifically, the proposed method achieves comparable results to the state-of-the-art method at the CVPR’21 LOVEU Challenge (Kang et al., 2021) with much faster running speed, as in Fig. 2.

In summary, we make the following contributions:

- We propose a spatial-channel attention module (SCAM), which also refines P-frame features with I-frame features and shows advantages with bidirectional information flow compared to spatial-channel compressed encoder (SCCE).
- Instead of using linear temporal contrastive module, we split frame sequences into successive local frames bags and compute frame differences using group similarity

in each local frames bag based on temporal context extracted by LSTM module.

- We provide additional ablation studies and qualitative analyses to demonstrate the superiority of each components comprehensively, including different backbones, different temporal feature extractor and learning scheme of different number of annotators.
- We achieve comparable results to the state-of-the-art methods at the CVPR'21 LOVEU Challenge (Kang et al., 2021) with much faster running speed and obtains about 2% absolute improvements compared with our preliminary version (Li et al., 2022), demonstrating its effectiveness. We also achieve 1.4% improvement on TAPOS (Shao et al., 2020) dataset compared to DDM-Net (Tang et al., 2022), which is developed on fully decoded RGB frames.

This paper is an extended version of a preliminary conference publication (Li et al., 2022). The main new contributions or differences include:

- 1) We improve the original SCCE by proposing the new SCAM that refines P-frame feature with I-frame features using bidirectional information flow.
- 2) We propose to utilize the LSTM module to capture temporal information for better performance.
- 3) We carry out more ablative studies to analyze each component of our approach in-depth.
- 4) Notable performance gains are achieved with the aforementioned new contributions in comparison with our preliminary version in Li et al. (2022).

The remainder of this paper is organized as follows. A brief review of related works is presented in Sect. 2. The details of the proposed end-to-end method for generic event boundary detection and experimental explanations are given in Sect. 3. Extensive experiments and ablation studies are given in Sect. 4. We conclude our method in Sect. 5.

## 2 Related Work

### 2.1 Temporal Action Localization (TAL)

TAL aims to localize the action segments from untrimmed videos. More specifically, for each action segment, the goal is to detect the start point, the end point and the action class it belongs to. Most approaches could be categorised into two groups, including two-stage methods (Richard & Gall, 2016; Ni et al., 2016; Caba Heilbron et al., 2017; Zhao et al., 2017; Chao et al., 2018) and single-stage methods (Lea et al., 2017; Lin et al., 2017; Alwassel et al., 2018; Long et al., 2019; Yuan et al., 2017; Ma et al., 2016; Yuan et al., 2017; Zhao et al.,

2020). In the two-stage method setting, the first stage generates action segment proposals. The actionness and the type of action for each proposal are then determined by the second stage, along with some post-processing methods such as grouping (Zhao et al., 2017) and Non-maximum Suppression (NMS) (Lin et al., 2019a) to eliminate redundant proposals. For one-stage methods, the classification is performed on the pre-defined anchors (Lin et al., 2017; Long et al., 2019) or video frames (Ma et al., 2016; Yuan et al., 2017).

TAL and GEBD are both tasks in video understanding, which solve the problem of boundary localization by extracting key information from videos. However, TAL is designed to locate specific action behaviors only, whereas GEBD can locate general events beyond action behaviors. In other words, TAL is a subtask within GEBD that can only locate actions. Research in TAL has provided many methods (Richard & Gall, 2016; Ni et al., 2016; Caba Heilbron et al., 2017; Lea et al., 2017; Lin et al., 2017; Alwassel et al., 2018; Long et al., 2019; Yuan et al., 2017; Ma et al., 2016) for action localization, and we have addressed the GEBD task by extending these methods to generic event localization.

### 2.2 Generic Event Boundary Detection

The goal of GEBD (Shou et al., 2021) is to localize the taxonomy-free event boundaries that break a long event into several short temporal segments. Different from Temporal Action Localization (TAL), GEBD only requires to predict the boundaries of each continuous segments. The current methods (Kang et al., 2021; Hong et al., 2021; Rai et al., 2021) all follow the similar fashion in Shou et al. (2021), which takes a fixed length of video frames before and after the candidate frame as input, and separately determines whether each candidate frame is the event boundary or not. Kang et al. (2021) use the temporal self-similarity matrix (TSM) as the intermediate representation and exploit the discriminative features with the popular contrastive learning approach for better performance. Hong et al. (2021) use the cascade classification heads and dynamic sampling strategy to boost both recall and precision. Rai et al. (2021) attempt to learn the spatiotemporal features using a two stream inflated 3D convolutions architecture. DDM-Net (Tang et al., 2022) present dense difference maps (DDM) to comprehensively characterize the motion pattern and exploit progressive attention on multi-level DDM to jointly aggregate appearance and motion clues. All these methods are developed on decoded RGB images and cannot benefit from almost compute-free motion vectors and residuals.

### 2.3 Attention Mechanism

Attention mechanism has been widely adopted in deep learning model design. The core of attention mechanism is to

recalibrate the origin input feature with different weights in different dimensions. Transformer (Vaswani et al., 2017a) and Non-local network (Wang et al., 2018b) can capture long-range dependencies by computing the response at a spatial position as a weighted sum of the features at all positions in the input feature maps. SENet (Hu et al., 2018) develops the “Squeeze-and-Excitation” (SE) block that adaptively recalibrates channel-wise feature responses by explicitly modelling interdependencies between channels. In addition to channel-wise recalibration, CBAM (Woo et al., 2018) sequentially infers attention maps along both channel and spatial dimensions, and then uses the attention maps to recalibrate the origin input feature. In contrast to the aforementioned methods, we attempt to refine the P-frame feature with the guidance of motion vectors and residuals by considering both spatial and channel dimensions of the features of I-frame. Different from SCCE (Li et al., 2022), we also refine the origin P-frame feature with the guidance of reference I-frame feature. This bidirectional information flow can filter noises and fully leverages the information in compressed video stream.

### 3 Method

The existing method (Shou et al., 2021) formulates the GEBD task as binary classification, which predicts the boundary labels for each frame by considering the temporal context information. That is, The preceding and succeeding frames of each video frame are feed into a neural network to detect the boundaries. It is inefficient due to the duplicated computation conducted in consecutive frames. To remedy this, we propose an end-to-end compressed video representation method for GEBD, which regards each video clip as a whole. Specifically, we use MPEG-4 encoded videos as our input. Each video clip  $\mathcal{V}$  is formed by  $N$  groups of pictures (GOPs), and each GOP contains one I-frame and  $T$  P-frames, *i.e.*,

$$\mathcal{V} = \{I_i, P_i^1, P_i^2, \dots, P_i^T\}_{i=1}^N, \quad (1)$$

where  $I_i \in \mathbb{R}^{3 \times \mathcal{H} \times \mathcal{W}}$  denotes the reference I-frame and  $P_i^t$  denotes the  $t$ -th P-frame of the  $i$ -th GOP, and  $\mathcal{H}$  and  $\mathcal{W}$  are the height and width of the video frame. For simplicity, we assume that there exists the same number of P-frames in all GOPs. The assumption of a fixed number of P-frames simplifies the process, as it allows us to standardize the temporal length of each video segment for processing. The P-frame  $P_i^t$  in the  $i$ -th GOP is formed by the initial motion vector  $\mathcal{M}_i^t \in \mathbb{R}^{2 \times \mathcal{H} \times \mathcal{W}}$  and initial residual  $\mathcal{R}_i^t \in \mathbb{R}^{3 \times \mathcal{H} \times \mathcal{W}}$ , which can be firstly obtained nearly cost-free from the compressed video stream, and then trace all motion vectors back until to the reference I-frame and accumulate the residual on the way to decouple the dependencies between the consec-

#### Algorithm 1 Local Compressed Video Stream Learning for GEBD

**Require:** Input the I-frame  $I$ ,  $T$  motion vectors  $\mathcal{M}^l$  and residuals  $\mathcal{R}^l$ ,  $l$  is the position of the candidate frame in the time sequence, and  $k$  is the number of adjacent frames.

```

1:  $\mathbf{x}_I = \text{ResNet-50}(I)$ 
2: for  $t = 1, \dots, T$  do
3:    $\mathbf{x}_M^t = \text{ResNet-18}(\mathcal{M}^t)$ 
4:    $\mathbf{x}_R^t = \text{ResNet-18}(\mathcal{R}^t)$ 
5:    $\tilde{\mathbf{v}}^t = \text{SCAM}(\mathbf{x}_I, \mathbf{x}_M^t, \mathbf{x}_R^t)$ 
6: end for
7:  $\mathcal{B}^l = \{\tilde{\mathbf{v}}^{l-k}, \dots, \tilde{\mathbf{v}}^l, \dots, \tilde{\mathbf{v}}^{l+k}\}$ 
8:  $\tilde{\mathcal{B}}^l = \text{LSTM}(\mathcal{B}^l)$ 
9:  $\mathcal{S}^l = \text{Group-Similarity}(\tilde{\mathcal{B}}^l, \tilde{\mathcal{B}}^l)$ 
10:  $\mathcal{P}^l = \text{Average-Pool}(\text{FCN}(\mathcal{S}^l))$ 
11:  $\mathcal{Y}^l = \text{Classifier-Head}(\mathcal{P}^l)$ 
12: return  $\mathcal{Y}^l$ 

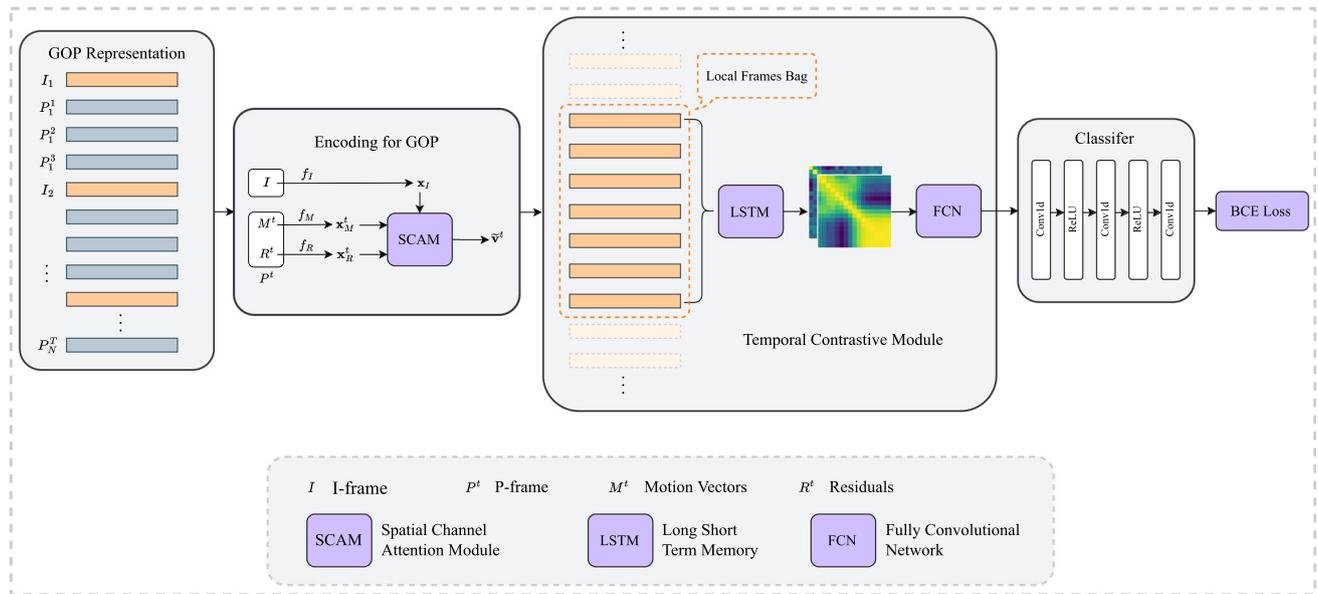
```

