

PR-Net: Preference Reasoning for Personalized Video Highlight Detection

Runnan Chen^{*1,2}, Penghao Zhou², Wenzhe Wang³, Nenglun Chen¹, Pai Peng²,
 Xing Sun^{†2}, Wenping Wang^{†1}

¹The University of Hong Kong ²Youtu Lab, Tencent
³Zhejiang University

Abstract

Personalized video highlight detection aims to shorten a long video to interesting moments according to a user's preference, which has recently raised the community's attention. Current methods regard the user's history as holistic information to predict the user's preference but negating the inherent diversity of the user's interests, resulting in vague preference representation. In this paper, we propose a simple yet efficient preference reasoning framework (PR-Net) to explicitly take the diverse interests into account for frame-level highlight prediction. Specifically, distinct user-specific preferences for each input query frame are produced, presented as the similarity weighted sum of history highlights to the corresponding query frame. Next, distinct comprehensive preferences are formed by the user-specific preferences and a learnable generic preference for more overall highlight measurement. Lastly, the degree of highlight and non-highlight for each query frame is calculated as semantic similarity to its comprehensive and non-highlight preferences, respectively. Besides, to alleviate the ambiguity due to the incomplete annotation, a new bi-directional contrastive loss is proposed to ensure a compact and differentiable metric space. In this way, our method significantly outperforms state-of-the-art methods with a relative improvement of 12% in mean accuracy precision.

1. Introduction

The short video has become an indispensable medium for people to acquire knowledge and share experiences in their daily lives. However, an unedited video cost minutes or hours for a person to collect meaningful moments by going through the whole video. How to automatically identify the most interesting moments for specific users has recently

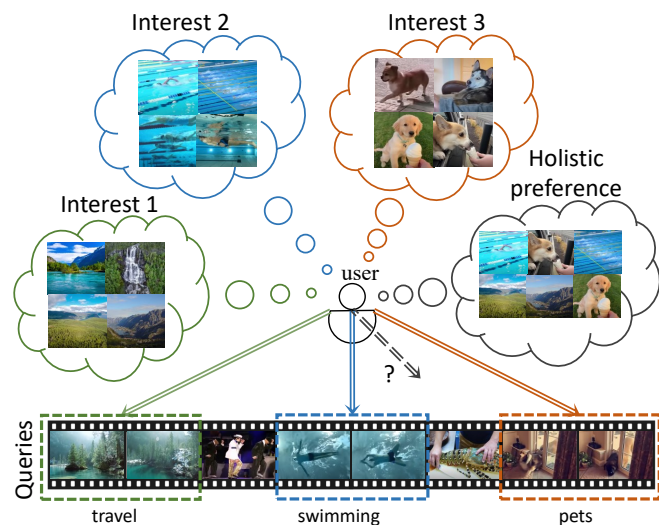


Figure 1. Considering the diversity of the user's interests is critical for guiding personalized video highlight detection. For example, if a user is interested in travel, swimming and pets, those query frames that are semantically similar to one of the user's interests can be marked as the highlight frames. On the other hand, predicting a holistic user's preference from history will lead to vague preference representation and unsatisfied performance.

drawn attention from the research community.

Unlike generic highlight detection methods [41, 13, 40], we focus on personalized video highlight detection, where the user's previously created historical highlight video segments are used as guidance to detect the interesting moments. To the best of our knowledge, only two methods are proposed to solve this problem. Molino and Gygli [7] directly concatenate the history feature into the input video feature to perform segment level prediction by using a ranking model. Rochan et al. [28] introduce a temporal-adaptive instance normalization layer that encodes history for frame-level highlight predictions. Although promising results they achieved, there are still some limitations. Firstly, They re-

*Work done during internship at Tencent Youtu Lab.

†Corresponding author

gard the user’s history as holistic information to predict the user’s preference based on the assumption that there is only one preference in its history. However, it is somehow against common sense in the real world because the user’s interests are inherently diverse. Suppose a user prefers to travel, swimming, pets, et al. It is not suitable to extract a holistic preference to represent such diverse interests (Fig. 1). Secondly, because there are many repeated or similar shots in a video while only a few positive samples are marked, some unlabeled samples are actually positive. However, the current methods treat all unlabeled samples as negative ones, leading to the wrong label assignment [5].

In this paper, we propose a preference reasoning framework (PR-Net) for frame-level personalized video highlight detection, which overcomes the above limitations. Intuitively, if a video frame is semantically similar to one of the user’s history highlights, it could be regarded as an interesting moment. Based on this observation, we first attend the query frame to history frames to obtain attention weights and then forming the user-specific preference embedding by the attention-weighted sum of history frame embeddings. Considering the generic preference is also an important reference to judge the degree of highlight, especially when the user’s history is missing. We set a learnable generic preference embedding, and it combines the user-specific preference embedding to generate the comprehensive preference embedding for more overall highlight detection. In the end, to ensure a more compact and differentiable metric space, a non-highlight preference is used to consider the degree of non-highlight for each query frame. We first infer the degree of highlight and non-highlight of each query frame based on the semantic similarity to its comprehensive preference embedding and non-highlight preference embedding, respectively. And then constraint the two similarities of all video frames by a new contrastive loss function. Note that we only train those most non-highlight frames as negative samples to alleviate the wrong label assignment problem.

