

For referencing this work, please cite

T. Li, Z. Sun and X. Xiao, "Unsupervised Modality-Transferable Video Highlight Detection With Representation Activation Sequence Learning," IEEE Transactions on Image Processing, 2024.

```
@article{li2024highlight,  
  title={Unsupervised Modality-Transferable Video Highlight Detection with  
  Representation Activation Sequence Learning},  
  author={Li, Tingtian and Sun, Zixun and Xiao, Xinyu},  
  journal={IEEE Transactions on Image Processing},  
  year={2024}  
}
```

Unsupervised Modality-Transferable Video Highlight Detection with Representation Activation Sequence Learning

Tingtian Li^{1*}, Zixun Sun¹, and Xinyu Xiao¹

Abstract—Identifying highlight moments of raw video materials is crucial for improving the efficiency of editing videos that are pervasive on internet platforms. However, the extensive work of manually labeling footage has created obstacles to applying supervised methods to videos of unseen categories. The absence of an audio modality that contains valuable cues for highlight detection in many videos also makes it difficult to use multimodal strategies. In this paper, we propose a novel model with cross-modal perception for unsupervised highlight detection. The proposed model learns representations with visual-audio level semantics from image-audio pair data via a self-reconstruction task. To achieve unsupervised highlight detection, we investigate the latent representations of the network and propose the representation activation sequence learning (RASL) module with k -point contrastive learning to learn significant representation activations. To connect the visual modality with the audio modality, we use the symmetric contrastive learning (SCL) module to learn the paired visual and audio representations. Furthermore, an auxiliary task of masked feature vector sequence (FVS) reconstruction is simultaneously conducted during pretraining for representation enhancement. During inference, the cross-modal pretrained model can generate representations with paired visual-audio semantics given only the visual modality. The RASL module is used to output the highlight scores. The experimental results show that the proposed framework achieves superior performance compared to other state-of-the-art approaches.

Index Terms—unsupervised, highlight detection, multimodal, representation activation sequence.

I. INTRODUCTION

WITH the prevalence of video media on internet platforms, there is a growing demand for automatically extracting highlight segments from sheer volumes of footage for quick browsing. Especially in the era of flourishing short videos, fast production means both timely and reducing labor costs. Thus, as a technique that automatically locates attractive segments, highlight detection has attracted increasing attention from researchers in recent years [1]–[13].

The majority of highlight detection approaches are supervised [1], [3], [7], [8]. Given videos with frame-level labels, supervised methods learn from frame features and temporally localize highlight segments. However, as frame-level labeling is labor intensive, datasets with frame-level labels are unable to include vast video categories. Supervised approaches thus become domain specific and show a weak ability to adapt to wild videos of unseen categories. Another branch of methods corresponds to weakly supervised methods [2], [4], [9], [10].

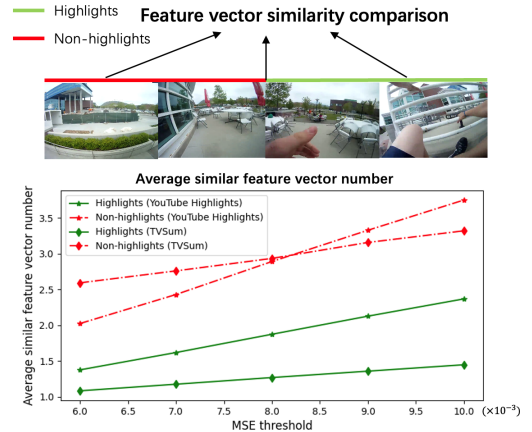



Fig. 1. We evaluate the similarity of feature vectors extracted by the feature extractor [14] from videos in the YouTube Highlights and TVSum datasets and show the number of feature vectors that exhibit similarity to others within the videos. We calculate the mean squared error (MSE) to assess the similarity between the target feature vector and others within a window (30 seconds). The sampling rate is set to 5 frames/second. If the MSE is below the threshold, we consider the two vectors to be similar. The results show that highlights have fewer similar feature vectors.

Weakly supervised methods do not rely on expensive frame-level labels. Usually, they are heuristic and learn from specific metadata, e.g., video duration [4] and video categories [2], [9], [10]. However, metadata still require special preparation and result in weak generalization. The association of visual and audio modalities has been employed in different applications [15]–[18]. The methods [9], [10] also aware that the audio modality is useful for highlight detection. However, the audio modality is unavailable in some situations. For example, a camera may be far from the sound source of an attractive scene, or the original sound may be obscured by irrelevant environmental noise or background music.


In this paper, we propose a novel unsupervised cross-modal highlight detection framework. Motivated by the saliency detection task, where salient regions that attract our attention in an image are distinct from its overall appearance [19], [20], we consider highlight frames, which attract our attention, to be highly likely to exhibit semantic distinctiveness in the footage compared to ordinary frames. To verify this, we analyze the similarity of feature vectors to others within the videos in the YouTube Highlights [3] and TVSum [21] datasets. We present the average numbers of similar feature vectors to others in Fig. 1. The results show that highlights have

¹Tencent, China.


*Email: tingtian.li@outlook.com.



(a) The training and inference flowchart of the weakly supervised visual-audio-based methods [9], [10]



(b) The pretraining flowchart of the proposed unsupervised method



(c) The inference flowchart of the proposed unsupervised method

Fig. 2. Flowcharts of the weakly supervised visual-audio-based methods [9], [10] and the proposed unsupervised method. [9], [10] need both the visual and the audio modalities as the input in the training and inference processes. The proposed method requires the visual and audio modalities as the input during pretraining. During inference, only the visual modality is needed to obtain representations with visual-audio-level semantics from the pretrained model, and the RASL module is used to output the highlight scores.

fewer similar feature vectors, indicating that they are more distinct than non-highlights are. From the view of information theory, we argue that highlights contain more meaningful information than ordinary snippets because they are more distinct. Inspired by this, we construct a network with temporal activations and propose the representation activation sequence learning (RASL) module to learn the significant representation activations through self-reconstruction during pretraining. The original purpose of training deep neural networks for signal reconstruction was to solve inverse problems [22]–[24]. Recently, self-reconstruction has become frequently used in self-supervised model pretraining for improving the performance of downstream tasks, such as image and video classification, object detection, and semantic segmentation [25]–[29]. However, the focus of these works is primarily on obtaining effective representations from intermediate layers of pretrained networks for downstream tasks, and the use of representation activations of self-reconstruction networks for direct application is often overlooked. In this study, we detect highlight moments using the representation activations of the pretrained network in a self-supervised manner. The proposed RASL module promotes larger values in the representation activations for distinct highlight snippets for better reconstruction, as they are more challenging to reconstruct than redundant

ordinary frames are. This enables the model to infer highlight moments from the representation activation sequence. Specifically, the RASL module aggregates the top- k representation activations with the highest responses and guides them to be more distinguishable. As the activations selected by the fixed hyperparameter k probably incorporate outliers, we propose k -point contrastive learning to suppress the outliers. The framework contains two branches of autoencoders operating on the visual and audio modalities. We use the symmetric contrastive learning (SCL) module to establish the connection between the paired visual and the audio representations via modal contrastive learning [30]. During inference, the visual branch can generate representations with paired visual-audio semantics and output highlight scores via the RASL module without the need for the audio modality. Figure 2 shows the flowcharts of the weakly supervised visual-audio methods [9], [10] and the proposed unsupervised method. [9], [10] require both visual frames and audio waves as inputs during training and inference. In contrast, the proposed method requires both modalities only during pretraining. During inference, given only the visual frames, the pretrained cross-modal network and the RASL module are used to output the highlight scores. To our knowledge, the proposed method is the first visual-audio-based method that does not require any audio as the inference input. Note that weakly supervised multimodal approaches [9], [10] also need video category labels, whereas the pretraining of the proposed methods is unsupervised without the need for any labels or metadata. We also utilize multitask learning (MTL) to improve the framework performance. MTL is an approach for improving the performance of a network by using the domain information contained in another simultaneously trained task. Here, we adopt masked feature vector sequence (FVS) reconstruction, which has been shown to improve representation performance in self-supervised learning studies [26], [27], [31], as an auxiliary task.