utive P-frames. In this way, each P-frame only depends on the reference I-frame rather than other P-frames. After that, we build our model based on the backtraced motion vectors and residuals and regard each GOP as a process unit. The overall network architecture is presented in Fig. 3. As shown in Fig. 3, the GOP is first encoded by the designed spatial-channel attention module (SCAM) to generate the unified video representation. After that, a temporal contrastive module is used to exploit the temporal context information to obtain the discriminative feature representations. Finally, a classifier is used to generate the accurate event boundaries. Our algorithm flow is shown in Algorithm 1.

#### 3.1 Learning from Spatial-Channel Attention Module

Motion, uncovered regions, and lighting variations frequently happen in video sequences. Modern codecs use macroblock as the basic unit for motion-compensated prediction in a number of mainstream visual coding standards such as MPEG-4, H.263, and H.264. Motion vectors record the moving direction of each macroblock with respect to its reference frame(s), describing the motion patterns of videos, which is important for the GEBD task. The residuals can be regraded as the compensations of the motion information, which contains the boundary information of moving objects and plays a crucial role in identifying the important regions in the I-frame. Thus, we propose applying the attention mechanism to different regions of I-frame with the guidance of motion vectors to enrich the features by considering both channel and spatial dimensions. For simplicity, we omit the index  $i$  of the GOP in the following sections.

Firstly, we use the convolutional neural network taking the decoded RGB image as input to extract the feature representation  $\mathbf{x}_I$  of the I-frame  $I$ , *i.e.*,  $\mathbf{x}_I = f_I(I)$ , where  $\mathbf{x}_I \in \mathbb{R}^{C \times H \times W}$  is the features of the I-frame  $I$ , and  $C$ ,  $H$  and  $W$  are the channel, height and width of the features  $\mathbf{x}_I$ ,



**Fig. 3** The architecture of our method. The spatial-channel attention module (SCAM) is designed to obtain the refined P-frame representation  $\tilde{v}^t$  based on reference I-frame feature  $\mathbf{x}_I$ , resized motion vectors  $M^t$  and resized residuals  $R^t$ . This module regards each GOP as a process unit, which is efficient and can be paralleled in a large batch size. Then

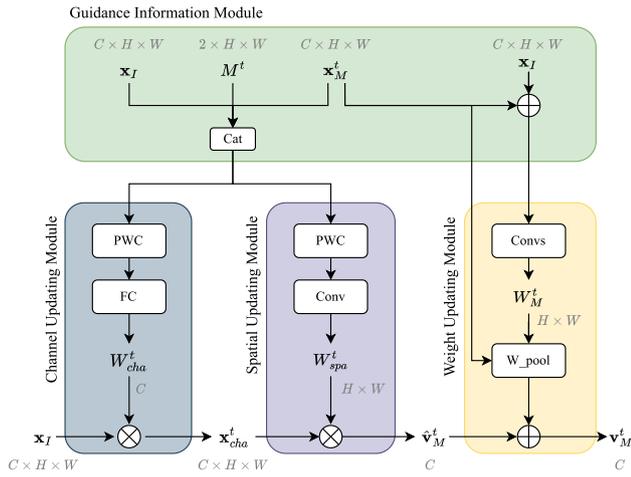
we use temporal contrastive module to capture temporal dependence explicitly based on unified representation  $\tilde{v}^t$ , which provides strong cues for boundary detection. After that, a simple classifier is applied to make the final predictions trained with the Gaussian smoothed soft labels

respectively.  $f_I(\cdot)$  denotes the model used to extract features for the I-frame, which is pre-trained on large-scale datasets (e.g., ResNet50 pre-trained on ImageNet). Meanwhile, we can similarly compute the features for the P-frames  $\{P^1, P^2, \dots, P^T\}$  with ResNet-18, by directly taking the initial motion vectors  $\mathcal{M}^t$  and initial residuals  $\mathcal{R}^t$ , i.e.,  $\mathbf{x}_M^t = f_M(\mathcal{M}^t)$ , and  $\mathbf{x}_R^t = f_R(\mathcal{R}^t)$  as input, where  $\mathbf{x}_M^t, \mathbf{x}_R^t \in \mathbb{R}^{C \times H \times W}$  denote the features of the motion vectors and residuals, respectively.

We choose the ResNet-18 to extract motion vectors and residuals features for two reasons: (1) Motion vectors and residuals, as types of compressed-domain information, are typically characterized by an uneven distribution with most positions being empty and only a few containing meaningful data. Therefore, for such sparse information, using a small network (such as ResNet-18) to extract features may be more effective in processing it. (2) ResNet-18 has a powerful generalization capability and the success it has achieved in many image analysis tasks. Although motion vectors and residuals may not directly correlate with natural images, the key concept is to leverage the high-level feature extraction ability of the pre-trained network to create useful representations. In this way, a considerable amount of time can be saved on extracting features for the P-frames. This simple strategy (Wu et al., 2018), in which the motion vector and residual are treated as separate branches without fusion, can only provide limited performance improvement. The method (Shou

et al., 2019) attempts to integrate the optical flow in the training phase, which can further improve the accuracy. However, there is still much room for improvement of the aforementioned methods. Specifically, the motion vectors record the motion patterns of both the scenes and objects in videos, and the residuals provide the compensation information. Both of them do not contain the context information of the scenes. To this end, we design the spatial channel compressed encoder module by integrating the features of the reference I-frame  $x_I$  in computing the features of P-frames.

We first compute the features  $\mathbf{x}_M^t$  of the motion vectors by refining the features of the reference I-frame  $\mathbf{x}_I$  in both the channel and spatial dimensions. As indicated by Zeiler and Fergus (2014), different regions on the feature maps focus on different parts of the images. Thus, we introduce the attention weight for each feature map of  $\mathbf{x}_I$  based on the information of the P-frame  $\mathbf{x}_M^t$ . Specifically, we concatenate I-frame feature  $\mathbf{x}_I$ , motion vector feature  $\mathbf{x}_M^t$  and resized motion vectors  $M^t$  (resized from  $\mathcal{M}^t$ ) together in the channel dimension to compute the channel weight  $W_{\text{cha}}^t$  using a lightweight PWC-Net (Sun et al., 2018). PWC-Net is chosen to maintain consistency with previous work (Shou et al., 2019), and because of the similarity between motion vectors and optical flow information, which makes it suitable for processing using an



**Fig. 4** The architecture of the proposed spatial-channel attention module (SCAM). W\_pool corresponds to the weighted pool function. We concatenate the features of I-frame  $\mathbf{x}_I$ , resized motion vectors  $M^t$  and the features of motion vectors  $\mathbf{x}_M^t$  to modulate the features of the reference I-frame  $\mathbf{x}_I$  in both channel and spatial dimensions. After that, we also use I-frame  $\mathbf{x}_I$  feature to modulate the origin P-frame feature  $\mathbf{x}_M^t$  with spatial weight  $W_M^t$ . Different from the previous method (Li et al., 2022) where the modulated features  $\hat{\mathbf{v}}_M^t$  is residually added with the features of motion vectors, our method can learn more robust representation with the bidirectional information flow

optical flow network. *i.e.*,

$$\begin{aligned} \mathbf{z}_{cha}^t &= \text{PWC}([\mathbf{x}_I; \mathbf{x}_M^t; M^t]) \\ \mathbf{h}_{cha}^t &= \text{avg\_pool}(\mathbf{z}_{cha}^t) \\ W_{cha}^t &= \sigma(W_2 \cdot \zeta(W_1 \mathbf{h}_{cha}^t + b_1) + b_2) \end{aligned} \tag{2}$$

where  $\sigma$  is the sigmoid function,  $\zeta$  is the ReLU function, and  $W_1, b_1, W_2, b_2$  are the learnable weights of the FC layers. After that, the features of the I-frame  $\mathbf{x}_I$  are updated based on  $W_{cha}^t$  as follows,

$$\mathbf{x}_{cha}^t = \mathbf{x}_I \otimes W_{cha}^t \tag{3}$$

where  $\otimes$  is the channel-wise multiplication. In this way, we can compute the channel-weighted feature  $\mathbf{x}_{cha}^t$  by updating  $\mathbf{x}_I$  in channel dimension, with the guidance of the motion vectors. Meanwhile, the channel-weighted feature  $\mathbf{x}_{cha}^t$  is further updated in the spatial dimension and the spatial dimension is reduced. That is, given the features  $\mathbf{x}_I$  of the reference I-frame, motion vector features  $\mathbf{x}_M^t$  and resized motion vectors  $M^t$ , we compute the 2D weight map  $W_{spa}^t$ , *i.e.*,

$$\begin{aligned} \mathbf{z}_{spa}^t &= \text{PWC}([\mathbf{x}_I; \mathbf{x}_M^t; M^t]) \\ \mathbf{h}_{spa}^t &= \text{2d\_conv}(\mathbf{z}_{spa}^t) \\ W_{spa}^t &= \text{softmax}(\mathbf{h}_{spa}^t) \end{aligned} \tag{4}$$