We conduct the experiments on the PHD-GIF [7] dataset, the only related large-scale dataset for this task. The results show that our method significantly outperforms the state-of-the-art methods, with the relative improvement of 12% in mean accuracy precision.

The contributions of our framework are as follows.

- We propose a novel preference reasoning framework named PR-Net for personalized video highlight detection.
- We propose a novel bi-directional contrastive loss to train a more compact and differentiable metric space for frame-level content understanding.
- Our method significantly outperforms state-of-the-art methods on a large-scale dataset.

2. Related Work

Personalized video highlight detection aims to detect the interesting moments given a new video based on the previously highlighted video segments the user is interested in. The task is closed to video summarization to find all key events that represent the whole video. Our method is also related to the personalized recommendation if treating the video frame as the item to recommend.

2.1. Video Highlight Detection

Many prior methods [16, 13, 33, 41, 42] have been proposed to solve generic video highlight detection. They typically learn a ranking model to measure the degree of the highlight of video segments without the user’s preference. Only two methods are predicting the user’s preference from history for personalized video highlight detection. Molino and Gygli [7] apply a shot detection algorithm [10] for sampling a set of video segments. And then, directly concatenate the history feature and the video segment feature into a fused feature for highlight prediction. However, the usage of the shot boundary detector makes their pipeline computationally complex and expensive. Rochan et al. [28] employ a temporal-adaptive instance normalization that explicitly considers the user’s history while generating highlights. These two methods share the same weakness when comparing our method. They regard the user’s history as holistic information to predict the user’s preference but negating the inherent diversity of the user’s interests. Besides, they train all unlabeled video frames as non-highlight samples, resulting in wrong label assignment because some unlabeled frames are semantic similar to highlight samples.

2.2. Video Summarization

Unlike video highlight detection that aims to find the most interesting moments, video summarization is to select a set of key events to represent the whole video semantically. Early approaches [18, 19, 21, 23, 24, 25, 26, 32, 47, 45, 29] select some keyframes from a video to satisfy the diversity and representativeness properties based on hand-crafted heuristics, which is not accurate and robust enough. Some learning-based method [8, 11, 12, 29, 43, 44, 46, 14] directly learn from human-annotated training data that show superior performance. However, these are generic methods without the user’s preference involved. And it is impractical to train a particular model for every new user. There are also some personalized video summarization methods. They either use meta-data [1, 15, 34] or user textual query [36, 31, 22] to consider user’s preference. While we adopt a more convenient way that exploits visual features from the user’s history to predict the user preference.

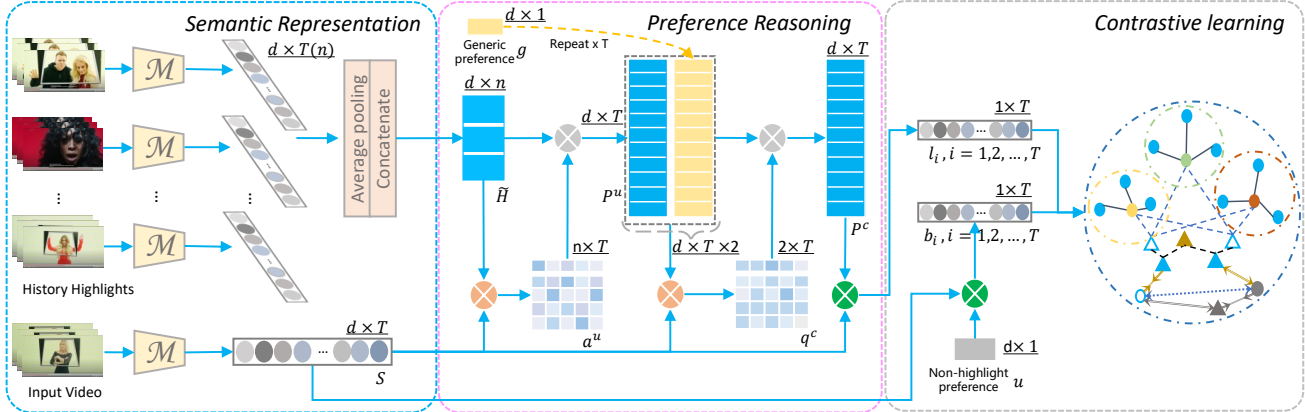


Figure 2. Illustration of our framework. Firstly, each frame’s context semantic representation in input video and history highlights are produced by an encoder \mathcal{M} (left). Secondly, a stacked self-attention module entailing two stages of attention. It first attends to history embedding to form user-specific preference embeddings $P^u = \{p_1^u, p_2^u, \dots, p_T^u\}$, one for each input video frame. The second stage attends to a user-specific preference embedding P^u and the generic preference embedding g to form comprehensive preference embeddings $P^c = \{p_1^c, p_2^c, \dots, p_T^c\}$ for each of the input video frames (middle). Lastly, a bi-directional contrastive loss is proposed to constraint the degree of highlight l_i and non-highlight b_i , which are similarities between each input frame embedding to its comprehensive preference embedding P^c and non-highlight preference embedding u , respectively(right).