To summarize, the main contributions of our proposed method are as follows:

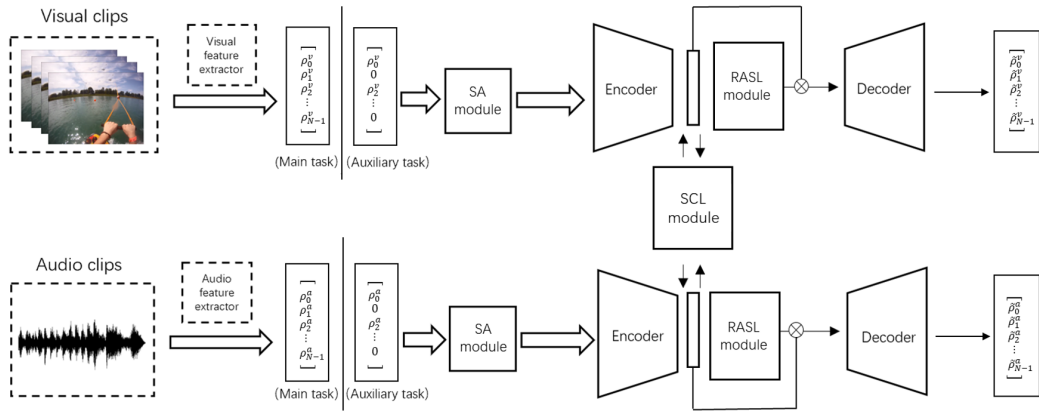
1) We propose a novel unsupervised cross-modal highlight detection framework. We use a self-reconstruction task to pretrain the model and build the visual-audio connection. Given only visual frames, the pretrained model can generate representations with cross-modal semantics according to the modal connection. It can directly infer highlights of wild videos without further training.

2) We propose the RASL module with k -point contrastive learning to guide the activations of significant representations to be more distinguishable and suppress the incorporated activation outliers without the need for frame-level annotated labels. During inference, the RASL module outputs the highlight scores from the learned representations.

3) We use the SCL module to build the modal connection with cross-modal contrastive learning. The SCL module enables the visual branch to generate representations with visual-audio-level semantics from learning to pair representations.

4) We build an MTL framework and add an auxiliary task of masked FVS reconstruction. The auxiliary task enhances the representations and improves the highlight detection performance.

Pre-training:



Inference:




Fig. 3. The framework of the proposed highlight detection method. During pretraining, we input the visual and audio clips into the two branches and extract their feature vectors to compose the visual and audio FVSs. Then, we enhance the FVSs of the two modalities via the SA modules. After that, the self-attended FVSs are fed to the autoencoders for self-reconstruction. The significant activations are learned from the RASL module. The paired visual-audio representations are learned through the SCL module. The auxiliary task of masked FVS reconstruction is conducted to improve the performance of the main highlight detection task. During inference, we use only the cross-modal pretrained visual branch and the RASL module to output highlight scores.

II. RELATED WORK

A. Video Highlight Detection

Early highlight detection works focused on solving the sporting video clipping problem [32]–[34]. Some later works localized attractive segments for social media production [3], [5]. [1], [3], [5], [7], [35], [36] formulated attractive segment localization as a segment ranking problem. They trained ranking networks using contrastive learning and assigned higher scores to highlight segments. These methods are supervised and require the availability of frame-level annotated labels. However, manually labeling highlight moments from footage is labor intensive. As a result, the number of annotated videos is limited, and supervised methods easily fall into domain-specific problems. To avoid the limited frame-level annotation problem, weakly supervised highlight detection approaches have been studied [2], [4], [9], [10]. Without frame-level annotation, weakly supervised approaches detect highlight segments with the assistance of metadata. [2] distinguished highlight frames from FVS reconstruction errors of specific video categories. [4] demonstrated that videos of shorter duration have higher probabilities of being highlight clips. This fact implicitly supervises the network to prefer segments from shorter videos. [9], [10] reported that the audio modality can assist in detecting highlight snippets. [10] proposed to adopt a ranking loss and uses multiple submodules to fuse the modalities. [9] utilized a hierarchical temporal encoder and a multimodal tensor fusion mechanism to fuse modalities.

However, existing visual-audio-based methods [9], [10] require visual frames and audio waves for both training and inference. They cannot handle situations where the original sound is lost or strongly disturbed. Moreover, although [2], [4], [9], [10] avoid frame-level annotation, they still need to prepare the dataset with metadata in specific ways, e.g., video topics [2], [9], [10] and annotations of video durations [4]. In contrast, our proposed method learns from the unsupervised self-reconstruction task without requiring any metadata. During inference, the proposed method requires only visual frames but can generate representations with learned visual-audio semantics to infer the highlight segments of videos in the wild.

B. Video Summarization

Video summarization is a task similar to video highlight detection. The goal is to integrate the important frames and generate a compact summary of a given video [37]–[48]. [44] proposed to learn the summaries from category-specific videos. [46] detected important frames using a probabilistic model for diverse sequential subset selection. [45] utilized long short-term memory (LSTM) [49] to model the variable-range dependencies for video summarization. [40] located the segments that cooccur most frequently across collected videos using a keyword. [41] learned semantic matching between the generated summaries and web videos. Several other methods [37]–[39] utilized the generative adversarial network (GAN) [50]–[52] to regularize the summarizer by validating the

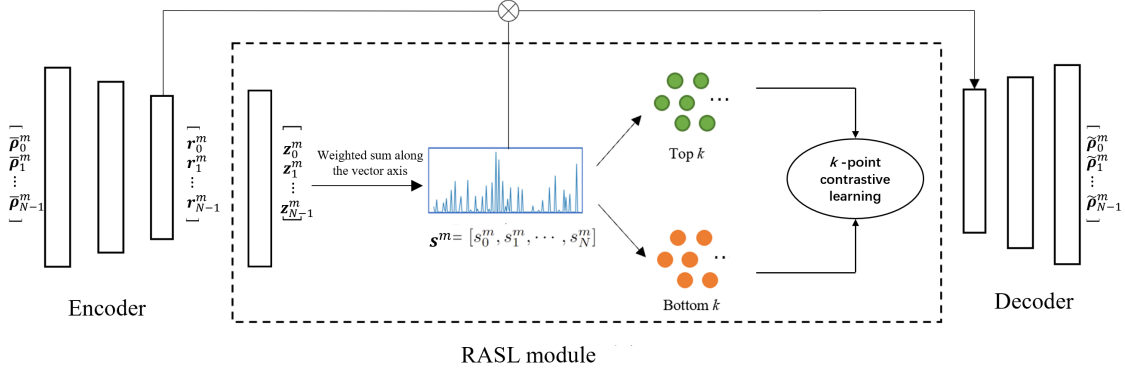


Fig. 4. The illustration of the proposed RASL module, which is shown in the dotted box. The representation vectors r^m output from the encoder are sent to the module, which learns s^m via k -point contrastive learning. Then, the product of s^m and z^m is fed to the decoder for FVS reconstruction.

consistency between the estimated summaries and the video features.

C. Multitask Learning and Masked Signal Reconstruction

The MTL paradigm aims to leverage useful knowledge contained in multiple related tasks that are trained simultaneously to improve the performance and data efficacy of all the tasks [53]–[55]. MTL is similar to transfer learning [56] in that it utilizes the information in pretrained tasks to improve the performance of downstream tasks. However, MTL trains multiple tasks simultaneously, leveraging their shared knowledge to improve model generalizability and performance. In this work, we adopt the MTL strategy for performance improvement.

Masked signal reconstruction, as a popular task of self-supervised learning, has been applied to different modal signals, e.g., vision [25], [26], language [31], and audio [27], [57]. This approach predicts the original signals from these intentional distortions and shares the learned knowledge with other tasks. Specifically, [26] demonstrated that autoencoders with downsampling and upsampling abilities are naturally suitable for self-supervised reconstruction of mask signals. Considering that the proposed framework also contains an autoencoder structure, we select masked FVS reconstruction as our auxiliary task.

III. METHOD

Fig. 3 shows the framework of our proposed method. During pretraining, we extract the FVSs of the visual and audio modalities and feed them into the autoencoders with the RASL modules for self-reconstruction. The representations with visual-audio semantics are learned by the SCL module. The auxiliary task of masked FVS reconstruction is used to improve the highlight detection performance. During inference, we use the cross-modal pretrained visual branch and the RASL module to output the highlight scores. In this section, we will illustrate the details of the proposed method as follows.