where  $W_{spa}^t \in \mathbb{R}^{H \times W}$  is the spatial weight map. The softmax function is applied to  $\mathbf{h}_{spa}^t$  across a 2D spatial map, guaranteeing that the sum of values at each spatial location in the

feature map equals 1. After that, we use  $W_{spa}^t$  to weight the features  $\mathbf{x}_{cha}^t$  in the spatial dimension to compute the enriched features of the motion vectors  $\hat{\mathbf{v}}_M^t \in \mathbb{R}^C$ , *i.e.*,

$$\hat{\mathbf{v}}_M^t = \sum_{i=1}^H \sum_{j=1}^W \mathbf{x}_{cha}^t(i, j) \cdot W_{spa}^t(i, j) \tag{5}$$

where  $i, j$  are the spatial positions of  $\mathbf{x}_{cha}^t$  and  $W_{spa}^t$ . Previous method (Li et al., 2022) obtains the refined features of the motion vectors  $\mathbf{v}_M^t \in \mathbb{R}^C$  use residual addition, which doesn't consider the information flow from I-frame feature to P-frame feature. Thus we propose to modulate P-frame feature in spatial dimension by using weighted pooling method. The method applies the weight value  $W_M^t$  (shape of  $H \times W$ ) to each channel at the same position in  $\mathbf{x}_M^t$  (shape of  $C \times H \times W$ ), thereby performing the weighted pooling, *i.e.*,

$$\mathbf{v}_M^t = \hat{\mathbf{v}}_M^t + \text{weighted\_pool}(\text{Convs}(\mathbf{x}_M^t + \mathbf{x}_I)) \tag{6}$$

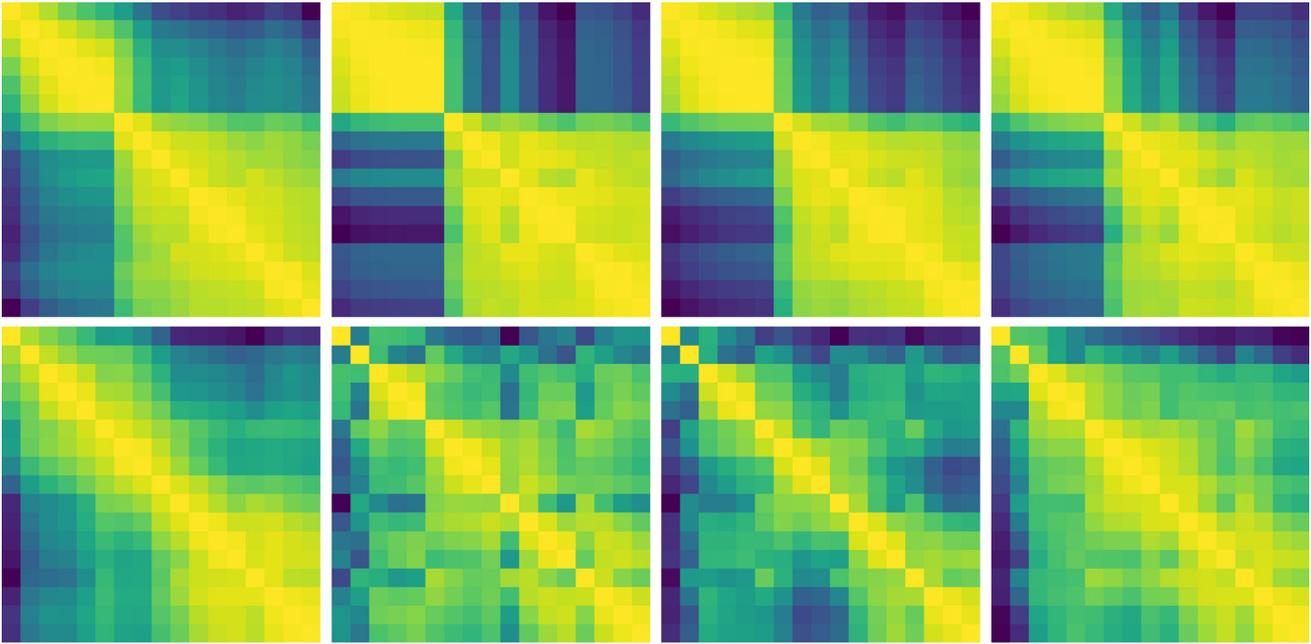
where Convs is a 3-layer Conv-ReLU network and the output channel of last layer is 1 for predicting a 2D spatial weight map. The overall computing process of  $\mathbf{v}_M^t$  is presented in Fig. 4. Similarly, we can compute the refined features for the residuals  $\mathbf{v}_R^t \in \mathbb{R}^C$ . The final feature representations for the P-frame is further computed as

$$\tilde{\mathbf{v}}^t = \mathbf{v}_M^t + \mathbf{v}_R^t \tag{7}$$

In this way, we can compute the features of the P-frames  $\{\tilde{\mathbf{v}}^1, \tilde{\mathbf{v}}^2, \dots, \tilde{\mathbf{v}}^T\}$  in the GOP by considering the reference I-frame  $I$  in both channel and spatial dimensions. The overall process is very efficient and can be processed in parallel in GOPs. After extracting the discriminative features for both the I-frames and P-frames in the same feature space, we can predict the event boundaries efficiently and accurately.

### 3.2 Learning from Local Frames Bag

Based on the extracted features of the video  $\mathcal{V}$ , we aim to design a temporal module to predict the event boundaries accurately. The existence of an event boundary in a video clip implies that there is a visual content change at that point, thus it is very difficult to infer the boundary from one single frame. As a result, the key clue for event boundary detection is to localize changes in the temporal domain. Inspired by humans, *i.e.*, look back and forth around the candidate boundary frames to determine event boundaries, we construct a local frames bag for each candidate frame and each local frames bag is responsible for providing context information to predict an event boundary. Specifically, for each candidate frame  $\tilde{\mathbf{v}}^l$  at time  $l$ , we construct a local frames bag for it by gathering adjacent  $k$  frames before candidate frame  $\tilde{\mathbf{v}}^l$  and adjacent  $k$  frames after  $\tilde{\mathbf{v}}^l$ , resulting in a local frames



**Fig. 5** Visualization of grouped similarity maps  $S_t^l$ ,  $G = 4$  in this example. First row indicates that there is a potential boundary in this local sequence while the second row shows no boundary in this sequence.

sequence, namely local frames bag:

$$\mathcal{B}^l = \{\tilde{\mathbf{v}}^{l-k}, \tilde{\mathbf{v}}^{l-(k-1)}, \dots, \tilde{\mathbf{v}}^l, \dots, \tilde{\mathbf{v}}^{l+(k-1)}, \tilde{\mathbf{v}}^{l+k}\}, \quad (8)$$

which can be implemented through efficient memory view method provided by modern deep learning framework and processed in parallel.

After obtaining local frames bag  $\mathcal{B}^l$  for candidate frame  $\tilde{\mathbf{v}}^l$ , we use a 2-layer long short-term memory (LSTM) to learn temporal relationships. Formally, for each frame  $\tilde{\mathbf{v}}^l$  in local frames bag  $\mathcal{B}^l$ , we compute the hidden state in each layer as follows:

$$\begin{aligned} i^t &= \sigma(W_{ii}\tilde{\mathbf{v}}^t + b_{ii} + W_{hi}h_{t-1} + b_{hi}) \\ f^t &= \sigma(W_{if}\tilde{\mathbf{v}}^t + b_{if} + W_{hf}h_{t-1} + b_{hf}) \\ g^t &= \tanh(W_{ig}\tilde{\mathbf{v}}^t + b_{ig} + W_{hg}h_{t-1} + b_{hg}) \\ o^t &= \sigma(W_{io}\tilde{\mathbf{v}}^t + b_{io} + W_{ho}h_{t-1} + b_{ho}) \\ c^t &= f^t \odot c^{t-1} + i^t \odot g^t \\ h^t &= o^t \odot \tanh(c^t), \end{aligned} \quad (9)$$

where  $h^t$  is the hidden state at time  $t$ ,  $c^t$  is the cell state at time  $t$ ,  $\tilde{\mathbf{v}}^t$  is the input frame at time  $t$ ,  $h^{t-1}$  is the hidden state of the layer at time  $t-1$  or the initial hidden state at time 0, and  $i^t$ ,  $f^t$ ,  $g^t$ ,  $o^t$  are the input, forget, cell, and output gates, respectively.  $\sigma$  is the sigmoid function, and  $\odot$  is the Hadamard product. While there are more options to learn temporal relations like famous Transformer (Vaswani et al.,

We can also observe slightly different patterns between the same group, which may imply that each group is learning in a different aspect

2017b), we use a simple 2-layer LSTM which works very well as shown in Table 5.

After learning the temporal relationships, we obtain temporal-enhanced local frames bag  $\tilde{\mathcal{B}}^l = \{h^{l-k}, h^{l-(k-1)}, \dots, h^l, \dots, h^{l+(k-1)}, h^{l+k}\}$ . The LSTM module aims at discovering relationships between frames and giving high level representation of frames sequences. However, event boundaries emphasize the differences between adjacent frames and neural networks tend to take shortcuts during learning (Geirhos et al., 2020). Thus classifying these frames directly into boundaries may lead to inferior performance due to non-explicit cues. Based on this intuition, we propose to guide classification with feature similarity of each frame pair in the local frames bag  $\tilde{\mathcal{B}}^l$ . Instead of performing similarity calculation with all  $C$ -dimensional channels, we found it beneficial to split the channels into several groups and calculate the similarity of each group independently. Formally, the group similarity map  $S^l$  is calculated as follows. First, split  $\tilde{\mathcal{B}}^l$  in the channel dimension to get  $\{\tilde{b}_i^l\}_{i=1}^G$ , where  $\tilde{b}_i^l \in \mathbb{R}^{\frac{C}{G} \times (2k+1) \times (2k+1)}$  and  $G$  is the number of groups. Second, we calculate the cosine similarity of each  $\tilde{b}_i^l$  as follows:

$$S^l = [\text{sim}(\tilde{b}_1^l, \tilde{b}_1^l), \dots, \text{sim}(\tilde{b}_G^l, \tilde{b}_G^l)], \quad (10)$$

where  $S^l \in \mathbb{R}^{G \times (2k+1) \times (2k+1)}$ . As the group similarity map  $S^l$  contains the similarity scores, it shows different patterns (as shown in Fig. 5) in different sequences, which are crit-

ical for boundary detection. Then we use a 4-layer fully convolutional network (Long et al., 2015) to learn the similarity patterns, which we found works very well and efficient enough. Then we average pool the output of FCN to get a vector representation  $\mathcal{P}^l$ , and this vector is used for final event boundary classification:

$$\mathcal{P}^l = \text{average-pool}(\text{FCN}(S^l)), \quad (11)$$

where  $\mathcal{P}^l \in \mathbb{R}^C$ . Then for the final classification, we use the contrastive representations  $\{\mathcal{P}^1, \mathcal{P}^2, \dots, \mathcal{P}^T\}$  to make the event boundary predictions.