2.3. Personalized Recommendation

Personalized recommendation systems are widely utilized in a variety of areas, including movies [27, 6], music [20, 6] and e-commerce [38, 39, 37, 2, 9]. Several methods focus on incorporating knowledge graphs (KGs) into recommender systems for coupling recommendation [27, 39, 37, 2]. Xian et al. [38] utilize reinforcement learning and a policy-guided graph search algorithm to sample reasoning paths for the recommendation. Gao et al. [6] profile the users’ dynamic preferences by a vector autoregressive model. Gu et al. [9] model users’ hierarchical real-time preferences via a pyramid recurrent neural network for e-commerce recommendation. In this paper, we explore a stacked-attention model for preference reasoning in video highlight detection.

3. Method

Given a video V with T frames, we aim at predicting the degree of highlight and non-highlight of each frame under the guidance of the user’s history. Here the history $H = \{h_1, h_2, \dots, h_n\}$ donates n video segments the user is interested in. Note that the frame numbers vary in different videos, and the number of history video segments varies in different users.

Our framework is illustrated in Fig. 2. Firstly, each frame’s contextual semantics embedding in history and the input video are processed by an encoder \mathcal{M} . Secondly, each input video frame’s comprehensive preference, represented as the weighted sum of history and generic preference embedding, is reasoned via a stacked self-attention module. Lastly, each input frame embedding’s similarity to its com-

prehensive preference embedding and non-highlight preference embedding, respectively reflecting the degree of highlight and non-highlight, is constrained by a new proposed bi-directional contrastive loss. In the following, we present details and give more insights into our framework.

3.1. Frame-level Context Semantic Representation

To alleviate the computational burden, following prior work [28], we utilize a feature extractor that pre-trained on Sports 1M dataset [17] to extract C3D [35] (conv5) layer features as the local feature representation of each frame in the input videos and user’s history. However, the C3D feature is fixed and only encodes local appearance and motion features during the training process. The contextual semantics of each frame is also important to estimate the interestingness. For example, a complete action or event is more interesting than the background and the shot boundary. Therefore, we design an encoder \mathcal{M} to capture both the short and long-range features of each frame.

The encoder \mathcal{M} is inspired by U-net [30], the network structure is shown in Fig. 3. The input is a video with T frames represented as C3D features $V \in \mathcal{R}^{d \times T}$, the output $S \in \mathcal{R}^{d \times T}$ is the contextual semantics of each frames in the video. Here d indicates the feature dimension.

$$S = \mathcal{M}(V). \quad (1)$$

3.2. Attention-guided Preference Reasoning

Considering the diversity of user’s preferences is beneficial for an accurate recommendation system. Unfortunately, current highlight detection methods [7, 28] directly learn the holistic preferences from the user’s history, which ignores

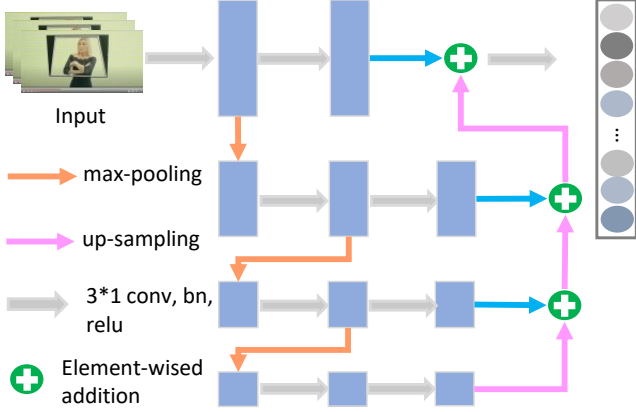


Figure 3. The architecture of the encoder \mathcal{M} . The dimension of all convolutional layers is d , the same as the frame feature.

the inherent diversity of the user's interests. We provide the preference reasoning module details on how to tackle this issue in the following.

Preference reasoning module expects two inputs: the contextual semantic embedding of the input video $S = \{s_1, s_2, \dots, s_T\}$; the contextual semantic embedding of the user's history $\tilde{H} = \mathcal{M}(H) = \{\tilde{h}_1, \tilde{h}_2, \dots, \tilde{h}_n\}$, where $\tilde{h}_j, j = 1, 2, \dots, n$ is the representation for the user history by averaging the features of frames in the video segment.

Considering that a video frame s_i could be interesting if the contextual semantic is similar to one of the user's histories. Therefore, we retrieve the user's most relevant history highlight to obtain the user-specific preference embedding p_i^u for this frame.

$$p_i^u = \sum_{j=1}^n \tilde{h}_j * a_{ij}^u, \quad (2)$$

where $*$ is the multiplication operation, a_{ij}^u is the attention weight and obtained as follows:

$$a_{ij}^u = \frac{\exp(\lambda q_{ij}^u)}{\sum_{j=1}^n \exp(\lambda q_{ij}^u)}, q_{ij}^u = \frac{s_i \cdot \tilde{h}_j}{\|s_i\| \|\tilde{h}_j\|}, \quad (3)$$

where \cdot is the dot product operation. λ is the inversed temperature of the softmax function [4] and set to be 9 empirically.