A. Representations of The Visual and The Audio Modalities

Our method learns from visual and audio representations to detect video highlights during pretraining. We first transform

the raw visual and audio signals into FVSs, which are also the targets of self-reconstruction. Here, we denote the visual and the audio modalities as v and a , respectively. For the visual modality v , we first extract the frames of the target video at a fixed interval, resulting in N frames. We map the extracted i -th frame to the feature vector $\rho_i^v \in R^{1 \times d_v}$ by the feature extractor [14], where d_v is the visual feature vector dimension. We denote the visual FVS as $\rho^v = [\rho_0^v, \rho_1^v, \dots, \rho_{N-1}^v]$. For the audio modality a , we extract the feature vectors using the feature extractor [57]. Since the audio signal temporal density is much greater than the visual signal temporal density, we further downsample the audio FVS via average pooling for network parameter reduction. We denote the obtained audio FVS as $\rho^a = [\rho_0^a, \rho_1^a, \dots, \rho_{N-1}^a]$, where $\rho_i^a \in R^{1 \times d_a}$ indicates the i -th audio feature vector with dimension d_a .

B. Intramodal Self-Attention

Highlighting moments within a video are not independent. They depend on the video structure of the contents and are semantically related. Here, we use the self-attention (SA) module [58] to model this relationship in our representation generation. The SA module learns temporal weights according to the feature vectors at other temporal locations. It effectively models the temporal dependency of the feature vectors. Specifically, given an FVS $\rho^m = [\rho_0^m, \rho_1^m, \dots, \rho_{N-1}^m]$, we project ρ^m to the value sequence V^m , the query sequence Q^m and the key sequence K^m , where $\rho_i^m \in R^{1 \times d_\rho^m}$ is the i th feature vector of the modality m with dimensions d_ρ^m and $m \in \{v, a\}$. We model the dependency by the weight V^m , which is calculated by the multiplication of Q^m and K^m . The calculation of the SA mechanism is shown as follows:

$$\begin{aligned} O_{\rho^m} &= \text{softmax}\left(\frac{Q^m(K^m)^T}{\sqrt{d_\rho^m}}\right)V^m \\ &= \text{softmax}\left(\frac{\rho^m W^{Q,m}(\rho^m W^{K,m})^T}{\sqrt{d_\rho^m}}\right)\rho^m W^{V,m}, \end{aligned} \quad (1)$$

$$\bar{\rho}^m = F^m(\rho^m) = \rho^m + O_{\rho^m} W^{O,m}, \quad (2)$$

where O_{ρ^m} is the self-attended FVS and F^m represents the SA process for modality m . $W^{Q,m} \in R^{d_\rho^m \times d_\rho^m}$, $W^{K,m} \in$

$R_{\rho}^{d_{\rho}^m \times d_{\rho}^m}$, $W^{V,m} \in R^{d_{\rho}^m \times d_{\rho}^m}$ and $W^{O,m} \in R^{d_{\rho}^m \times d_{\rho}^m}$ are four learnable projection matrices. $\sqrt{d_{\rho}^m}$ is a scaling factor. The temporal dependency of the FVS is modeled by the second part of Eq. (2). We combined O_{ρ^m} with the original FVS ρ^m to obtain the enhanced FVS $\bar{\rho}^m = [\bar{\rho}_0^m, \bar{\rho}_1^m, \dots, \bar{\rho}_{N-1}^m]$.

C. Representation Activation Selection Learning

Fig. 1 shows that highlight snippets are often more distinct from overall videos while ordinary frames are more semantically similar and redundant. We argue that highlights contain more meaningful information, as in information theory, the information content of an event is measured by $-\log(P(A))$, where the information content increases as the probability $P(A)$ of an event A decreases. Compared to the redundant ordinary FVS, the highlight FVS with distinct information is more difficult to reconstruct. Inspired by this, we build a network with temporal representation activations for highlight detection. During pretraining, the network tends to learn high activations on distinct representations to minimize the reconstruction error. To guide the activations corresponding to highlight moments to be more distinguishable, we propose the RASL module, which learns larger activations on the significant representation vectors. Thus, we can recognize the highlight moments via the values of representation activations during inference. The structure of the RASL module is shown in Fig. 4. The module learns the representation activation sequence $s^m = [s_0^m, s_1^m, \dots, s_{N-1}^m]$ of modality m from the representation vector sequence r^m . The representation vector sequence is obtained by

$$r^m = [r_0^m, r_1^m, \dots, r_{N-1}^m] = E^m(\bar{\rho}^m) = E(F^m(\rho^m)), \quad (3)$$

where E^m is the encoder of the autoencoder of modality m . The i -th activation for modality m is calculated as follows:

$$\begin{aligned} s_i^m &= \sum_c w_c^m z_i^m[c] \\ &= \sum_c w_c^m \psi^m(r_i^m)[c], \end{aligned} \quad (4)$$

where ψ^m represents a convolutional layer. $z_i^m = \psi^m(r_i^m)$ is the representation vector processed by the convolutional layer ψ^m . The representation activation s_i^m is obtained by the weighted sum of z_i^m along the vector axis. c is the index of the vector axis. w_c^m is the result of the weighted sum process. We implement the weighted sum process using a bias-free fully connected layer, where w_c^m serves as a weight of this layer. The products of the sequences s^m and r^m are then sent to the decoder for FVS reconstruction. We reconstruct the FVS of modality m by minimizing the reconstruction loss

$$\begin{aligned} L_e^m &= |\rho^m - \bar{\rho}^m| \\ &= |\rho^m - D^m(s^m \cdot r^m)| \\ &= |\rho^m - D^m([s_0^m \cdot r_0^m, s_1^m \cdot r_1^m, \dots, s_{N-1}^m \cdot r_{N-1}^m])|, \end{aligned} \quad (5)$$

where D^m is the decoder of modality m . $\bar{\rho}^m$ is the reconstructed FVS of modality m . In Eq. (5), the representation activation s_i^m serves as an attention weight for r_i^m . Therefore, s^m can be used to denote the highlight temporal locations and




Fig. 5. The demonstration of the SCL module. The module maximizes and minimizes the multiplied values of the paired and the unpaired representation vectors, respectively.

is regarded as the highlight score sequence in inference. We also guide s^m to be more distinguishable at the highlight moments. We use the top- k pooling to aggregate the k activations with the highest probabilities and enlarge their values. The sorted top- k activation index is obtained as

$$\begin{aligned} \Phi^k &= \{i \in \{0, \dots, k-1\} | l_i \in \Omega^k \text{ and} \\ &\quad s_{l_i}^m \leq s_{l_{i+1}}^m \text{ for } i < k-1\}, \\ \text{s.t. } \Omega^k &= \underset{\Gamma \subseteq \{1, \dots, N\}, |\Gamma|=k}{\operatorname{argmax}} \sum_{i \in \Gamma} s_i^m. \end{aligned} \quad (6)$$

However, the hyperparameter k is fixed, and we cannot guarantee that all the selected k activations correspond to highlight moments. To solve this noisy selection problem, we propose k -point contrastive learning to suppress the outliers of the selected activations. This mechanism assigns attention weights to the top activations according to the similarity of the top and the bottom processed representation vectors. The sorted bottom- k index set can be obtained as

$$\begin{aligned} \Pi^k &= \{i \in \{0, \dots, k-1\} | l_i \in \Psi^k \text{ and} \\ &\quad s_{l_i}^m \leq s_{l_{i+1}}^m \text{ for } i < k-1\}, \\ \text{s.t. } \Psi^k &= \underset{\Gamma \subseteq \{1, \dots, N\}, |\Gamma|=k}{\operatorname{argmin}} \sum_{j \in \Gamma} s_j^m. \end{aligned} \quad (7)$$

The objective function of the k -point contrastive learning can be expressed as follows:

$$\begin{aligned} L_r^m &= -\log \frac{1}{1 + e^{-\eta}}, \\ \eta &= \frac{1}{k} \sum_{i=0}^{k-1} \hat{s}_i^m \cdot [1 - \operatorname{sim}(\sigma(\hat{W}_i \hat{z}_i^m), \sigma(\check{W}_i \check{z}_i^m))], \end{aligned} \quad (8)$$

where $\hat{s}_i^m \in \{s_i^m | i \in \Phi^k\}$, $\hat{z}_i^m \in \{z_i^m | i \in \Phi^k\}$, $\check{z}_i^m \in \{z_i^m | i \in \Pi^k\}$, and sim represent the cosine similarity functions. \hat{W}_i and \check{W}_i are two projection matrices for \hat{z}_i^m and \check{z}_i^m , respectively. σ denotes the layer normalization function. Through Eq. (8), the top activations \hat{s}_i^m are enlarged. If

\hat{z}_i^m is similar to \tilde{z}_i^m in the projected space, \hat{s}_i^m will be assigned a low weight. Thus, the impact of noisy activations is suppressed. When L_r is minimized, the distance between \hat{z}_i^m and \tilde{z}_i^m increases. Through k -point contrastive learning, the top- k activation points for the highlight moments become more distinguishable, while the incorporated noisy outliers are suppressed.