### 3.3 Loss Function

Given feature representations  $\{\mathcal{P}^1, \mathcal{P}^2, \dots, \mathcal{P}^T\}$  of each video frame and the corresponding ground-truth labels, the event boundary detection task is intuitively formulated as the binary classification task. However, the ambiguities of annotations disrupt the learning process, which leads to poor convergence. To solve this issue, we use the Gaussian kernel to preprocess the ground-truth event boundaries to obtain the soft labels instead of using the “hard labels” of boundaries. Specifically, for each annotated boundary, the intermediate label of the neighboring position  $i$  is computed as:

$$g_i^l = \exp\left(-\frac{(l-i)^2}{2\alpha^2}\right) \quad (12)$$

where  $g_i^l$  indicates the intermediate label at time  $i$  corresponding to the annotated boundaries at time  $l$ . We set  $\alpha = 1$  in all our experiments. The final soft labels are computed as the summation of all intermediate labels. Finally, a simple nonlinear Conv1D classifier is applied to predict the boundary score  $S^l$  and the binary cross-entropy loss is used to guide the training process.

## 4 Experiments

We conduct our experiments on the Kinetics-GEBD (Shou et al., 2021) and TAPOS (Shao et al., 2020) datasets. The Kinetics-GEBD dataset contains the largest number of temporal boundaries, including 54, 691 videos and 1, 290, 000 event boundaries, spans a broad spectrum of video domains in the wild and is open-vocabulary rather than building on a pre-defined taxonomy. The TAPOS dataset contains Olympics sport videos with 21 actions. The training set contains 13,094 action instances and the validation set contains 1, 790 action instances. Since it is not suitable for GEBD task, following (Shou et al., 2021), we re-purpose TAPOS for GEBD task by trimming each action instance with its action label hidden and conducting experiments on each action instance.

Furthermore, to verify the generality and effectiveness of our method, we also conducted experiments on the popular action recognition datasets UCF101 (Soomro et al., 2012) and HMDB51 (Kuehne et al., 2011). UCF101 consists of 101 action classes in 13, 320 videos, and HMDB51 contains 51 distinct action categories with a total of 6, 766 video clips.

To quantitatively evaluate the results of the generic event boundary detection task, the F1 score is used as the measurement metric. As described in Shou et al. (2021), Rel.Dis. (Relative distance, the error between the detected and ground truth timestamps, divided by the length of the corresponding whole action instance) is used to determine whether a detection is correct (*i.e.*,  $\leq$  threshold) or incorrect (*i.e.*,  $>$  threshold). A detection result is compared with each rater’s annotation, and the highest F1 score is treated as the final result. We report F1 scores of different thresholds range from 0.05 to 0.5 with a step of 0.05. In particular, all experimental results are the average of multiple experimental results.

### 4.1 Implementation Detail

We implement our method with the popular deep learning framework PyTorch (Paszke et al., 2019). ResNet50 and ResNet18 (He et al., 2016) pretrained on ImageNet (Deng et al., 2009) are used to extract the features for I-frames and P-frames in all experiments if not particularly indicated. Our method is implemented based on the MPEG-4 Part 2 specifications (Gall, 1991), where each GOP contains 1 I-frame and 11 P-frames. We sample 3 P-frames in each GOP to reduce the redundancy, *i.e.*,  $T = 3$  in (1). We use the standard SGD with momentum set to 0.9, weight decay set to  $10^{-4}$ , and learning rate set to  $10^{-2}$ . We set the batch size to 4 for each GPU and train the network on 8 NVIDIA Tesla V100 GPUs, resulting in a total batch size of 32. The network is trained for 30 epochs with a learning rate drop by a factor of 10 after 16 epochs and 24 epochs, respectively. We test the running speed of all methods on 1 NVIDIA Tesla V100 GPU. All the source code of our method will be made publicly available after the paper is accepted.

We also evaluate our method with different backbones including CSN (Tran et al., 2019), ViT-Base (Dosovitskiy et al., 2021a), and Swin-Tiny (Liu et al., 2021) to compare with the state-of-the-art methods in the LOnG-form VidEo Understanding Challenge (LOVEU).

### 4.2 Results and Analysis

We first train and evaluate the proposed method on the Kinetics-GEBD (Shou et al., 2021) train-validation split. The evaluation protocol presented in (Shou et al., 2021) uses Relative Distance (*i.e.*, **Rel.Dis.**, the error between the predicted and ground truth timestamps) to determine whether a prediction is correct or not and then use the precision, recall, and F1

**Table 1** The evaluation results on the Kinetics-GEBD validation set with different Rel.Dis. thresholds

Rel.Dis. Threshold	0.05	0.1	0.15	0.2	0.25	0.3	0.35	0.4	0.45	0.5	avg
BMN Lin et al. (2019b)	0.186	0.204	0.213	0.220	0.226	0.230	0.233	0.237	0.239	0.241	0.223
BMN-StartEnd Shou et al. (2021)	0.491	0.589	0.627	0.648	0.660	0.668	0.674	0.678	0.681	0.683	0.640
TCN-TAPOS Shou et al. (2021)	0.464	0.560	0.602	0.628	0.645	0.659	0.669	0.676	0.682	0.687	0.627
TCN Lea et al. (2016)	0.588	0.657	0.679	0.691	0.698	0.703	0.706	0.708	0.710	0.712	0.685
PC Shou et al. (2021)	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	0.646	0.776	0.818	0.842	0.856	0.864	0.868	0.874	0.877	0.879	0.830
E2E Li et al. (2022)	0.743	0.830	0.857	0.872	0.880	0.886	0.890	0.893	0.896	0.898	0.865
DDM-Net Tang et al. (2022)	0.764	0.843	0.866	0.880	0.887	0.892	0.895	0.898	0.900	0.902	0.873
Ours	<b>0.768</b>	<b>0.848</b>	<b>0.872</b>	<b>0.885</b>	<b>0.892</b>	<b>0.896</b>	<b>0.899</b>	<b>0.901</b>	<b>0.903</b>	<b>0.906</b>	<b>0.877</b>

Bold numbers indicate the best performance under the corresponding metrics

scores as the evaluation metrics. We present all results with Rel.Dis. threshold set from 0.05 to 0.5 with 0.05 interval as shown in Table 1. Our method improves the F1 score over all thresholds by a large margin. Compared to the previous baseline method PC (Shou et al., 2021), our method achieves an absolute improvement of 14.3% while running 10× faster.

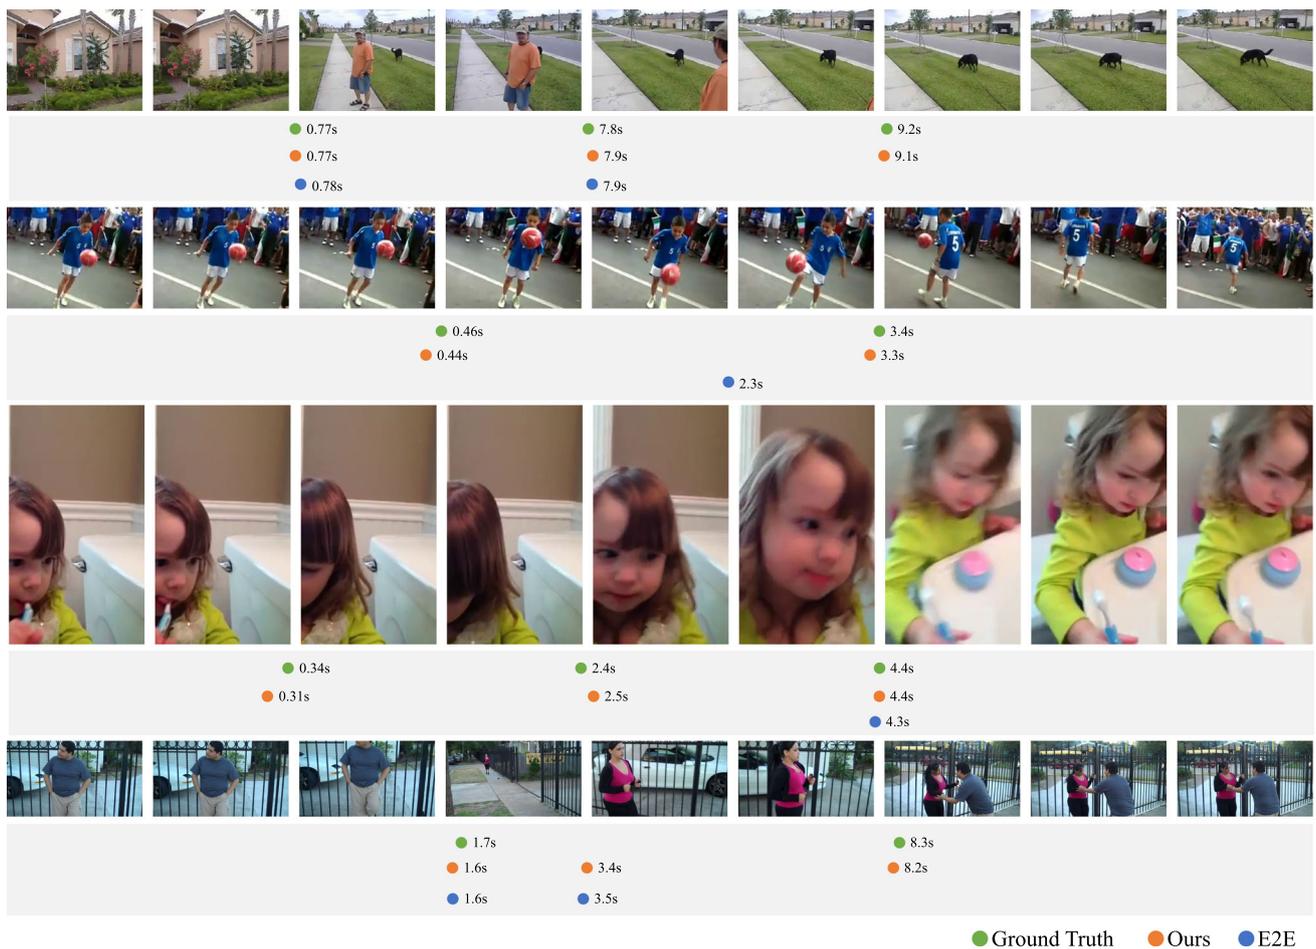
Compared to the previous end-to-end method E2E (Li et al., 2022), our method achieves an absolute improvement of 2.5% while maintaining almost the same running speed. Improvements come mainly from the advanced spatial-channel compressed encoder and the local frames bag. Compared to the baseline method PC (Shou et al., 2021) with optical flow input stream, our advanced spatial-channel compressed encoder can be a better alternative to learning temporal information from cost-free motion vectors and residuals in compressed videos. Our local frames bag can also explicitly provide strong temporal signals for GEBD, giving 2.1% absolute improvements compared to E2E (Li et al., 2022). Example qualitative results on Kinetics-GEBD are shown in Fig. 6. It's worth noting that our method has linear computational complexity with respect to the video length thus can scale well. We also evaluate our method on the TAPOS (Shao et al., 2020) train validation split. The results are shown in Table 2. Compared to DDM-Net (Tang et al., 2022), we increase the F1 score @ 0.05 from 0.604 to 0.618. Note that DDM-Net is not fully end-to-end and uses decoded RGB images as input, which is slow in both decoding and inference stages.