Next, considering the generic preference is also an important reference to estimate the degree of highlight, especially when the user's history is missing. Therefore, we take the user-specific preference embedding p_i^u and the generic preference embedding g to form the comprehensive preference embedding p_i^c for more overall highlight detection.

$$p_i^c = \frac{p_i^u * q_1^c}{q_1^c + q_2^c} + \frac{g * q_2^c}{q_1^c + q_2^c}, \quad (4)$$




Figure 4. Illustration of the bi-directional contrastive learning. The history embedding is the average feature of frames, represented as the yellow, green and red circle. The user-specific preference is the weighted sum of history embeddings (represented by the blue dotted lines). The attention weights are the similarity between the query frame and the history embeddings. The comprehensive preference is the weighted sum of the user-specific and generic preferences (represented by the black dotted lines). As described in section 3.3, we constraint the relationship of the above notations (represented by the solid lines).

where the generic preference embedding g is learnable and is shared for all frames. The weight terms q_1^c and q_2^c are calculated by following:

$$q_1^c = \exp\left(\frac{\lambda s_i \cdot p_i^u}{\|s_i\| \|p_i^u\|}\right), q_2^c = \exp\left(\frac{\lambda s_i \cdot g}{\|s_i\| \|g\|}\right). \quad (5)$$

We obtain T comprehensive preference embeddings $p_i^c, i = 1, 2, \dots, T$, one for each input frame. The degree of highlight l_i for i -th frame is represented as the cosine similarity of the comprehensive preference embedding p_i^c and its contextual semantic embedding s_i . Note that we omit the normalization term to reduce the computational burden.

$$l_i = s_i \cdot p_i^c. \quad (6)$$

We show a feasible explanation of the results. If the i -th frame is predicted to be interesting, we could check the attention weight q_1^c and q_2^c to see how generic preference and user-specific preference affect the prediction. Moreover, the attention weights a_{ij}^u reflect the importance of history highlights for the i -th frame prediction.

3.3. Bi-directional Contrastive Learning

To ensure a compact and differentiable embedding space, we introduce a learnable non-highlight preference embedding u to measure the degree of non-highlight b_i for the i -th frame.

$$b_i = s_i \cdot u. \quad (7)$$

We constraint the relationship between positive frames and negative frames from two aspects. On the one hand, we set the comprehensive preference embeddings as the anchors, pull in the positive frame embeddings while pushing away those negative frame embeddings. On the other hand, we set the non-highlight embeddings as the anchors, pull in the negative frame embeddings while pushing away those positive frame embeddings (Fig. 4). The loss function is formulated as the following:

$$\mathcal{L} = \sum_{y \in \Omega} \log(\tilde{l}_y + \tilde{b}_y) + \sum_{x \in \mathcal{U}} \log(\tilde{l}_x + \tilde{b}_x) - \left(\sum_{y \in \Omega} \tilde{l}_y + \sum_{x \in \mathcal{U}} \tilde{b}_x \right), \quad (8)$$

where Ω is the set of frames with the positive label. \mathcal{U} is the set of unlabeled frames with the top K highest value of b_i . K is set to be five times the size of Ω empirically. The \tilde{l}_i and \tilde{b}_i is the relative degree of highlight and non-highlight among all frames, respective, calculated as follows.

$$\tilde{l}_i = \frac{\exp(l_i)}{\sum_{i=1}^T \exp(l_i)}, \tilde{b}_i = \frac{\exp(b_i)}{\sum_{i=1}^T \exp(b_i)}. \quad (9)$$

3.4. Implementation Details

To employ multi-GPU training, we cut a video into fixed-length segments (256 frames) and only train those segments that contain highlight frames, and the batch size is set to be 32 (32 users). We take the whole video as input for testing, and the batch size is set to be 1. The entire framework is built on PyTorch 1.50. The training time is about 7 hours for 150 epochs on 8 NVIDIA Tesla V100 GPUs. It is optimized by the adam optimizer using the default configuration. We train our model from scratch except for the C3D feature extractor.

4. Experiments

In this section, the experiments conducted on a large-scale video highlight detection dataset show that our method significantly outperforms current state-of-the-art methods. We also discuss its advantages and limitations. In the end, we conduct ablation experiments to highlight the contributions of each module and present their insights.

4.1. Experiments Configuration

Dataset PHD-GIFs [7] provides a large-scale dataset that contains the user’s history information for video highlight detection. The released dataset consists of 11,972 users for training, 1000 users for validation, and 850 users for testing. Note that there is no overlap among users in these three subsets. For training, each user has at most twenty annotated videos, with one or more segments marked as the highlight in each video. In summary, there are 222,015 annotations in 119,938 videos.



Figure 5. Visualization of the semantically similar frame in the same video. The frames in the yellow box are ground truth highlights, while others are unlabeled frames. We use the width of the yellow line to indicate the similarity between the ground truth and the unlabeled frame. The three videos are swimming, football, and Songs, respectively, show that a video contains similar frames is a common phenomenon in this dataset.

The dataset has some notable properties. First, the annotations cover various topics, such as cartoons, pets, music videos and beautiful scenes. Moreover, the user has consistent interests in one or more topics, which provide reliable guidance for personalized video highlight detection. Secondly, only a small proportion of frames are marked as the highlights in the dataset. It is about 1:50 for positive vs unlabeled. Lastly, we observe that there are many similar shots in a video, but only one of them is annotated as the highlight (Fig 5). The incomplete annotation may cause ambiguity when training the unlabeled frames as negative samples.