D. Symmetric Contrastive Learning Module

The visual and audio modalities are highly correlated in highlight detection. For example, when a basketball player has a brilliant goal on the court, there are cheers from the crowd. To build the connection between the two modalities, we use the SCL module, as shown in Fig. 5, to learn the paired representations r^v and r^a . r^v and r^a are the representation vector sequences generated from the encoders of the visual and the audio modalities, respectively. The learned representations with visual-audio-level semantics will continue to be processed separately by the two branch networks. The SCL module is inspired by [30], which uses symmetric contrastive learning to predict the labels of image and text pairs for unsupervised learning. We utilize the SCL module to enable representations to have paired cross-modal semantics. This is achieved by minimizing the term

$$L_s = -\Gamma \left(\sum_i \log((r_i^a)^T r_i^v) - \sum_i \sum_{j \neq i} \log((r_i^a)^T r_j^v) \right), \quad (9)$$

where Γ is a learnable temperature parameter. It serves as a learnable weight that is independent of the network structure and optimized during pretraining. It controls the strength of symmetric contrastive learning, which constructs the connection between the visual and audio modalities. In Eq. (9), we use multiplication to calculate the similarity values of two modal representation vectors. The paired visual and audio representation vectors are associated by maximizing their similarity values, while the similarity values of unpaired representations are suppressed by minimization. Through the SCL module, the representation vectors of the audio and the visual modalities interact. These vectors are then processed separately through their respective branches without interlacing their structures. During inference, a single modality can be input to its corresponding branch using cross-modal pretrained knowledge to generate the highlight score sequence. Here, we adopt the cross-modal pretrained visual branch and use its RASL module to generate the highlight score sequence s^v as our final highlight score sequence $s = [s_0, s_1, \dots, s_N]$.

E. The Auxiliary Task of Masked FVS Reconstruction

As our RASL module highly relies on the representations, improving the efficiency of the representations is essential for unsupervised highlight detection. Unsupervised masked signal reconstruction has shown its ability to improve the latent representation performance for downstream applications [25], [26], [58]. Specifically, [26] demonstrated that the autoencoder is suitable for the masked signal reconstruction task. Inspired by

Algorithm 1 Pretraining of the proposed method

- 1: **Preprocessing:** Initialize the network parameters θ
 - 2: **while** θ has not converged **do**
 - 3: calculating ρ^v and ρ^a from the sampled video visual frames and the audio waves. Randomly generating the mask M and obtain $\hat{\rho}^m$;
 - 4: **for** $m = v, a$ **do**
 - 5: Compute r^m based on Eqs. (1), (2), (3);
 - 6: Compute s^m based on Eq. (4);
 - 7: Compute L_e^m based on Eq. (5);
 - 8: Obtain Φ^k and Π^k based on Eqs. (6), (7);
 - 9: Compute L_r^m based on Eq. (8);
 - 10: Compute L_{au}^m based on Eq. (10);
 - 11: **end for**
 - 12: Compute L_s based on Eq. (9);
 - 13: Compute L_{total} based on Eq. (11);
 - 14: Update θ via gradient back-propagation from L_{total} .
 - 15: **end while**
-

this, we incorporate an auxiliary masked FVS reconstruction task within the framework. This task leverages the existing autoencoder structure and does not introduce additional complexity to the network. By reconstructing the masked FVS, the model learns to make predictions based on existing video information and understands the video semantics, enabling it to learn meaningful representations and effectively handle scenarios where certain video information is missing. This makes the model more robust and adaptable to real-world scenarios. The training of the auxiliary task relies only on the following reconstruction loss:

$$L_{au}^m = | \rho^m - D^m(s^m \cdot E^m(F^m(M \cdot \rho^m))) | \quad (10)$$

$$= | \rho^m - D^m(s^m \cdot E^m(F^m(\hat{\rho}^m))) |,$$

where $\hat{\rho}^m$ represents the masked FVS of modality m , and M masks a portion of the input FVS ρ^m . The auxiliary task intends to reconstruct ρ^m from $\hat{\rho}^m$. Since the SA module allows the model to selectively attend to different parts of the input sequence, enhancing its reconstruction capabilities and capturing meaningful dependencies within the data, it serves as a shared component that applies to both the main and auxiliary tasks. The difference between the reconstruction loss for the main and auxiliary tasks lies solely in the input. The main task utilizes the original FVS, whereas the auxiliary task uses the masked FVS to reconstruct the original FVS. The network leverages the knowledge from both tasks and improves the highlight detection performance.

F. Implementation

Based on the descriptions of the RASL module, the SCL module and the auxiliary task, the total loss of the proposed framework is shown as follows:

$$L_{total} = \sum_{m=v,a} (L_e^m + L_r^m + L_{au}^m) + L_s \quad (11)$$

The main and auxiliary tasks are trained simultaneously. The gradients of the network are backpropagated after the total loss

L_{total} is obtained. The pretraining procedure of the proposed method is organized in Algorithm 1. We use the RMSprop optimizer [59] with a learning rate of 0.001. The mask M in Eq. (10) masks 50% feature vectors of the input FVS ρ^m . The feature extractors [14] and [57] are used to build the visual and the audio FVSs from visual frames and audio waves, respectively. The visual and audio feature extractors contain transformer structures and use attention mechanisms to learn semantics, enabling the generation of high-quality representations. The visual and audio branches of the network have the same architectures. The encoder E and the decoder D contain 3 convolutional and 3 deconvolutional layers, respectively, with a stride of 2. There are ReLU activation functions that follow each convolutional and deconvolutional layer. The hyperparameter k in Eq. (6), (7) and (8) is set to 10. The initial value of Γ is set to 3.1. We implement the FVS generation and model pretraining steps sequentially using a server equipped with a 10-core CPU and a Tesla T4 GPU, which has 16 GB of memory capacity. The batch size for model pretraining was 8. We clip the input video length with a fixed length of 30 seconds. When the input or the remaining clipped video length is less than 30 seconds, we concatenate the video repeatedly until it is longer than 30 seconds and clip it. For the visual modality, we sampled the video every 0.2 seconds; thus, the temporal length of the visual FVS was 150. For the audio modality, the temporal density of the feature vectors extracted from audio waves is much greater than that of the visual feature vectors. To reduce the network parameters, we use average pooling to downsample the temporal length to 150. Linear interpolation is used to resample the vector length ρ_i^a to be equal to ρ_i^v . During inference, if the video is longer than 30 seconds, we first split it into clips of 30 seconds and sequentially input the clips into the networks. Then, we concatenate and normalize the highlight score sequences obtained from the network as the final result. If the input or the clipped video length is shorter than 30 seconds, we repeatedly concatenate and clip the video in the same way as in the training data preparation and then use the portion of the highlight score sequence corresponding to the original video as the result. The network is pretrained with the videos of the large-scale dataset ActivityNet [60]. The pretrained model can be directly used for inferring highlights of wild videos. Users can determine the highlight score threshold to segment

the videos according to the desired highlight length.

IV. EXPERIMENTS

In this section, we conduct extensive experiments and comparisons to evaluate the performance of the proposed method. Following other methods [4], [9], the experiments were conducted on the popular datasets YouTube Highlights [3] and TVSum [21]. There are various methods for comparison, including supervised, weakly supervised and unsupervised methods.

A. Evaluation datasets

We evaluate our method on YouTube Highlights and TVSum with the model pretrained on ActivityNet. The YouTube Highlights dataset contains 6 specific categories: surfing, skating, skiing, gymnastics, parkour, and dog. Each category consists of 100 segment-level annotated videos, and the accumulated time is approximately 1,430 minutes. TVSum contains 10 specific categories, including changing vehicle tires, grooming an animal, making sandwiches, parades, flash mob gatherings, and others, with 5 frame-level annotated videos in each category. Following [4], [47], [48], we average the importance scores of the manual labels of every segment to achieve segment-level highlight scores. To evaluate the proposed method, we report the mean average precision (mAP) and top 5 mAPs for YouTube Highlights and TVSum, respectively, as in [4], [47], [48].