We also conduct experiments on the UCF-101 and HMDB-51 action recognition datasets to validate the effectiveness of our method as in Li et al. (2022). We follow the same settings as E2E (Li et al., 2022) except that we use spatial-channel attention module (SCAM) to process the motion vectors and residuals instead of spatial-channel compressed encoder (SCCE). Our spatial-channel attention module gives more shortcuts to learning refined P-frame features from I-frame with motion vectors and residuals as guidance. Note that our local frames bag is also designed

to capture temporal dependency, which is more suitable for event boundary detection. Thus, it is not applied in the action recognition task. We have two configurations for the action recognition task experiments, one with the same backbone as the event boundary, and the other replaced ResNet-18 with I3D to extract features from motion vectors and residuals, as shown in Table 3. Our method achieves competitive results compared to state-of-the-art methods in compressed domain, *i.e.*, EMV-CNN (Zhang et al., 2016), DTMV-CNN (Zhang et al., 2018), CoViAR (Wu et al., 2018) and DMC-Net (Shou et al., 2019). Compared to E2E (Li et al., 2022), we obtain about **0.5%** improvements on both UCF-101 and HMDB-51 datasets with different backbones. The design principle of spatial-channel attention module (SCAM) is similar to spatial-channel compressed encoder (SCCE) (Li et al., 2022) in that making model generate more discriminative P-frame representations with the guidance of the compressed information (motion vectors and residuals). However, motion vectors and residuals in compressed domain can be very noisy and it is tough to learn a beneficial representation for P-frames, as presented in Fig. 7. Our SCAM uses a gating mechanism to filter out noisy information in both spatial and channel dimensions and uses a bidirectional information flow to refine the origin P-frame feature. In this way, the noisy information from the features of I-frame, motion vectors, and residuals could be effectively and selectively fused together to generate high-quality P-frame features with little overhead. Compared to DMC-Net (Shou et al., 2019), our method can directly learn discriminative features for P-frame with the spatial-channel attention module, which avoids extra optical flow as supervision during training phase.

### 4.3 Failure Case Analysis

Based on our analysis of the Kinetics-GEBD validation split, we have identified two specific categories where our method tends to make mistakes. The first category includes videos with minimal visual changes, as shown in Fig. 8 (above):



**Fig. 6** Example qualitative results on Kinetics-GEBD validation split. Compared with E2E (Li et al., 2022), our method can generate more accurate boundaries which are consistent with ground truth

**Table 2** The evaluation results on the TAPOS validation set with different Rel.Dis. thresholds

Rel.Dis. threshold	0.05	0.1	0.15	0.2	0.25	0.3	0.35	0.4	0.45	0.5	avg
ISBA Ding and Xu (2018)	0.106	0.170	0.227	0.265	0.298	0.326	0.348	0.369	0.382	0.396	0.302
TCN Lea et al. (2016)	0.237	0.312	0.331	0.339	0.342	0.344	0.347	0.348	0.348	0.348	0.330
CTM Huang et al. (2016)	0.244	0.312	0.336	0.351	0.361	0.369	0.374	0.381	0.383	0.385	0.350
TransParser Shao et al. (2020)	0.289	0.381	0.435	0.475	0.500	0.514	0.527	0.534	0.540	0.545	0.474
PC Shou et al. (2021)	0.522	0.595	0.628	0.646	0.659	0.665	0.671	0.676	0.679	0.683	0.642
DDM-Net Tang et al. (2022)	0.604	0.681	0.715	0.735	0.747	0.753	0.757	0.760	0.763	0.767	0.728
Ours	<b>0.618</b>	<b>0.694</b>	<b>0.728</b>	<b>0.749</b>	<b>0.761</b>	<b>0.767</b>	<b>0.771</b>	<b>0.774</b>	<b>0.777</b>	<b>0.780</b>	<b>0.742</b>

Bold numbers indicate the best performance under the corresponding metrics

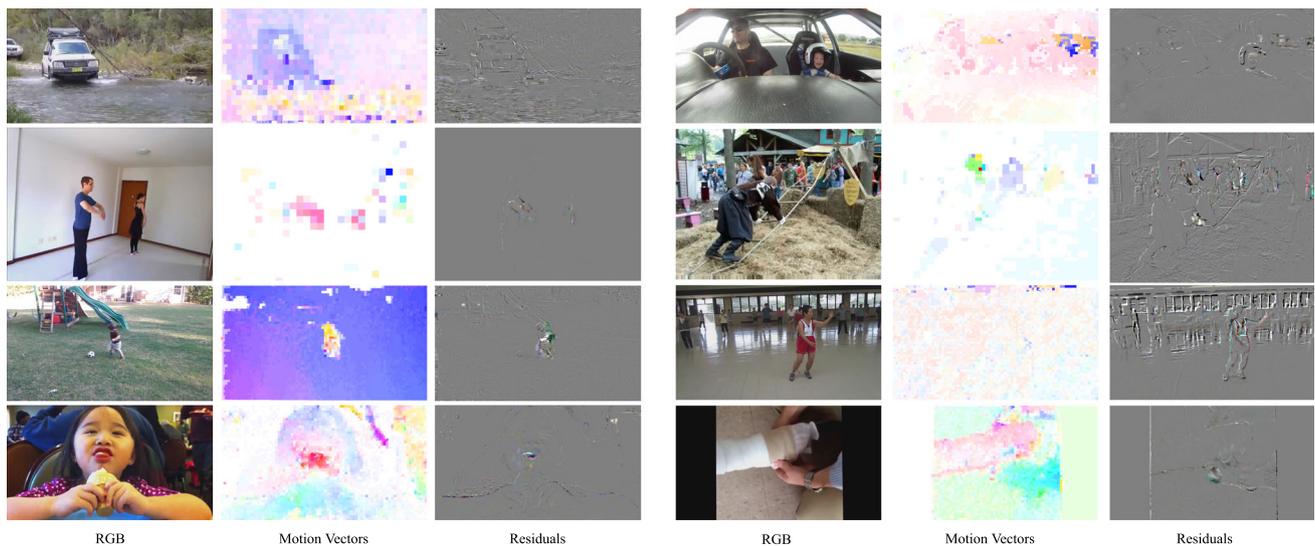
mold ceramics with clay, play the violin, and rotate a suspended object in one direction. These videos pose a challenge for our compression-based method as it struggles to detect event boundaries due to the loss of detailed information in the compressed domain. The second category comprises sports-related activity videos, exemplified by high jump, long jump, and pole vault in Fig. 8 (below). These movements typically involve three events: approach run, takeoff, and landing.

However, our model tends to segment these actions into finer stages, such as dividing the action change in the event of takeoff into multiple events. After thinking and research, we believe that the first type of problems can make up for the missing information by introducing new features such as optical flow. The second type of problem can be solved by introducing external common sense knowledge. For example, before performing event detection on high jump-related

**Table 3** Accuracy on the HMDB-51 and UCF-101 datasets

	HMDB-51	UCF-101
<i>Decoded video based methods (RGB only)</i>		
ResNet-50 He et al. (2016)	48.9	82.3
ResNet-152 He et al. (2016)	46.7	83.4
ActionFlowNet (2-frames) Ng et al. (2018)	42.6	71.0
ActionFlowNet Ng et al. (2018)	56.4	83.9
PWC-Net + CoViAR Sun et al. (2018)	62.2	90.6
TVNet Fan et al. (2018)	71.0	94.5
C3D Tran et al. (2015)	51.6	82.3
Res3D Tran et al. (2017)	54.9	85.8
ARTNet Wang et al. (2018a)	70.9	94.3
MF-Net Chen et al. (2018)	74.6	96.0
S3D Xie et al. (2017)	75.9	96.8
I3D RGB Carreira and Zisserman (2017)	74.8	95.6
<i>Compressed video based methods</i>		
EMV-CNN Zhang et al. (2016)	51.2 (split1)	86.4
DTMV-CNN Zhang et al. (2018)	55.3	87.5
CoViAR Wu et al. (2018)	59.1	90.4
DMC-Net(ResNet-18) Shou et al. (2019)	62.8	90.9
DMC-Net(I3D) Shou et al. (2019)	71.8	92.3
E2E (ResNet-18) Li et al. (2022)	63.3	91.0
E2E (I3D) Li et al. (2022)	72.1	92.5
Ours (ResNet-18)	<b>63.8</b>	<b>91.4</b>
Ours (I3D)	<b>72.7</b>	<b>93.1</b>

Bold numbers indicate the best performance under the corresponding metrics

**Fig. 7** Visualization of the compressed information. The decoded RGB frames, motion vectors, and residuals are presented in different columns

