

HighlightMe: Detecting Highlights from Human-Centric Videos

Uttaran Bhattacharya^{*1}, Gang Wu², Stefano Petrangeli², Viswanathan Swaminathan², and Dinesh Manocha¹

¹University of Maryland, College Park, MD, USA.

{uttaranb|dmanocha}@umd.edu

²Adobe Research, San Jose, CA, USA.

{gawu|petrange|vishy}@adobe.com

Abstract

We present a domain- and user-preference-agnostic approach to detect highlightable excerpts from human-centric videos. Our method works on the graph-based representation of multiple observable human-centric modalities in the videos, such as poses and faces. We use an autoencoder network equipped with spatial-temporal graph convolutions to detect human activities and interactions based on these modalities. We train our network to map the activity- and interaction-based latent structural representations of the different modalities to per-frame highlight scores based on the representativeness of the frames. We use these scores to compute which frames to highlight and stitch contiguous frames to produce the excerpts. We train our network on the large-scale AVA-Kinetics action dataset and evaluate it on four benchmark video highlight datasets: DSH, TVSum, PHD², and SumMe. We observe a 4–12% improvement in the mean average precision of matching the human-annotated highlights over state-of-the-art methods in these datasets, without requiring any user-provided preferences or dataset-specific fine-tuning.

1. Introduction

Human-centric videos focus on human activities, tasks, and emotions [62, 50]. These videos form a major part of the rapidly growing volume of online media [8], coming from multiple domains, such as amateur sports and performances, lectures, tutorials, video weblogs (vlogs), and individual or group activities, e.g., cookouts and holiday trips. However, unedited human-centric videos also tend to contain large chunks of irrelevant and uninteresting content, requiring them to be edited for efficient browsing [47].

To address this problem, researchers have developed multiple techniques for detecting highlightable excerpts and summarizing videos [11, 53, 42, 63, 44, 67]. Given unedited footage, highlight detection obtains the moments of interest,

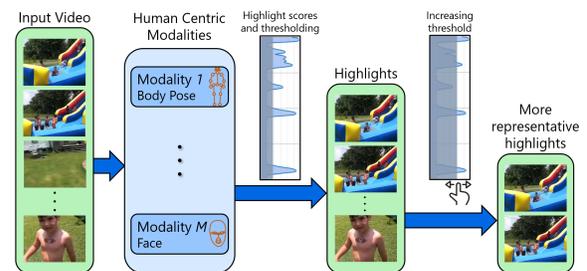


Figure 1: **Detecting highlight excerpts using human-centric modalities.** Our method leverages multiple human-centric modalities, e.g., body poses and faces, observable in videos focusing on human activities, to detect highlights. We use a 2D or 3D interconnected point representation of each modality to construct a spatial-temporal graph representation to compute the highlight scores.

and summarization computes the most relevant and representative set of excerpts. Detecting effective highlights not only expedites browsing, but also improves the chances of those highlights being shared and recommended [53]. Current methods can learn to detect these excerpts given annotated highlights [47, 11], or sets of exemplars for different highlight categories, e.g., learning from skiing images to detect skiing excerpts from videos [23, 25]. Other methods obviate the need for supervision by learning the representativeness of each frame or shot with respect to the original video [36] and exploiting video metadata such as duration [53] and relevance of shots [67, 64]. All these methods either assume or benefit from some domain-specific knowledge of the unedited footage, e.g., running and jumping may be more relevant in a parkour video, whereas sliding maneuvers may be more relevant in a skiing video. Alternative methods do not consider domain-specific knowledge but consider the pre-recorded preferences of multiple users instead to detect personalized highlights [42].

Whether they assume domain-specific knowledge or user-preferences, existing methods work in the 2D image space of the frames or shots constituting the videos. State-

^{*}Work done while Uttaran an intern at Adobe Research

of-the-art image-based networks can learn rich semantic features capturing the interrelations between the various detected objects in the images, leading to efficient highlight detection. However, these approaches do not explicitly model human activities or inter-person interactions that are the primary focus of human-centric videos. Developing methods for human-centric videos, meanwhile, has been essential for a variety of tasks, including expression and emotion recognition [34, 2, 38], activity recognition [56], scene understanding [50, 32], crowd analysis [51], video super-resolution [32], and text-based video grounding [48]. These methods show that human-centric videos need to be treated separately from generic videos, by leveraging human-centric modalities such as poses and faces. Therefore, there is both the scope and the need to bring the machineries of human-centric video understanding to the task of highlight detection as well.