B. Comparisons

We compare our proposed framework to other competing approaches for evaluation. The compared supervised methods include LSVM [3], GIFs [1], KVS [44], sLstm [46], DPP [45], SMRS [42] and SM [61]. The compared weakly supervised approaches include RARE [2], MBF [40], CVS [48], DSN [47], VESD [41], SG [37], Quasi [43], LM [4], MINI-Net [10], and LR [9], and the compared unsupervised method is CHD [62]. Among them, LR and MINI-Net also utilize visual and audio modalities. CHD is another unsupervised method that does not require video labels. In addition to highlight detection methods, several of these methods involve video summarization approaches. Since most of the important frames detected belong to highlights, we still compare these

TABLE I
EXPERIMENTAL RESULTS ON THE YOUTUBE HIGHLIGHTS (YTH) DATASET (mAP)

Topic	Supervised (frame-level annotations)		Weakly supervised (video-level annotations)				Unsupervised (trained on YTH) (w/o using YTH data)	
	GIFs	LSVM	RARE	LM	MINI-Net	LR	CHD	Ours
dog	0.308	0.60	0.49	0.579	0.577	0.554	0.606	<u>0.642</u>
gymnast.	0.335	0.41	0.35	0.417	0.574	0.623	0.711	0.758
parkour	0.540	0.61	0.50	0.670	0.698	0.701	0.742	0.709
skating	0.554	0.62	0.25	0.578	0.522	0.691	0.498	0.456
skiing	0.328	0.36	0.22	0.486	0.539	0.601	0.682	0.665
surfing	0.541	0.51	0.49	0.651	0.593	0.598	0.685	0.667
Average	0.464	0.54	0.38	0.564	0.584	0.630	0.654	<u>0.651</u>

The best and the second best overall performances are marked in bold and underlined respectively.

TABLE II
EXPERIMENTAL RESULTS ON THE TVSUM DATASET (TOP 5 MAP)

Topic	Supervised (frame-level annotations)					Weakly supervised (video-level annotations)								Unsupervised (trained on TVSum) (w/o using TVSum data)		
	KVS	DPP	sLSTM	SM	SMRS	Quasi	MBF	CVS	SG	VESD	DSN	LM	MINI-Net	LR	CHD	Ours
VT	0.353	0.399	0.411	0.415	0.272	0.336	0.295	0.328	0.423	0.447	0.373	0.559	0.785	0.850	-	0.897
VU	0.441	0.453	0.462	0.467	0.324	0.369	0.357	0.413	0.472	0.493	0.441	0.429	0.566	0.714	-	0.657
GA	0.402	0.457	0.463	0.469	0.331	0.342	0.325	0.379	0.475	0.496	0.428	0.612	0.736	0.819	-	0.925
MS	0.417	0.462	0.477	0.478	0.362	0.375	0.412	0.398	0.489	0.503	0.436	0.540	0.753	0.786	-	0.583
PK	0.382	0.437	0.448	0.445	0.289	0.324	0.318	0.354	0.456	0.478	0.411	0.604	0.769	0.802	-	0.847
PR	0.403	0.446	0.461	0.458	0.276	0.301	0.334	0.381	0.473	0.485	0.417	0.475	0.633	0.755	-	0.830
FM	0.397	0.442	0.452	0.451	0.302	0.318	0.365	0.365	0.464	0.487	0.412	0.432	0.612	0.716	-	0.697
Bk	0.342	0.395	0.406	0.407	0.297	0.295	0.313	0.326	0.417	0.441	0.368	0.663	0.756	0.773	-	0.833
BT	0.419	0.464	0.471	0.473	0.314	0.327	0.365	0.402	0.483	0.492	0.435	0.691	0.756	0.786	-	0.875
DS	0.394	0.449	0.455	0.453	0.295	0.309	0.357	0.378	0.466	0.488	0.416	0.626	0.656	0.681	-	0.890
Average	0.398	0.447	0.451	0.461	0.306	0.329	0.345	0.372	0.462	0.481	0.424	0.563	0.702	<u>0.768</u>	0.528	0.783

The best and the second best overall performances are marked in bold and underlined, respectively.




Fig. 6. The performance (mAP) varies by changing the k value of the k -point contrastive learning.

summarization approaches using the same metrics. The results of the compared methods evaluated on the two datasets are reported in [4], [9], [62].

Table I shows the comparison results on YouTube Highlights. Compared to the weakly supervised approaches MINI-Net and LR, which also use visual and audio modalities, our unsupervised method achieves average gains of 6.7% and 2.1%, respectively. Notably, our method infers the highlight segments without requiring the audio modality as the input, while MINI-Net and LR require both visual frames and audio waves as the input. Our method also achieves better average performance than do the supervised methods LSVM and GIFs. This is because supervised methods easily fall into the problem

of overfitting frame-level data. Our method and CHD demonstrated very similar overall performances. However, CHD is trained on the YouTube Highlights training split. It leverages the inclusion of similar videos in both the training and test splits, allowing it to easily align with the test video distribution. In contrast, our method trains on an unrelated dataset. This allows our method to be directly applicable and practical for real-world applications, especially in scenarios where data collection is laborious or data availability is constrained. We also find that the proposed method shows relatively weak performance in terms of the skating video category. This is because the cameras of most skating videos of YouTube Highlights shake severely and do not face skaters stably, which is different from the usual visual presentation forms in the unlabeled pretraining dataset ActivityNet. This deteriorates the performance of our method for skating videos. Table II shows the experimental results on TVSum. TVSum consists of 10 categories of diverse long videos. As another unsupervised method, CHD, only reports its overall performance on TVSum but not across TVSum categories [62]; we only compare its overall performance. The proposed method achieves the best overall performance compared to all the competing methods. We also observe that the proposed method outperforms CHD with a large gain of 25.5%, even though CHD is trained on the TVSum training split. This shows the robust and adaptive performance of the proposed method. We find that the proposed method shows relatively weak performance in

TABLE III
THE ABLATION STUDY RESULTS ON THE DATASETS

Method	YouTube Highlights	TVSum
Ours w/o RASL	0.560	0.616
Ours w/o SA	0.638	0.711
Ours w/o auxiliary	0.644	0.724
Ours w/o vision	0.625	0.614
Ours w/o audio	0.630	0.704
Ours	0.651	0.783

TABLE IV
THE COMPARISONS OF DIFFERENT AUDIO-VISUAL FUSION SCHEMES

Method	YouTube Highlights	TVSum
Summation	0.630	0.700
Concatenation	0.628	0.716
Submodule MLP [10]	0.589	0.604
low-rank fusion [9]	0.621	0.736
Ours	0.651	0.783

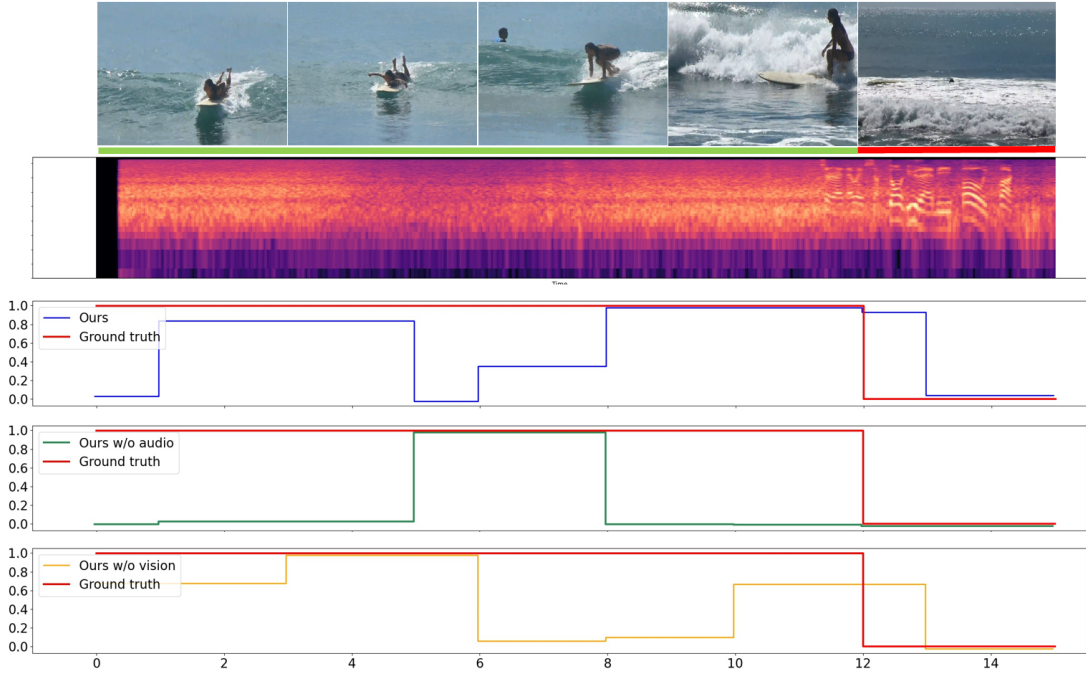


Fig. 7. Illustration of the SCL module that improves the highlight detection performance. The subfigures in the first and the second rows are the representative visual frames and the audio mel-spectrogram of the input video. The subfigures in the third, fourth and last rows are the highlight score curves estimated by the fully proposed method, the proposed method pretrained without the audio modality and the method trained without the visual modality, respectively. The proposed method pretrained without the audio modality leaves out the beginning and the tail parts, while the proposed method pretrained without the visual modality leaves out the middle part of the highlights. In contrast, the full proposed method with the SCL module connecting the two modalities yields the best overall results.