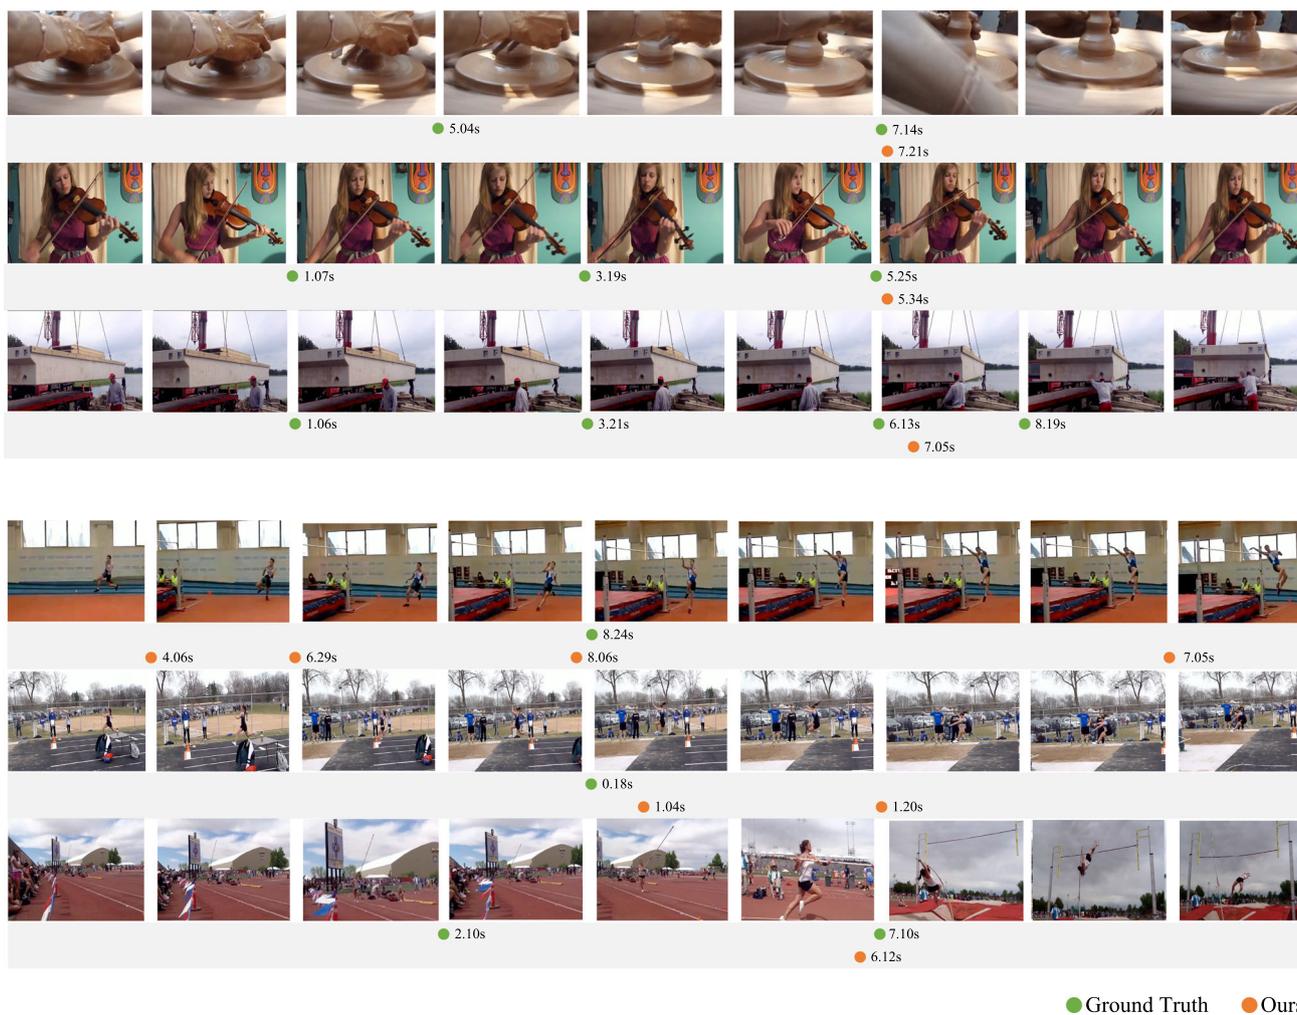


Fig. 8 Examples of failure cases taken from the validation split of the Kinetics-GEBD dataset

videos, let the model learn the knowledge that “high jump is generally divided into three steps: approach run, takeoff and landing”.

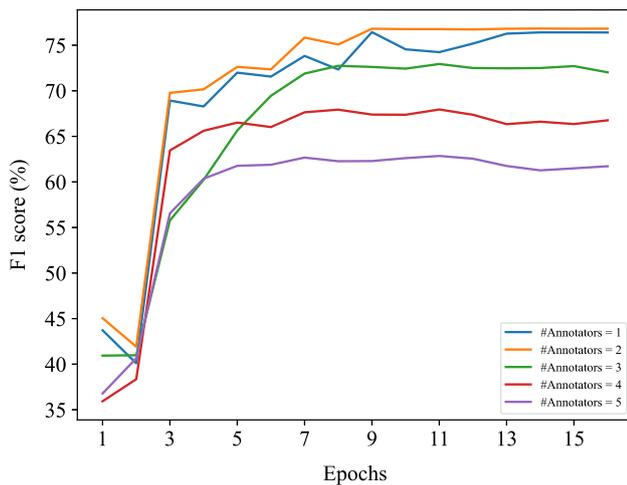
#### 4.4 Ablation Study

In this section, we conduct several ablation studies to demonstrate the effectiveness of different components in the proposed method. All experiments are conducted on the Kinetics-GEBD train split with ResNet50 backbone and tested on a local minval split to reduce the computation cost. The local minval split is constructed from the Kinetics-GEBD validation split by randomly sampling 2, 000 videos. Default settings are marked in gray .

1) *Influence of Number of Annotators:* To remedy the ambiguities of the event boundaries based on human perception, five different annotators are used for each video to label the boundaries based on predefined principles. To analyze the influence of number of annotators used in training

phrase, we conduct 5 experiments by selecting top-1 to top-5 ground truth labels with respect to F1 consistency for training. The results are presented in Fig. 9. We can see that the F1 scores increase when using top-2 ground truth labels compared to only using top-1 labels. This is because top-2 ground truth labels introduce more training samples and the annotation qualities are often consistent with each other. However, when using annotations from more than three annotators, the F1 score is decreased. This can be interpreted as different annotators have labeled event boundaries very differently based on their own subjectivity, which makes model confused and hard to converge. We utilize top-2 ground truth labels as the default setting in our experiments.

2) *Influence of the Backbone for Compressed Information:* A intuitive benefit to use compressed information (*i.e.*, motion vectors and residuals) is that we can learn discriminative features with a very lightweight backbone (*e.g.*, ResNet-18), which saving model forwarding time naturally. However, *can we further improve the F1 score if we use a*



**Fig. 9** Influence of number of annotators. F1 scores get increased when using top-2 ground truth labels compared to only using top-1 labels. When using annotations from more than 3 annotators, the F1 score is decreased

**Table 4** Influence of the backbone for compressed information

Backbone	Rec	Prec	F1	Speed(ms)
ResNet-18	0.799	0.740	0.768	4.5
ResNet-34	0.797	0.738	0.766	7.2
ResNet-50	0.783	0.725	0.753	12.1
ResNet-101	0.781	0.722	0.750	18.3

*more complex backbone for compressed information without considering time cost?* To explore this, we replace residuals and motion vectors' feature extractors  $f_R$  and  $f_M$  in Sect. 3.1 with different backbones. The results are presented in Table 4. Interestingly, we do not observe obvious improvements when using much more power feature extractors and the performances even get slightly decreased when using ResNet-50 and ResNet-101. We attribute this to that motion vectors and residuals are noisy in nature and contains less useful information compared with fully decoded RGB frame, thus a more powerful backbone may overfit simple data distribution and cannot generalize well. We use ResNet-18 as default compressed information backbone for better speed and accuracy trade-off.

3) *Influence of the Temporal Feature Extractor:* After constructing local frames bag for each candidate frame as presented in Sect. 3.2, we are able to extract temporal context information using different modules. We analyze the effectiveness of various of temporal modules, the results are shown in Table 5. Average pooling and max pooling indicate that we accumulate local frames bag  $B^l$  into vector representation with avg-pool operator and max-pool operator, respectively. Thus the following group similarity module and FCN module are not adapted. Since avg-pool operator

**Table 5** Influence of the temporal feature extractor

Method	Rec	Prec	F1
Average pooling	0.728	0.672	0.699
Max pooling	0.720	0.669	0.694
$\emptyset$	0.754	0.731	0.742
LSTM	0.799	0.740	0.768
GRU	0.790	0.732	0.760
Transformer	0.778	0.725	0.751

**Table 6** Influence of the local frames bag size  $k$

Bag size	Rec	Prec	F1
$k = 0$	0.723	0.650	0.685
$k = 2$	0.734	0.748	0.741
$k = 4$	0.752	0.749	0.750
$k = 6$	0.775	0.751	0.763
$k = 8$	0.799	0.740	0.768
$k = 10$	0.743	0.752	0.747
$k = 12$	0.732	0.758	0.745

and max-pool operator only aggregate information linearly, which only can capture little temporal context, we observe obvious performance drop.  $\emptyset$  indicates that the features of local frames bag  $B^l$  are directly fed into following group similarity module and FCN module, which cannot explicitly learn temporal context information, either. As for LSTM and GRU, we observe similar performance. When using more powerful module Transformer (Vaswani et al., 2017b), we find that the F1 score is decreased. We infer that the Transformer may learn deleterious cues for generic event boundary detection. We use LSTM module if not specified.

4) *Influence of the Local Frames Bag Size:* Besides the discriminative features of P-frames, the temporal dependencies are also important to predict the accurate event boundaries. To validate the effectiveness of the temporal contrastive module, we conduct several experiments, shown in Table 6. As shown in Table 6, without the temporal contrastive module (*i.e.*,  $k = 0$ ), the overall accuracy (F1 score) decreased dramatically. After adapting the proposed temporal module, the F1 score improves sharply, *i.e.*, 0.741 vs. 0.685 at  $k = 2$ . To further analyze the effective of different window size in model accuracy, we also perform several experiments with different  $k$  values. Table 6 shows that the recall starts to drop when  $k > 8$ . We believe that it is because larger window size mixes temporal information cross boundaries, resulting in the combination of multiple different predictions and decreasing the recall value. Considering the performance, we set  $k = 8$  in our experiments as the default setting.