Only the YouTube video ID of videos are provided in the dataset for download from the web. Because some video links are no longer available, we only download 95,111 videos with a total duration of 10,212 hours. Following Rochan et al. [28], we extract the frame-level C3D feature from all videos by a pre-trained model. The processing time is about two weeks on 8 NVIDIA Tesla V100 GPUs. We process 9478 users for training, 750 users for validation and 675 users for testing in the end. Since the data processing is tedious, time-consuming and computationally expensive, we are willing to release the C3D features of the whole dataset to the community for future works.

Evaluation Metric The mean Average Precision (mAP) is adopted to measure the performance of our method, which is also commonly used by previous works [41, 28]. Follow Rochan et al. [28] we treat videos separately when calculating the mAP.



Figure 6. Qualitative evaluation of different configurations of our method. We show examples of our method with only generic preference (PR-Net_G), with only user-specific preference (PR-Net_U) and our full method (PR-Net_{full}) on four videos. In addition, we show the user’s history (multiple GIFs) and a few sampled frames from the highlight predictions of the three models for each video. These visualizations indicate that the user-specific preference well captures the user’s diverse interests from history, resulting in more reasonable predictions than only generic preference.

Baselines Our method is compared with other state-of-the-art methods on the dataset, including FCSN [29], Video2GIF [13], PHD-GIFs [7] and A-VHD [28]. FCSN is the state-of-the-art video summarization method, which is generic video highlight detection for comparison. Video2GIF is a state-of-the-art model for video highlight detection. Following Rochan et al. [28], we use the publicly available pre-trained model for comparison. PHD-GIFs is a state-of-the-art personalized video highlight detection method. We aggregate frames into the five-second shots in every video to report the performance for a fair comparison. A-VHD is also a state-of-the-art method for personalized video highlight detection.

4.2. Results and Discussions

We quantitatively compare our method with other state-of-the-art methods. Moreover, its explainability, advantages and limitations are presented in the following.

Comparison results The results show that our method significantly outperforms current state-of-the-art methods (Table 1). Compared with two generic video highlight detection methods (Video2GIF and FCSN) without using

Table 1. Comparison with state-of-the-art methods on the PHD-GIFs dataset. The performance of Video2GIF, FCSN* and A-VHD* are reported in [28]. While others are evaluated on our dataset. PR-Net is our full method.

Method	mAP(%)
Video2GIF [13]	14.75
FCSN* [29]	15.22
A-VHD* [28]	16.73
FCSN [29]	15.15
PHD-GIFs [7]	16.25
A-VHD [28]	16.68
PR-Net	18.66

the user’s history, our method efficiently utilizes the user’s preference for guiding the highlight detection and achieves much more accurate results. The mAP is relatively improved by 28.24% and 23.17%, respectively. PHD-GIFs directly concatenates the average history feature into the input video feature to perform segment-level prediction. A-VHD predicts the affine parameters of the adaptive instance normalization from the user’s history, where the predicted affine parameters can be regarded as the user’s preference

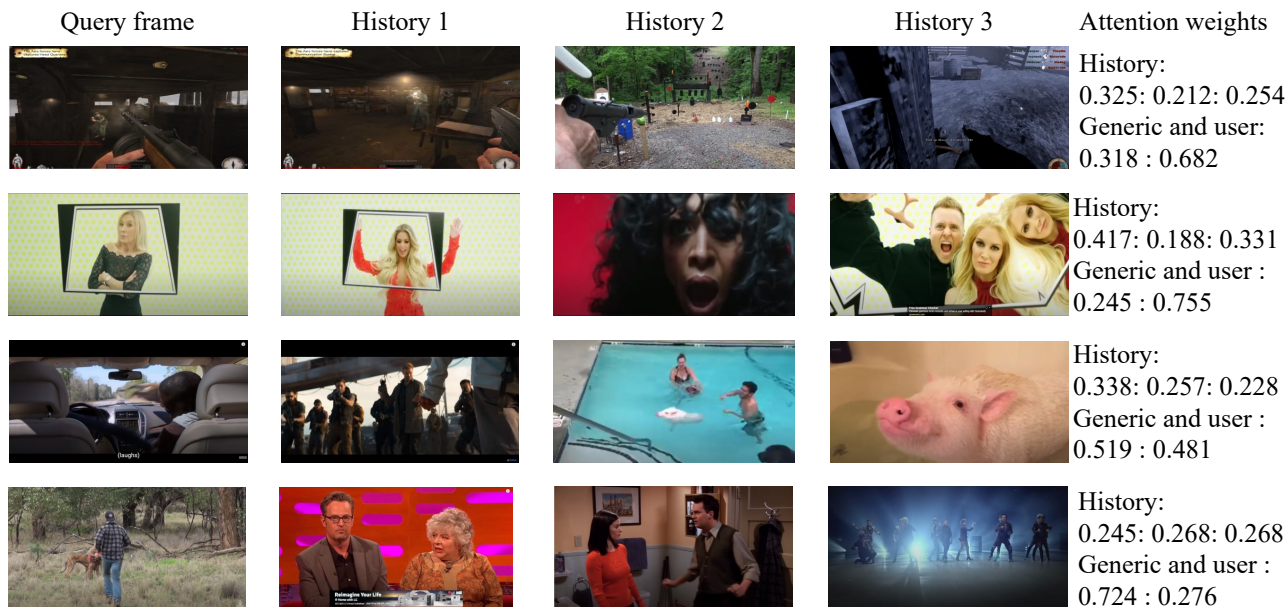


Figure 7. Attention visualization for the detection results. We select four cases under different scenarios to present three types of attention weights (for history embeddings, user-specific preference embedding and generic preference embedding). Each row stands for a case. Note that we only show the top three history highlights for each case, and the remaining histories are not meaningful than the top 3. In case 1 (the first row), all history highlights are semantically related to the query frame. Case 2 presents a highly subjective scenario in which the semantics are ambitious in the history and query frame. Case 3 and Case 4 show that the user’s interests are diverse. However, there is no relevant history with the query frame in case 4.