the making sandwich (MS) category. This is because videos in this category include segments where chefs use exaggerated statements and actions to introduce the sandwich and capture the attention of viewers. However, these segments are not considered key segments based on prior knowledge of the topic of making sandwiches. In contrast, the weakly supervised trained methods MINI-Net and LR learn highlight segments given the video topics, resulting in better performance for videos in the sandwich category. Overall, the unsupervised proposed method achieves 8.1% and 1.5% gains in terms of the average score by the most competitive methods, MINI-Net and LR, respectively. This demonstrates the superior performance of the proposed method.

C. Ablation study

In this subsection, we conduct ablation studies to investigate the effects of the proposed components in the model. We also evaluate the impacts of the modality fusion schemes and the values of k in the k -point contrastive learning and the reconstruction targets.

Impacts of the model components: We evaluate the model components and show the ablation study results in Table III. We first evaluate the effect of the proposed RASL module, which uses k -point contrastive learning. In this case, we set the value of k to 0, and the loss term L_r^m is excluded. Without the k -point contrastive learning, the RASL module degrades to a reweighting module. The performance of the proposed method drops by 9.1% and 16.7% on YouTube Highlights and TVSum, respectively. This finding certifies

TABLE V
THE COMPARISONS OF RECONSTRUCTION TARGETS

Visual target	Audio target	YouTube Highlights	TVSum
Pixel	mel-spectrogram	0.639	0.603
Pixel	Audio FVS [57]	0.648	0.608
Visual FVS [14]	mel-spectrogram	0.649	0.700
Visual FVS [14]	Audio FVS [57]	0.651	0.783

that the proposed RASL module using k -point contrastive learning can emphasize important representation activations and improve highlight detection performance. We also drop the SA module to evaluate the effect of SA enhancement on FVS. Without the SA module, the performance drops by 2.3% and 7.2% on YouTube Highlights and TVSum, respectively. This finding certifies that the SA module can improve the representation performance of FVS. We also ablate the auxiliary masked FVS reconstruction task to evaluate the effect of the MTL on our model learning. In this case, the proposed method without MTL also falls behind the full version. The results indicate that the MTL framework improves highlight detection performance by learning knowledge across the main and auxiliary tasks. We also show the experimental results of pretraining the model using only a single modality. The performance degradation indicates that the association between the paired visual and audio perceptions in the pretrained model improves the performance.

Impact of the modality fusion scheme: We compare 4 other modality fusion schemes to our proposed SCL module.

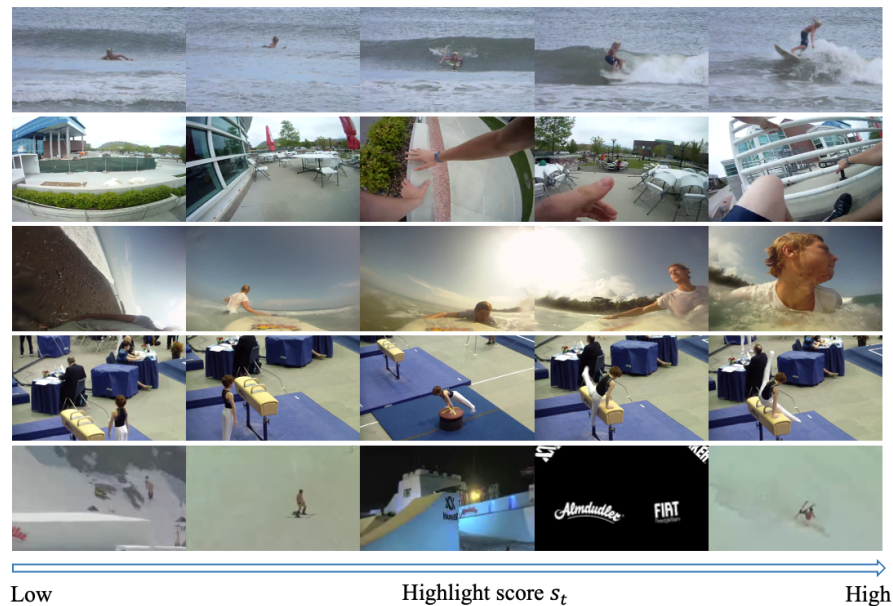


Fig. 8. The example highlighting the prediction ability of the proposed method on YouTube Highlights. The highlight score s_t increases along the axis from the left to the right.

Summation and concatenation are the two most commonly used approaches for modal fusion. Submodule MLP and low-rank fusion are two modal fusion schemes proposed in the highlight detection methods [10] and [9], respectively. As all the compared fusion schemes directly fuse the required visual and audio modalities, we fuse the self-attended FVS $\bar{\rho}^v$ and $\bar{\rho}^a$ using the compared fusion schemes and send the fused FVS to a network with the same architecture as one branch of the proposed network for self-reconstruction. Table IV shows the comparison results. Our method achieves the best performance with the proposed SCL module. This is because some sounds of the attractive scenes in YouTube Highlights and TVSum are disrupted by irrelevant environmental sounds or background music, which decreases the performance of the above direct fusion schemes. The submodule MLP even suffers from overfitting, as the complex network has a large number of parameters. Compared to other schemes that directly fuse the representations of the two modalities, the proposed model trained employing the SCL module requires only the visual modality as the input during inference. The visual branch of our model generates the highlight score sequence by inferring representation vectors with paired visual-audio semantics without the probable audio noise, as the overall correct connection of the related visual and audio semantics has been established by the SCL module during pretraining on ActivityNet, where the existing audio noise varies and does not consistently lead the model to converge toward a fixed erroneous modal connection.

Impact of the k -point contrast learning: In addition, we investigate the effect of the k -point contrast learning and the value k on our framework. We ablate the contrast learning by using $\eta = \frac{1}{k} \sum_{i=0}^k \hat{s}_i^m$ as a substitute in Eq. (8) for comparison. Fig. 6 shows the results of our framework with and without k -point contrastive learning with varying values

TABLE VI
THE COMPARISONS OF THE EFFICIENCY

Method	Training method	Parameters	FLOPs
GIFS	Supervised	440.31M	823.44G
sLSTM	Supervised	2.69M	0.07G
SG	Weakly supervised	18.78M	9.16G
Ours	Unsupervised	8.53 M	0.82G

of k . The performance of the proposed method with k -point contrastive learning is better overall than that without k -point contrast learning. Notably, the gaps between the curves are relatively larger when the value of k is high than when k is small. This is because our k -point contrastive learning suppresses noisy activations when using a large k .

Impact of the reconstruction targets: We also demonstrated that the use of FVS as the reconstruction target is more beneficial for highlight detection than the use of visual pixels and audio mel-spectrograms. Table V shows comparisons of the proposed model trained on different reconstruction targets for highlight detection. The results in Table V demonstrate that utilizing the visual and audio FVSs extracted by deep networks [14] and [57] as the reconstruction targets achieves the best performance for highlight detection. Replacing one or both of the deep feature extractors with visual pixels or mel-spectrograms results in a degradation in performance. This is because features extracted by deep neural networks have better representations than traditional pixels and mel-spectrograms. It is also possible that the performance of the proposed highlight detection method can be further improved with other advanced features due to the rapid development of deep feature extractors. Notably, our results using very simple feature pixels and a mel-spectrogram still outperform 6 of the 7 methods in Table I and 13 of the 15 methods in

Table II for YouTube Highlights and TVSum, respectively, even other methods use complex features. This demonstrates the robustness of the proposed framework to different input features.