**Table 7** Effectiveness of different designs

Component				
With SCAM?	✓	✓	✓	✓
With local frames bag?		✓	✓	✓
With Gaussian smoothing?				✓
F1 score	0.653	0.685	0.760	0.768

5) *Effectiveness of Various Model Design:* As discussed before, the Spatial Channel Attention Module (SCAM) generates discriminative features of P-frames with the guidance of motion vectors and residuals, which is critical for learning a good representation from compressed information. As shown in Table 7, the significant improvement in F1 score is obtained by using local frames bag, *i.e.*, 0.760 vs. 0.685. The constructed local frames bag in Sect. 3.2 provides rich context information for event boundary detection. Compared to E2E (Li et al., 2022) that only uses linear weighted summation representation, our local frame bag is more flexible and adaptive for learning discriminative features leveraging group similarity and FCN. As shown in Table 7, adapting SCAM gives 0.32 improvement. Besides, using the soft labels generated by Gaussian kernel provides further 0.8% absolute improvements. Using the ambiguous “hard labels” disrupt the learning process, which leads to poor convergence. Our soft label strategy effectively solve this issue and speeds up the training process.

6) *Rationality of SCAM module design:* The spatial-channel attention module (SCAM) is designed to refine the feature representations of the P-frames based on compressed information with bidirectional information flow. From the Fig. 4, the SCAM is composed of four modules: Guidance Information Module (GIM), Channel Updating Module (CUM), Spatial Updating Module (SUM), and Weight Updating Module (WUM). GIM processes input information, CUM and SUM introduce channel and spatial features, respectively, while WUM handles feature weighting. To verify the importance of these modules in SCAM, we conduct ablation experiments from two aspects of different inputs and different submodules, as shown in Table 8. From the Table 8 (a), we can find that inputs I-frame feature  $x_I$ , motion vector feature  $x_M^t$  and resized motion vectors  $M^t$  together work best. The results in Table 8 (b) also prove the rationality and effectiveness of the design of each module in SCAM.

7) *Influence of the fusing different modalities:* Compressed video contains three kinds of I-frame feature  $x_I$ , motion vectors  $M^t$  and residuals  $R^t$ . In order to explore the effect of fusing different compressed domain information on model performance, we conducted three sets of comparative experiments and each consisting of two experiments: one with fusion and one without fusion. The experimental results in

**Table 8** The ablation studies of SCAM on the Kinetics-GEBD minval split

(a) Ablation of different inputs.					
$x_I$	$M^t$	$x_M^t$	Rec	Prec	F1
✓	-	-	0.795	0.722	0.757
-	✓	-	0.787	0.728	0.756
-	-	✓	0.795	0.717	0.754
✓	-	✓	0.781	0.73	0.755
-	✓	✓	0.764	0.739	0.752
✓	✓	-	0.781	0.729	0.754
✓	✓	✓	<b>0.799</b>	<b>0.740</b>	<b>0.768</b>
(b) Ablation of different submodules.					
CUM	SUM	WUM	Rec	Prec	F1
✓	-	-	<b>0.809</b>	0.664	0.729
-	✓	-	0.796	0.721	0.757
-	-	✓	0.743	0.354	0.480
✓	-	✓	0.808	0.678	0.737
-	✓	✓	0.789	0.729	0.758
✓	✓	-	0.787	0.731	0.758
✓	✓	✓	0.799	<b>0.740</b>	<b>0.768</b>

Bold numbers indicate the best performance under the corresponding metrics

**Table 9** Ablation of fusing different modalities on the Kinetics-GEBD minval split

$x_I$	$M^t$	$R^t$	Fuse	Rec	Prec	F1
✓	✓			0.709	0.525	0.603
✓	✓		✓	0.682	0.651	0.667
✓		✓		0.809	0.558	0.660
✓		✓	✓	0.801	0.711	0.753
✓	✓	✓		0.753	0.572	0.650
✓	✓	✓	✓	0.799	0.740	0.768

**Table 10** Ablation of sampling different P-frames in each GOP on the Kinetics-GEBD minval split

T	Rec	Prec	F1
1	0.749	<b>0.759</b>	0.754
2	0.800	0.711	0.753
3	0.799	0.740	<b>0.768</b>
4	<b>0.801</b>	0.722	0.760
5	0.797	0.728	0.761

Bold numbers indicate the best performance under the corresponding metrics

the Table 9 show that modality fusion can further improve the performance of the model. Specifically, fusing can bring about a 0.64-0.12 improvement. This proves that fusing

**Table 11** Comparisons with other state-of-the-art methods on LOVEU Challenge. † indicates that the results come from our implementations since the test server is unavailable now. The speed (ms) is computed by averaging per-frame decoding and inference time

Method	Rec	Prec	F1	Speed
†CLA Kang et al. (2021)	0.815	0.768	0.791	90.2
†CASTANET Hong et al. (2021)	0.838	0.732	0.781	93.9
E2E (CSN+R18) Li et al. (2022)	0.813	0.761	0.786	20.4
E2E (R50+R18) Li et al. (2022)	0.751	0.742	0.746	4.7
Ours (CSN+R18)	0.831	0.792	0.812	20.3
Ours (R50+R18)	0.799	0.740	0.768	4.7
Ours (ViT-B+R18)	0.798	0.727	0.761	5.9
Ours (Swin-T+R18)	0.783	0.729	0.755	5.1

information from different compressed domains is beneficial for reconstructing the P-frame.

8) *Influence of the sampling different P-frames in each GOP*: In the MPEG-4 encoding format, each GOP usually contains 1 I-frame and 11 P-frames. If all P-frames are used as input, it will increase the amount of computation and slow down the inference speed. Therefore, we adopted the strategy of sampling part of the P-frames. However, too few P-frames will lose a lot of information and affect the performance of the model. In order to determine the optimal number of P-frames, we conducted experiments as shown in Table 10. Specifically, we experimented with sampling 1 frame to 5 frames respectively, and the experimental results show that sampling 3 frames of P-frames works best. In addition, the result of sampling 5 frames is slightly reduced, 0.761 VS 0.768, which proves that too many P-frames will only bring redundant information and cannot further improve model performance.

#### 4.5 Comparisons with State-of-the-Arts on LOVEU Challenge

We also compare the proposed method with the state-of-the-art methods at CVPR'21 LOnG-form VidEo Understanding (LOVEU) Challenge<sup>2</sup> as shown in Table 11. The winners' solutions are complicated and running slowly, *e.g.*, CLA (Kang et al., 2021) relies on pre-extracted features and uses global similarity matrix which cannot scale well, CASTANET (Hong et al., 2021) is not fully end-to-end and introduces redundant computations between nearby frames. E2E (Li et al., 2022) remedies this by using motion vectors and residuals in the compressed domain and achieves competitive results while running  $20\times$  faster than CLA (Kang et al., 2021). In addition, our method obtains absolute improvements of 2.3% and 4.2% compared to the preliminary version E2E (Li et al., 2022) when using CSN (Tran et al., 2019) and ResNet50 as backbones, respectively, while running almost at the same speed.

<sup>2</sup> <https://sites.google.com/view/loveucvpr21>.

#### 4.6 Limitation of the Compressed Videos

Working with compressed videos has several shortcomings that need to be considered, including potential loss of information during compressing videos into compressed streams. Since compressed data is a condensed version of uncompressed data, some data may be lost or altered, leading to inaccurate results and compromising data quality. Additionally, compression artifacts present a challenge when dealing with compressed data, as certain elements of the original data may be removed or distorted, particularly in visual tasks like image recognition. Finally, processing compressed data is generally more complex than uncompressed data because it must first be decompressed before analysis or manipulation. This added step can increase processing time and require specialized tools and techniques for accurate results. In addition, the limitations of directly using compressed video compared to decoding RGB video frames is that existing deep learning-based models are specifically designed to handle RGB videos. Therefore, directly applying these architecture may lead to suboptimal performance. A future direction is to develop a series of tailored network architectures directly suitable for compressed vision tasks.

Despite these challenges, working with compressed data remains crucial for research and development in various applications.

#### 5 Conclusion and Future work

In this work, we propose an end-to-end compressed video representation learning method for GEBD. Specifically, we convert the video input into successive frames and use the Gaussian kernel to preprocess the annotations. Meanwhile, we design a spatial-channel attention module (SCAM) to make full use of the motion vectors and residuals to learn discriminative feature representations for P-frames with bidirectional information flow. After that, we propose a temporal contrastive module that uses local frames bag as representation to model the temporal dependency between frames

and generate accurate event boundaries with group similarity. Extensive experiments conducted on the Kinetics-GEBD and TAPOS datasets demonstrate that the proposed method performs favorably against the state-of-the-art methods.

While our method has shown promising results, there is still room for improvement. Currently, the model leverages only the high-level semantic information extracted by the backbone network. However, for event boundary detection tasks, the importance of low-level detail information cannot be overlooked. Thus, an improved approach could involve the backbone network initially extracting multiscale features. Subsequently, a multiscale feature fusion module could be used to process these multiscale features. Regarding future research on further improvements, there are three possible directions: (1) The first direction for future research involves expanding current SCAM modules to support additional encoding formats beyond the current MPEG-4 standard. While the current SCAM can only handle one encoding format, extending it to include a range of general encoding formats would greatly enhance its utility and practicality. (2) The second direction involves enhancing the temporal module to allow independent and flexible selection of temporal modules based on specific scenarios. Currently, the temporal module is limited in its ability to select the most appropriate temporal module for different scenarios. (3) The last possible direction is to incorporate more information, such as audio and knowledge graphs. Audio information can provide a new basis for judgment and can assist the model in determining event boundary points. The knowledge information contained in the knowledge graph can help the model understand events to better determine the beginning and end of the boundary.

**Acknowledgements** Libo Zhang was supported by the Key Research Program of Frontier Sciences, CAS, Grant No. ZDBS-LY-JSC038, High-end Research Institutions Innovation Special Funds introduced by Zhongshan Science and Technology Bureau (No.2020AG011) and Youth Innovation Promotion Association, CAS (2020111). Heng Fan and his employer received no financial support for the research, authorship, and/or publication of this article.