for guiding personalized video highlight detection. However, the limitation of the two methods is that the user preference is determined by the holistic history highlights, implying only one interest in the user’s history. Therefore, they may not be suitable solutions when user’s interests are diverse. While in our method, we overcome the limitation by reasoning the most relevant history highlight for each input frame, results in more accurate detection results, with the relative improvement of 12%.

Attention Visualization The personalized video highlight detection is subjective and highly depends on the user’s preference. Intuitively, the query video frame could be regarded as a highlight if it is semantically similar to one of the user’s highlight histories. Our method utilizes the attention mechanism to find the most relevant highlight history for the query video frame. After checking many cases, we discovered that the learned attention weights almost reflect the semantic similarity between the query and history frames, which indicates a potential way to explain the detection results. As shown in Fig. 7, we visualize the results with the top three most relevant history highlights of four cases.

In case 1 of the query frame, it is shooting an enemy with a machine gun. Observing the attention weight, we find the history with the highest attention weight is most semanti-

cally similar to the query frame, shooting an enemy with a pistol. For history 2 and 3, shooting without the enemy, the attention weight is lower than history 1.

Case 2 is highly subjective, and the semantics is ambitious. It seems that all history highlights are relevant to the query frame. However, we can observe that the most relevant history share the same background with the query frame.

The user’s interests are diverse in case 3. It seems that the user prefers a movie, swimming and pets. We can see that the attention weight reflects semantic similarities. Moreover, the generic weight is raised compared with case 1 and case 2 because the query frame is not highly relevant to the historical highlights.

There are no relevant history highlights in case 4 from our view. In this case, the generic preference significantly affects the degree of the query frame’s highlight.

Limitations and Future Work The detection results of our method are sometimes confusing, which is hard to tell it is interesting or not. We think the subjectiveness of the user annotation mainly causes it as the model is directly learning from the training data. Further enhancing the semantic representation by multiple feature extractors may be a promising way to solve the problem. We left it in future work.

Table 2. Ablation Study Experiments. PR-Net(*) stand for the different modification based on our method. (Ablation experiments section). PR-Net_{full} is our full method.

ablation target	method	mAP(%)
Representation	PR-Net _{residual}	16.69
	PR-Net _{ranking}	16.85
Loss Strategy	PR-Net _{CE}	17.09
	PR-Net _{triplet}	17.26
Attention Strategy	PR-Net _G	16.78
	PR-Net _U	17.57
	PR-Net _{meanU}	16.87
	PR-Net _{meanU+G}	17.21
History Size	PR-Net _{h0}	16.78
	PR-Net _{h5}	17.48
-	PR-Net	18.66

4.3. Ablation Experiments

We conduct extensive experiments to highlight the effectiveness of different modules in our framework.

PR-Net_{residual}: Replacing the encoder \mathcal{M} with a residual block, in which the number of the parameters is larger than \mathcal{M} . The residual block is applied to the C3D feature of each frame.

PR-Net_{ranking}, PR-Net_{CE} and PR-Net_{triplet}: We respectively replace the bi-directional contrastive loss with the ranking loss (PR-Net_{ranking}) and the Cross-Entropy loss (PR-Net_{CE}) as they are also used in other highlight detection method [41, 28]. Note that for the implementation of PR-Net_{ranking}, we abandon the non-highlight preference and only measure each frame’s degree of highlight. We also present the performance when using the triplet loss [3] (PR-Net_{triplet}).

PR-Net_G and PR-Net_{meanU}: To evaluate the effect of user-specific preference, two variants of the PR-Net are studied: Directly removing the user-specific preference (PR-Net_G), using the mean history embeddings to replace the user-specific preference (PR-Net_{meanU}).

PR-Net_U and PR-Net_{meanU+G}: For the evaluation of the generic preference, we explore the two variants of the PR-Net: without generic preference (PR-Net_U), adding generic preference on PR-Net_{meanU} variant (PR-Net_{meanU+G}).

PR-Net_{h0} and PR-Net_{h5}: Two experiments are conducted to evaluate the effect of history: no history information involved (PR-Net_{h0}) and restrict the history size with five PR-Net_{h5}.

Effect of frame-level contextual semantics representation To verify how contextual semantic representation affects the performance, we only consider the short-range features of each frame (PR-Net_{residual}). Compared with our full method (PR-Net_{full}), the performance is dramatically

decreased (from 18.66% to 16.69%). The result shows that contextual semantic representation is vital for highlight detection.

Effect of bi-directional contrastive loss Compared with PR-Net_{ranking}, PR-Net_{CE} and PR-Net_{triplet} that respectively use ranking loss, Cross entropy loss and triplet loss, our carefully designed bi-directional contrastive loss provide more reasonable constraints in the metric space for video highlight detection, resulting in better performance.