D. Qualitative evaluation

This subsection presents the qualitative evaluation results for further study. In Fig. 7, we show representative frames, an audio mel-spectrogram of an input video and the estimated highlight score curves. The highlight score curves are estimated using three models: the proposed network, the proposed network pretrained without using the audio modality, and the proposed network pretrained without using the visual modality. The proposed method, pretrained without the audio modality, yields high estimated scores in the middle, where the surfer strongly moves from lying on the surfboard to standing on the surfboard. However, this approach ignores the beginning and the tail highlight parts, where the surfer is motionless, either lying or standing on the surfboard. The proposed method, pretrained without the visual modality, covers a longer range of correct high highlight scores. However, relatively low scores were observed in the middle, where degradation of the mel-spectrogram was observed. It is worth noting that both our method pretrained without the audio and the visual modalities give no credit at the end of the video. This is because, visually, there are only sea waves and no surfing motion at the tail. Acoustically, the pattern of the mel-spectrogram at the tail was different from that of the overall mel-spectrogram. With the SCL module, the proposed method considering cross-modal semantics yields the best overall results.

We also demonstrate the representative frames from five highlight detection results on YouTube Highlights in Fig. 8. In the first four instances, we find that the representative frames with higher values of s_t represent attractive segments that contain more information, while the frames with low s_t values are less meaningful. This indicates that our unsupervised method effectively identifies highlight moments. However, in the last instance, where abrupt video transitions occur, the network mistakenly identifies these transitions as highlight frames. Although these sudden transitions draw viewer attention, they are not relevant to the video topic and hence should not be considered highlights. This shows the limitation of our proposed method, as it might classify some irrelevant frames that capture attention as highlights.

E. Efficiency evaluation

We also conduct an efficiency evaluation of the proposed method and compare it to two supervised methods and one weakly supervised method that have available codes. As different official implementation platforms with various efficiencies are used, we count the number of parameters and floating point operations (FLOPs) for independently assessing the computational complexity of each algorithm. Considering that the inputs for all the methods are sustainable video features, we ensure a fair comparison by excluding the feature extractors and focusing solely on the network architectures for all the methods. The experimental results are shown in Table VI. The

supervised method GIFs has the highest number of parameters and FLOPs, while the sLSTM has the lowest. Notably, the proposed method also yields a low FLOP value, specifically lower than 1G, and has half the number of parameters compared to SG.

V. CONCLUSION

In this paper, we present a novel unsupervised cross-modal highlight detection framework. We propose the RASL module with the k -point contrastive learning mechanism to learn the significant activations through a self-reconstruction task. To enable the network to connect the visual and audio modalities, we propose the SCL module to learn paired representations. Given only the visual frames, the cross-modal pretrained network can generate representations with visual-audio-level semantics and directly infer the highlight scores. An auxiliary task of masked FVS reconstruction is used to enhance the representation. The experimental results demonstrate the effectiveness and superior performance of the proposed approach compared to other state-of-the-art highlight detection approaches.

REFERENCES

- [1] M. Gygli, Y. Song, and L. Cao, "Video2gif: Automatic generation of animated gifs from video," in *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2016, pp. 1001–1009.
- [2] H. Yang, B. Wang, S. Lin, D. Wipf, M. Guo, and B. Guo, "Unsupervised extraction of video highlights via robust recurrent auto-encoders," in *Proceedings of the IEEE international conference on computer vision*, 2015, pp. 4633–4641.
- [3] M. Sun, A. Farhadi, and S. Seitz, "Ranking domain-specific highlights by analyzing edited videos," in *European conference on computer vision*. Springer, 2014, pp. 787–802.
- [4] B. Xiong, Y. Kalantidis, D. Ghahyaram, and K. Grauman, "Less is more: Learning highlight detection from video duration," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2019, pp. 1258–1267.
- [5] T. Yao, T. Mei, and Y. Rui, "Highlight detection with pairwise deep ranking for first-person video summarization," in *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2016, pp. 982–990.
- [6] W. Liu, T. Mei, Y. Zhang, C. Che, and J. Luo, "Multi-task deep visual-semantic embedding for video thumbnail selection," in *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2015, pp. 3707–3715.
- [7] Y. Jiao, X. Yang, T. Zhang, S. Huang, and C. Xu, "Video highlight detection via deep ranking modeling," in *Pacific-Rim Symposium on Image and Video Technology*. Springer, 2017, pp. 28–39.
- [8] M. Rochan, M. K. Krishna Reddy, L. Ye, and Y. Wang, "Adaptive video highlight detection by learning from user history," in *European conference on computer vision*. Springer, 2020, pp. 261–278.
- [9] Q. Ye, X. Shen, Y. Gao, Z. Wang, Q. Bi, P. Li, and G. Yang, "Temporal cue guided video highlight detection with low-rank audio-visual fusion," in *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 2021, pp. 7950–7959.
- [10] F.-T. Hong, X. Huang, W.-H. Li, and W.-S. Zheng, "Mini-net: Multiple instance ranking network for video highlight detection," in *European Conference on Computer Vision*. Springer, 2020, pp. 345–360.
- [11] A. Sharghi, B. Gong, and M. Shah, "Query-focused extractive video summarization," in *Computer Vision—ECCV 2016: 14th European Conference, Amsterdam, The Netherlands, October 11–14, 2016, Proceedings, Part VIII 14*. Springer, 2016, pp. 3–19.
- [12] R. Zeng, W. Huang, M. Tan, Y. Rong, P. Zhao, J. Huang, and C. Gan, "Graph convolutional networks for temporal action localization," in *Proceedings of the IEEE/CVF international conference on computer vision*, 2019, pp. 7094–7103.