**Data Availability Statement** The data that support the findings of this study are openly available in “GEBD” at <https://github.com/StanLei52/GEBD>, which are included in this published article (Shou et al., 2021).

## References

- Alwassel, H., Heilbron, F. C., & Ghanem, B. (2018). Action search: Spotting actions in videos and its application to temporal action localization. In: *ECCV*.
- Arnab, A., Deghani, M., Heigold, G., et al. (2021). Vivit: A video vision transformer. In: *ICCV*.
- Caba Heilbron, F., Barrios, W., Escorcia, V., et al. (2017). Scc: Semantic context cascade for efficient action detection. In: *CVPR*.
- Carreira, J., & Zisserman, A. (2017). Quo Vadis, action recognition? A new model and the kinetics dataset. In: *CVPR*.
- Chao, Y. W., Vijayanarasimhan, S., Seybold, B., et al. (2018). Rethinking the faster R-CNN architecture for temporal action localization. In: *CVPR*.
- Chen, Y., Kalantidis, Y., Li, J., et al. (2018). Multi-fiber networks for video recognition. In: *ECCV*.
- Deng, J., Dong, W., Socher, R., et al. (2009). Imagenet: A large-scale hierarchical image database. In: *CVPR*.
- Ding, L., & Xu, C. (2018). Weakly-supervised action segmentation with iterative soft boundary assignment. In: *CVPR*.
- Dosovitskiy, A., Beyer, L., Kolesnikov, A., et al. (2021a). An image is worth 16x16 words: Transformers for image recognition at scale. In: *ICLR*.
- Dosovitskiy, A., Beyer, L., Kolesnikov, A., et al. (2021b). An image is worth 16x16 words: Transformers for image recognition at scale. In: *ICLR*.
- Fan, H., Xiong, B., Mangalam, K., et al. (2021). Multiscale vision transformers.
- Fan, L., Huang, W., Gan, C., et al. (2018). End-to-end learning of motion representation for video understanding. In: *CVPR*.
- Feichtenhofer, C., Pinz, A., & Zisserman, A. (2016). Convolutional two-stream network fusion for video action recognition. In: *CVPR*.
- Feichtenhofer, C., Fan, H., Malik, J., et al. (2019). Slowfast networks for video recognition. In: *ICCV*.
- Gall, D. L. (1991). MPEG: A video compression standard for multimedia applications. *Communications of the ACM*, 34(4), 46–58.
- Geirhos, R., Jacobsen, J., Michaelis, C., et al. (2020). Shortcut learning in deep neural networks. *Nature Machine Intelligence*, 2(11), 665–673.
- He, K., Zhang, X., Ren, S., et al. (2016). Deep residual learning for image recognition. In: *CVPR*.
- Hong, D., Li, C., Wen, L., et al. (2021). Generic event boundary detection challenge at CVPR 2021 technical report: Cascaded temporal attention network (CASTANET). arXiv.
- Hu, J., Shen, L., & Sun, G. (2018). Squeeze-and-excitation networks. In: *CVPR*.
- Huang, D., Fei-Fei, L., Niebles, J. C. (2016). Connectionist temporal modeling for weakly supervised action labeling. In: *ECCV*.
- Huang, L., Liu, Y., Wang, B., et al. (2021). Self-supervised video representation learning by context and motion decoupling. In: *CVPR*.
- Ji, S., Xu, W., Yang, M., et al. (2013). 3d convolutional neural networks for human action recognition. *IEEE transactions on pattern analysis and machine intelligence*, 35(1), 221–231.
- Kang, H., Kim, J., Kim, K., et al. (2021). Winning the CVPR’2021 kinetics-GEBD challenge: Contrastive learning approach. arXiv.
- Kuehne, H., Jhuang, H., Garrote, E., et al. (2011). HMDB: A large video database for human motion recognition. In: *ICCV*.
- Lea, C., Reiter, A., Vidal, R., et al. (2016). Segmental spatiotemporal CNNs for fine-grained action segmentation. In: *ECCV*.
- Lea, C., Flynn, M. D., Vidal, R., et al. (2017). Temporal convolutional networks for action segmentation and detection. In: *CVPR*.
- Li, C., Wang, X., Wen, L., et al. (2022). End-to-end compressed video representation learning for generic event boundary detection. In: *CVPR*.
- Li, J., Wei, P., Zhang, Y., et al. (2020). A slow-i-fast-p architecture for compressed video action recognition. In: *ACM MM*.
- Lin, T., Zhao, X., & Shou, Z. (2017). Single shot temporal action detection. In: *ACM MM*.
- Lin, T., Liu, X., Li, X., et al. (2019a). BMN: Boundary-matching network for temporal action proposal generation. In: *ICCV*.
- Lin, T., Liu, X., Li, X., et al. (2019b). BMN: Boundary-matching network for temporal action proposal generation. In: *ICCV*.
- Liu, Z., Lin, Y., Cao, Y., et al. (2021). Swin transformer: Hierarchical vision transformer using shifted windows. In: *Proceedings of the IEEE/CVF international conference on computer vision*, pp 10012–10022.
- Liu, Z., Ning, J., Cao, Y., et al. (2022). Video Swin transformer.

- Long, F., Yao, T., Qiu, Z., et al. (2019). Gaussian temporal awareness networks for action localization. In: *CVPR*.
- Long, J., Shelhamer, E., & Darrell, T. (2015). Fully convolutional networks for semantic segmentation. In: *CVPR*.
- Ma, S., Sigal, L., & Sclaroff, S. (2016). Learning activity progression in lstms for activity detection and early detection. In: *CVPR*.
- Ng, J. Y., Choi, J., Neumann, J., et al. (2018). Actionflownet: Learning motion representation for action recognition. In: *WACV*.
- Ni, B., Yang, X., & Gao, S. (2016). Progressively parsing interactional objects for fine grained action detection. In: *CVPR*.
- Paszke, A., Gross, S., Massa, F., et al. (2019). Pytorch: An imperative style, high-performance deep learning library. In: *NeurIPS*.
- Rai, A. K., Krishna, T., Dietmeier, J., et al. (2021). Discerning generic event boundaries in long-form wild videos. arXiv.
- Richard, A., & Gall, J. (2016). Temporal action detection using a statistical language model. In: *CVPR*.
- Shao, D., Zhao, Y., Dai, B., et al. (2020). Intra- and inter-action understanding via temporal action parsing. In: *CVPR*.
- Shou, M. Z., Lei, S. W., Wang, W., et al. (2021). Generic event boundary detection: A benchmark for event segmentation. In: *ICCV*.
- Shou, Z., Lin, X., Kalantidis, Y., et al. (2019). Dmc-net: Generating discriminative motion cues for fast compressed video action recognition. In: *CVPR*.
- Simonyan, K., & Zisserman, A. (2014). Two-stream convolutional networks for action recognition in videos. In: *NIPS*.
- Soomro, K., Zamir, A. R., & Shah, M. (2012). UCF101: A dataset of 101 human actions classes from videos in the wild. arXiv.
- Sun, D., Yang, X., Liu, M., et al. (2018). Pwc-net: Cnns for optical flow using pyramid, warping, and cost volume. In: *CVPR*.
- Tang, J., Liu, Z., Qian, C., et al. (2022). Progressive attention on multi-level dense difference maps for generic event boundary detection. In: *CVPR*.
- Taylor, G. W., Fergus, R., LeCun, Y., et al. (2010). Convolutional learning of spatio-temporal features. In: *ECCV*.
- Tran, D., Bourdev, L. D., Fergus, R., et al. (2015). Learning spatiotemporal features with 3d convolutional networks. In: *ICCV*.
- Tran, D., Ray, J., Shou, Z., et al. (2017). Convnet architecture search for spatiotemporal feature learning. arXiv.
- Tran, D., Wang, H., Feiszli, M., et al. (2019). Video classification with channel-separated convolutional networks. In: *ICCV*.
- Varol, G., Laptev, I., & Schmid, C. (2018). Long-term temporal convolutions for action recognition. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 40(6), 1510–1517.
- Vaswani, A., Shazeer, N., Parmar, N., et al. (2017a). Attention is all you need. In: *NIPS*.
- Vaswani, A., Shazeer, N., Parmar, N., et al. (2017b). Attention is all you need. In: *NIPS*.
- Wang, L., Li, W., Li, W., et al. (2018a). Appearance-and-relation networks for video classification. In: *CVPR*.
- Wang, S., Lu, H., & Deng, Z. (2019). Fast object detection in compressed video. In: *ICCV*.
- Wang, X., Girshick, R. B., Gupta, A., et al. (2018b). Non-local neural networks. In: *CVPR*.
- Woo, S., Park, J., Lee, J., et al. (2018). CBAM: Convolutional block attention module. In: *ECCV*.
- Wu, C., Zaheer, M., Hu, H., et al. (2018). Compressed video action recognition. In: *CVPR*.
- Xie, S., Sun, C., Huang, J., et al. (2017). Rethinking spatiotemporal feature learning for video understanding. arXiv.
- Yu, Y., Lee, S., Kim, G., et al. (2021). Self-supervised learning of compressed video representations. In: *ICLR*.
- Yuan, Z., Stroud, J. C., Lu, T., et al. (2017). Temporal action localization by structured maximal sums. In: *CVPR*.
- Zeiler, M. D., & Fergus, R. (2014). Visualizing and understanding convolutional networks. In: *ECCV*.
- Zhang, B., Wang, L., Wang, Z., et al. (2016). Real-time action recognition with enhanced motion vector CNNs. In: *CVPR*.
- Zhang, B., Wang, L., Wang, Z., et al. (2018). Real-time action recognition with deeply transferred motion vector CNNs. *IEEE Transactions on Image Processing*, 27(5), 2326–2339.
- Zhang, H., Hao, Y., & Ngo, C. (2021). Token shift transformer for video classification. In: *ACM MM*.
- Zhao, P., Xie, L., Ju, C., et al. (2020). Bottom-up temporal action localization with mutual regularization. In: *ECCV*.
- Zhao, Y., Xiong, Y., Wang, L., et al. (2017). Temporal action detection with structured segment networks. In: *ICCV*.

**Publisher's Note** Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.

Springer Nature or its licensor (e.g. a society or other partner) holds exclusive rights to this article under a publishing agreement with the author(s) or other rightsholder(s); author self-archiving of the accepted manuscript version of this article is solely governed by the terms of such publishing agreement and applicable law.