Effect of generic preference We remove the generic preference embedding from our method and only utilize the user-specific preference to measure the degree of highlight (PR-Net_U, and PR-Net_{meanU}). Note that they are already outperforming the the-state-of-the-art method. However, because they do not consider the generic preference, especially when the relevant history is missing, the performance is worse than that with generic preference (PR-Net_{meanU+G} and PR-Net_{full})

Effect of user-specific preference It becomes the generic highlight detection model if without user-specific preference (PR-Net_G). As shown in Fig. 6, the detection results of using user-specific preference are obviously closer to history than those without it. Experiment results show that frame-level user-specific preference is vital for boosting performance and lifting the mAP from 16.78% to 18.66%. Compared with PR-Net_{meanU} that adopt mean history embeddings as the user-specific preference (holistic preference), the usage of attention mechanism to produce user-specific preference is functional, which implies that our method offers a more efficient way to utilize the user’s history for personalized video highlight detection.

Effect of history size Theoretically, the performance of our method is benefits from more history highlights for reference. Therefore, we also show the performances when restrict the maximal size of history to be zero (PR-Net_{h0}) and five (PR-Net_{h5}), respectively, in which the performance is to meet our expectation.

5. Conclusion

In this paper, we explore a new paradigm incorporating the user’s history for personalized video highlight detection. Unlike prior methods to predict a user preference from holistic history information, our attention-guided preference reasoning takes the diversity of user’s interests into account to predict frame-level preferences. Experiments conducted on the large-scale dataset show that our method is much more accurate than state-of-the-art methods.

References

- [1] N. Babaguchi, K. Ohara, and T. Ogura. Learning personal preference from viewer's operations for browsing and its application to baseball video retrieval and summarization. *IEEE transactions on multimedia*, 9(5):1016–1025, 2007. [2](#)
- [2] Y. Cao, X. Wang, X. He, Z. Hu, and T.-S. Chua. Unifying knowledge graph learning and recommendation: Towards a better understanding of user preferences. In *The world wide web conference*, pages 151–161, 2019. [3](#)
- [3] D. Cheng, Y. Gong, S. Zhou, J. Wang, and N. Zheng. Person re-identification by multi-channel parts-based cnn with improved triplet loss function. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 1335–1344, 2016. [8](#)
- [4] J. Chorowski, D. Bahdanau, D. Serdyuk, K. Cho, and Y. Bengio. Attention-based models for speech recognition. In *Proceedings of the 28th International Conference on Neural Information Processing Systems-Volume 1*, pages 577–585, 2015. [4](#)
- [5] B. Fréney and M. Verleysen. Classification in the presence of label noise: a survey. *IEEE transactions on neural networks and learning systems*, 25(5):845–869, 2013. [2](#)
- [6] L. Gao, J. Wu, C. Zhou, and Y. Hu. Collaborative dynamic sparse topic regression with user profile evolution for item recommendation. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 31, 2017. [3](#)
- [7] A. Garcia del Molino and M. Gygli. Phd-gifs: personalized highlight detection for automatic gif creation. In *Proceedings of the 26th ACM international conference on Multimedia*, pages 600–608, 2018. [1](#), [2](#), [3](#), [5](#), [6](#)
- [8] B. Gong, W.-L. Chao, K. Grauman, and F. Sha. Diverse sequential subset selection for supervised video summarization. *Advances in neural information processing systems*, 27:2069–2077, 2014. [2](#)
- [9] Y. Gu, Z. Ding, S. Wang, and D. Yin. Hierarchical user profiling for e-commerce recommender systems. In *Proceedings of the 13th International Conference on Web Search and Data Mining*, pages 223–231, 2020. [3](#)
- [10] M. Gygli. Ridiculously fast shot boundary detection with fully convolutional neural networks. In *2018 International Conference on Content-Based Multimedia Indexing*, pages 1–4. IEEE, 2018. [2](#)
- [11] M. Gygli, H. Grabner, H. Riemenschneider, and L. Van Gool. Creating summaries from user videos. In *European conference on computer vision*, pages 505–520. Springer, 2014. [2](#)
- [12] 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*, pages 3090–3098, 2015. [2](#)
- [13] 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*, pages 1001–1009, 2016. [1](#), [2](#), [6](#)
- [14] 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*, pages 345–360. Springer, 2020. [2](#)
- [15] A. Jaimes, T. Echigo, M. Teraguchi, and F. Satoh. Learning personalized video highlights from detailed mpeg-7 metadata. In *Proceedings. International Conference on Image Processing*, volume 1, pages I–I. IEEE, 2002. [2](#)
- [16] 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*, pages 28–39. Springer, 2017. [2](#)
- [17] A. Karpathy, G. Toderici, S. Shetty, T. Leung, R. Sukthankar, and L. Fei-Fei. Large-scale video classification with convolutional neural networks. In *Proceedings of the IEEE conference on Computer Vision and Pattern Recognition*, pages 1725–1732, 2014. [3](#)
- [18] A. Khosla, R. Hamid, C.-J. Lin, and N. Sundaresan. Large-scale video summarization using web-image priors. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 2698–2705, 2013. [2](#)
- [19] G. Kim and E. P. Xing. Reconstructing storyline graphs for image recommendation from web community photos. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 3882–3889, 2014. [2](#)
- [20] H.-G. Kim, G. Y. Kim, and J. Y. Kim. Music recommendation system using human activity recognition from accelerometer data. *IEEE Transactions on Consumer Electronics*, 65(3):349–358, 2019. [3](#)
- [21] Y. J. Lee, J. Ghosh, and K. Grauman. Discovering important people and objects for egocentric video summarization. In *2012 IEEE conference on computer vision and pattern recognition*, pages 1346–1353. IEEE, 2012. [2](#)
- [22] 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*, pages 3707–3715, 2015. [2](#)
- [23] Z. Lu and K. Grauman. Story-driven summarization for egocentric video. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 2714–2721, 2013. [2](#)
- [24] 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*, pages 202–211, 2017. [2](#)
- [25] C.-W. Ngo, Y.-F. Ma, and H.-J. Zhang. Automatic video summarization by graph modeling. In *Proceedings Ninth IEEE International Conference on Computer Vision*, pages 104–109. IEEE, 2003. [2](#)
- [26] 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*, pages 7083–7092, 2017. [2](#)
- [27] F. Pecune, S. Murali, V. Tsai, Y. Matsuyama, and J. Cassell. A model of social explanations for a conversational movie recommendation system. In *Proceedings of the 7th International Conference on Human-Agent Interaction*, pages 135–143, 2019. [3](#)
- [28] M. Rochan, M. K. K. Reddy, L. Ye, and Y. Wang. Adaptive video highlight detection by learning from user history. In *European Conference on Computer Vision*, pages 261–278. Springer, 2020. [1](#), [2](#), [3](#), [5](#), [6](#), [8](#)