- [13] C. Gan, N. Wang, Y. Yang, D.-Y. Yeung, and A. G. Hauptmann, "Devnet: A deep event network for multimedia event detection and evidence recounting," in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2015, pp. 2568–2577.
- [14] Z. Liu, Y. Lin, Y. Cao, H. Hu, Y. Wei, Z. Zhang, S. Lin, and B. Guo, "Swin transformer: Hierarchical vision transformer using shifted windows," in *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 2021, pp. 10012–10022.
- [15] V. Sanguineti, P. Morerio, A. Del Bue, and V. Murino, "Audio-visual localization by synthetic acoustic image generation," in *Proceedings of the AAAI Conference on Artificial Intelligence*, vol. 35, no. 3, 2021, pp. 2523–2531.
- [16] H. Akbari, L. Yuan, R. Qian, W.-H. Chuang, S.-F. Chang, Y. Cui, and B. Gong, "Vatt: Transformers for multimodal self-supervised learning from raw video, audio and text," *Advances in Neural Information Processing Systems*, vol. 34, 2021.
- [17] T. Li, Z. Sun, H. Zhang, J. Li, Z. Wu, H. Zhan, Y. Yu, and H. Shi, "Deep music retrieval for fine-grained videos by exploiting cross-modal-encoded voice-overs," in *Proceedings of the 44th International ACM SIGIR Conference on Research and Development in Information Retrieval*, 2021, pp. 1880–1884.
- [18] T. Afouras, J. S. Chung, A. Senior, O. Vinyals, and A. Zisserman, "Deep audio-visual speech recognition," *IEEE transactions on pattern analysis and machine intelligence*, 2018.
- [19] M.-M. Cheng, N. J. Mitra, X. Huang, P. H. Torr, and S.-M. Hu, "Global contrast based salient region detection," *IEEE transactions on pattern analysis and machine intelligence*, vol. 37, no. 3, pp. 569–582, 2014.
- [20] Z. Wang, D. Xiang, S. Hou, and F. Wu, "Background-driven salient object detection," *IEEE transactions on multimedia*, vol. 19, no. 4, pp. 750–762, 2016.
- [21] Y. Song, J. Vallmitjana, A. Stent, and A. Jaimes, "Tvsum: Summarizing web videos using titles," in *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2015, pp. 5179–5187.
- [22] K. H. Jin, M. T. McCann, E. Froustey, and M. Unser, "Deep convolutional neural network for inverse problems in imaging," *IEEE Transactions on Image Processing*, vol. 26, no. 9, pp. 4509–4522, 2017.
- [23] K. Zhang, W. Zuo, Y. Chen, D. Meng, and L. Zhang, "Beyond a gaussian denoiser: Residual learning of deep cnn for image denoising," *IEEE transactions on image processing*, vol. 26, no. 7, pp. 3142–3155, 2017.
- [24] T. Li, Y.-H. Chan, and D. P. K. Lun, "Improved multiple-image-based reflection removal algorithm using deep neural networks," *IEEE Transactions on Image Processing*, vol. 30, pp. 68–79, 2020.
- [25] A. Dosovitskiy, L. Beyer, A. Kolesnikov, D. Weissenborn, X. Zhai, T. Unterthiner, M. Dehghani, M. Minderer, G. Heigold, S. Gelly *et al.*, "An image is worth 16x16 words: Transformers for image recognition at scale," in *International Conference on Learning Representations (ICLR)*, 2021.
- [26] K. He, X. Chen, S. Xie, Y. Li, P. Dollár, and R. Girshick, "Masked autoencoders are scalable vision learners," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2022, pp. 16 000–16 009.
- [27] A. Baevski, Y. Zhou, A. Mohamed, and M. Auli, "wav2vec 2.0: A framework for self-supervised learning of speech representations," *Advances in Neural Information Processing Systems*, vol. 33, pp. 12 449–12 460, 2020.
- [28] C. Zhuang, T. She, A. Andonian, M. S. Mark, and D. Yamins, "Unsupervised learning from video with deep neural embeddings," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2020, pp. 9563–9572.
- [29] H. Chen, Y. Wang, T. Guo, C. Xu, Y. Deng, Z. Liu, S. Ma, C. Xu, C. Xu, and W. Gao, "Pre-trained image processing transformer," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2021, pp. 12 299–12 310.
- [30] A. Radford, J. W. Kim, C. Hallacy, A. Ramesh, G. Goh, S. Agarwal, G. Sastry, A. Askell, P. Mishkin, J. Clark *et al.*, "Learning transferable visual models from natural language supervision," in *International Conference on Machine Learning*. PMLR, 2021, pp. 8748–8763.
- [31] J. Devlin, M.-W. Chang, K. Lee, and K. Toutanova, "Bert: Pre-training of deep bidirectional transformers for language understanding," in *the Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, 2019.
- [32] H. Tang, V. Kwatra, M. E. Sargin, and U. Gargi, "Detecting highlights in sports videos: Cricket as a test case," in *2011 IEEE International Conference on Multimedia and Expo*. IEEE, 2011, pp. 1–6.
- [33] J. Wang, C. Xu, E. Chng, and Q. Tian, "Sports highlight detection from keyword sequences using hmm," in *2004 IEEE International Conference on Multimedia and Expo (ICME)(IEEE Cat. No. 04TH8763)*, vol. 1. IEEE, 2004, pp. 599–602.
- [34] Z. Xiong, R. Radhakrishnan, A. Divakaran, and T. S. Huang, "Highlights extraction from sports video based on an audio-visual marker detection framework," in *2005 IEEE International Conference on Multimedia and Expo*. IEEE, 2005, pp. 4–pp.
- [35] Y. Yu, S. Lee, J. Na, J. Kang, and G. Kim, "A deep ranking model for spatio-temporal highlight detection from a 360o video," in *Proceedings of the AAAI Conference on Artificial Intelligence*, vol. 32, no. 1, 2018.
- [36] L. Wang, D. Liu, R. Puri, and D. N. Metaxas, "Learning trailer moments in full-length movies with co-contrastive attention," in *European Conference on Computer Vision*. Springer, 2020, pp. 300–316.
- [37] B. Mahasseni, M. Lam, and S. Todorovic, "Unsupervised video summarization with adversarial lstm networks," in *Proceedings of the IEEE conference on Computer Vision and Pattern Recognition*, 2017, pp. 202–211.
- [38] H. Kanafani, J. A. Ghauri, S. Hakimov, and R. Ewerth, "Unsupervised video summarization via multi-source features," in *Proceedings of the 2021 International Conference on Multimedia Retrieval*, 2021, pp. 466–470.
- [39] L. Yuan, F. E. Tay, P. Li, L. Zhou, and J. Feng, "Cycle-sum: Cycle-consistent adversarial lstm networks for unsupervised video summarization," in *Proceedings of the AAAI Conference on Artificial Intelligence*, vol. 33, no. 01, 2019, pp. 9143–9150.
- [40] W.-S. Chu, Y. Song, and A. Jaimes, "Video co-summarization: Video summarization by visual co-occurrence," in *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2015, pp. 3584–3592.
- [41] S. Cai, W. Zuo, L. S. Davis, and L. Zhang, "Weakly-supervised video summarization using variational encoder-decoder and web prior," in *Proceedings of the European conference on computer vision (ECCV)*, 2018, pp. 184–200.
- [42] E. Elhamifar, G. Sapiro, and R. Vidal, "See all by looking at a few: Sparse modeling for finding representative objects," in *2012 IEEE conference on computer vision and pattern recognition*. IEEE, 2012, pp. 1600–1607.
- [43] G. Kim, L. Sigal, and E. P. Xing, "Joint summarization of large-scale collections of web images and videos for storyline reconstruction," in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2014, pp. 4225–4232.
- [44] D. Potapov, M. Douze, Z. Harchaoui, and C. Schmid, "Category-specific video summarization," in *European conference on computer vision*. Springer, 2014, pp. 540–555.
- [45] B. Gong, W.-L. Chao, K. Grauman, and F. Sha, "Diverse sequential subset selection for supervised video summarization," *Advances in neural information processing systems*, vol. 27, 2014.
- [46] K. Zhang, W.-L. Chao, F. Sha, and K. Grauman, "Video summarization with long short-term memory," in *European conference on computer vision*. Springer, 2016, pp. 766–782.
- [47] R. Panda, A. Das, Z. Wu, J. Ernst, and A. K. Roy-Chowdhury, "Weakly supervised summarization of web videos," in *Proceedings of the IEEE International Conference on Computer Vision*, 2017, pp. 3657–3666.
- [48] R. Panda and A. K. Roy-Chowdhury, "Collaborative summarization of topic-related videos," in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2017, pp. 7083–7092.
- [49] A. Graves, "Long short-term memory," *Supervised sequence labelling with recurrent neural networks*, pp. 37–45, 2012.
- [50] A. Radford, L. Metz, and S. Chintala, "Unsupervised representation learning with deep convolutional generative adversarial networks," 2016.
- [51] M. Arjovsky, S. Chintala, and L. Bottou, "Wasserstein generative adversarial networks," in *International conference on machine learning*. PMLR, 2017, pp. 214–223.
- [52] T. Li and D. P. K. Lun, "Image reflection removal using the wasserstein generative adversarial network," in *ICASSP 2019-2019 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2019, pp. 1–5.
- [53] J. Baxter, "A bayesian/information theoretic model of learning to learn via multiple task sampling," *Machine learning*, vol. 28, no. 1, pp. 7–39, 1997.
- [54] L. Duong, T. Cohn, S. Bird, and P. Cook, "Low resource dependency parsing: Cross-lingual parameter sharing in a neural network parser," in *Proceedings of the 53rd annual meeting of the Association for Computational Linguistics and the 7th international joint conference on natural language processing (volume 2: short papers)*, 2015, pp. 845–850.
- [55] Y. Yang and T. M. Hospedales, "Trace norm regularised deep multi-task learning," *arXiv preprint arXiv:1606.04038*, 2016.

- [56] K. Weiss, T. M. Khoshgoftaar, and D. Wang, "A survey of transfer learning," *Journal of Big data*, vol. 3, no. 1, pp. 1–40, 2016.
- [57] A. T. Liu, S.-w. Yang, P.-H. Chi, P.-c. Hsu, and H.-y. Lee, "Mockingjay: Unsupervised speech representation learning with deep bidirectional transformer encoders," in *ICASSP 2020-2020 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2020, pp. 6419–6423.
- [58] A. Vaswani, N. Shazeer, N. Parmar, J. Uszkoreit, L. Jones, A. N. Gomez, Ł. Kaiser, and I. Polosukhin, "Attention is all you need," *Advances in neural information processing systems*, vol. 30, 2017.
- [59] A. Graves, "Generating sequences with recurrent neural networks," *arXiv preprint arXiv:1308.0850*, 2013.
- [60] F. Caba Heilbron, V. Escorcia, B. Ghanem, and J. Carlos Niebles, "Activitynet: A large-scale video benchmark for human activity understanding," in *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2015, pp. 961–970.
- [61] M. Gygli, H. Grabner, and L. Van Gool, "Video summarization by learning submodular mixtures of objectives," in *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2015, pp. 3090–3098.
- [62] T. Badamdorj, M. Rochan, Y. Wang, and L. Cheng, "Contrastive learning for unsupervised video highlight detection," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2022, pp. 14 042–14 052.