- [29] M. Rochan, L. Ye, and Y. Wang. Video summarization using fully convolutional sequence networks. In *Proceedings of the European Conference on Computer Vision*, pages 347–363, 2018. 2, 6
- [30] O. Ronneberger, P. Fischer, and T. Brox. U-net: Convolutional networks for biomedical image segmentation. In *International Conference on Medical image computing and computer-assisted intervention*, pages 234–241. Springer, 2015. 3
- [31] A. Sharghi, B. Gong, and M. Shah. Query-focused extractive video summarization. In *European Conference on Computer Vision*, pages 3–19. Springer, 2016. 2
- [32] 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*, pages 5179–5187, 2015. 2
- [33] M. Sun, A. Farhadi, and S. Seitz. Ranking domain-specific highlights by analyzing edited videos. In *European conference on computer vision*, pages 787–802. Springer, 2014. 2
- [34] Y. Takahashi, N. Nitta, and N. Babaguchi. User and device adaptation for sports video content. In *2007 IEEE International Conference on Multimedia and Expo*, pages 1051–1054. IEEE, 2007. 2
- [35] D. Tran, L. Bourdev, R. Fergus, L. Torresani, and M. Paluri. Learning spatiotemporal features with 3d convolutional networks. In *Proceedings of the IEEE international conference on computer vision*, pages 4489–4497, 2015. 3
- [36] D. Ulyanov, A. Vedaldi, and V. Lempitsky. Improved texture networks: Maximizing quality and diversity in feed-forward stylization and texture synthesis. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 6924–6932, 2017. 2
- [37] X. Wang, D. Wang, C. Xu, X. He, Y. Cao, and T.-S. Chua. Explainable reasoning over knowledge graphs for recommendation. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 33, pages 5329–5336, 2019. 3
- [38] Y. Xian, Z. Fu, S. Muthukrishnan, G. De Melo, and Y. Zhang. Reinforcement knowledge graph reasoning for explainable recommendation. In *Proceedings of the 42nd international ACM SIGIR conference on research and development in information retrieval*, pages 285–294, 2019. 3
- [39] Y. Xian, Z. Fu, H. Zhao, Y. Ge, X. Chen, Q. Huang, S. Geng, Z. Qin, G. De Melo, S. Muthukrishnan, et al. Cafe: Coarse-to-fine neural symbolic reasoning for explainable recommendation. In *Proceedings of the 29th ACM International Conference on Information & Knowledge Management*, pages 1645–1654, 2020. 3
- [40] B. Xiong, Y. Kalantidis, D. Ghadiyaram, 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*, pages 1258–1267, 2019. 1
- [41] 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*, pages 982–990, 2016. 1, 2, 5, 8
- [42] Y. Yu, S. Lee, J. Na, J. Kang, and G. Kim. A deep ranking model for spatio-temporal highlight detection from a 360-degree video. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 32, 2018. 2
- [43] K. Zhang, W.-L. Chao, F. Sha, and K. Grauman. Summary transfer: Exemplar-based subset selection for video summarization. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 1059–1067, 2016. 2
- [44] K. Zhang, W.-L. Chao, F. Sha, and K. Grauman. Video summarization with long short-term memory. In *European conference on computer vision*, pages 766–782. Springer, 2016. 2
- [45] K. Zhang, K. Grauman, and F. Sha. Retrospective encoders for video summarization. In *Proceedings of the European Conference on Computer Vision*, pages 383–399, 2018. 2
- [46] B. Zhao, X. Li, and X. Lu. Hierarchical recurrent neural network for video summarization. In *Proceedings of the 25th ACM international conference on Multimedia*, pages 863–871, 2017. 2
- [47] K. Zhou, Y. Qiao, and T. Xiang. Deep reinforcement learning for unsupervised video summarization with diversity-representativeness reward. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 32, 2018. 2