Main contributions. We develop an end-to-end learning system that detects highlights from human-centric videos without requiring domain-specific knowledge, highlight annotations, or exemplars. Our approach utilizes the human activities and interactions that are expressed through multiple sensory channels or modalities, including faces, eyes, voices, body poses, and hand gestures [1, 38]. We use graph-based representations for all the human-centric modalities to sufficiently represent how the inherent structure of each modality evolves with various activities and interactions over time. Our network learns from these graph-based representations using spatial-temporal graph convolutions and maps the per-frame modalities to *highlight scores* using an autoencoder architecture. Our highlight scores are based on the representativeness of all the frames in the videos, and we stitch together contiguous frames to produce the final excerpts. Our novel contributions include:

- **Highlight detection with human-centric modalities.** Our method identifies the observable modalities, such as poses and faces, in each input video and encodes their inter-relations, across both time and different persons, into *highlight scores* for highlight detection.
- **Annotation-free training of highlight scores.** We do not require highlight annotations, exemplars, user-preferences, or domain-specific knowledge. Instead, we only need to detect of one or more human-centric modalities using off-the-shelf modality detection techniques to train our highlight scores.
- **Domain- and user-agnostic performance.** Our trained network achieves state-of-the-art performance in highlight detection over a diverse range of domains and user preferences, evaluated over multiple benchmark datasets consisting of human-centric videos.

Our method achieves a mean average precision of 0.64 and 0.20 of matching human-annotated highlight excerpts on the benchmark domain-specific video highlight (DSH) dataset [47] and the personal highlight detection dataset (PHD²) [11] dataset, respectively, and outperform the corresponding state-of-the-art methods by 7% and 4% (absolute).

We also achieve state-of-the-art performance on the smaller benchmark datasets of TVSum [46] and SumMe [15], outperforming the current state-of-the-art baselines by 12% and 4% (absolute) on the mean average precision and mean F-score, respectively. Even for domains that are not fully human-centric (*e.g.*, dog shows) or videos where human-centric modalities are sparsely detected, the performance of our method is comparable to the current state-of-the-art.

2. Related Work

Both highlight detection and the closely related problem of video summarization have been well-studied in computer vision, multimedia, and related fields. Early methods utilized a variety of techniques including visual-content-based clustering, scene transition graphs, temporal variance of frames [59, 6, 49], and hand-crafted features representing semantic information such as facial activities [20]. On the other hand, recent approaches have capitalized on an impressive range of deep learning tools and techniques to perform highlight detection and video summarization.

Highlight Detection. The goal of highlight detection is to detect interesting moments or excerpts from unedited videos [49, 47]. A large contingent of methods pose this as a supervised ranking problem, such that the highlightable excerpts are ranked higher than all other excerpts [47, 17, 58, 18, 11, 60, 19, 52]. These methods assume the availability of human-annotated labels of the highlightable excerpts and train networks to learn either generic or domain-specific ranking metrics that correlate with these labels. On the other hand, weakly-supervised and unsupervised highlight detection methods eliminate label dependencies by leveraging exemplars or video metadata. Exemplars include scraped web images depicting domain-specific actions such as gymnastics and skiing [25]. Video metadata include information on video categories [57], or properties useful in differentiating unedited videos from edited videos, *e.g.*, duration [53]. Some approaches also take user preferences into account to generate personalized highlights [42]. All these methods perform computations in the 2D image space of the frames and do not utilize human-centric modalities.

Video Summarization. Video summarization aims to provide succinct synopses of videos in a variety of formats, including storyline graphs [24, 54], keyframe sequences [30], clips [15, 64], and their mixtures based on user requirements [14]. It is commonly posed as a subsequence estimation task satisfying coherence [35], diversity, and representativeness [40, 67]. Existing unsupervised approaches build on multiple concepts, such as visual co-occurrence [7], temporal relevance between frames and shots [23, 36, 44, 64], learning category-aware classifiers [41] and category-aware feature learning [66, 46]. Weakly supervised approaches use exemplar web images and videos [24, 22, 4, 43], and category descriptions [41, 40] as priors. Other approaches use supervised learning with human-annotated summaries, using subset selec-

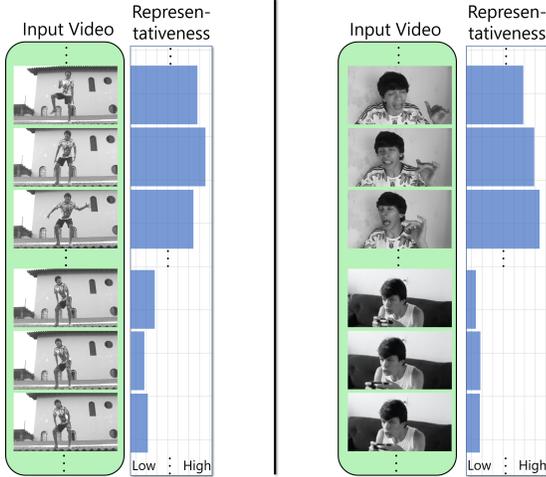


Figure 2: **Representativeness.** We show frames with different values of representativeness calculated in the space of poses (*left*) and face landmarks (*right*). We learn highlight scores based on the representativeness.

tion [13], visual importance scores [30, 15], submodular mixtures [16, 55], and temporal inter-relations [63, 64, 65]. While our objective is highlight detection, our approach is inspired by these summarization methods. Particularly, we ensure that our highlight score captures the representativeness in the videos and satisfies robust feature reconstruction.

Multimodal Learning. A wide body of work has focused on multimodal action recognition [5, 45, 33, 10] and emotion recognition [3, 26, 61, 38, 39]. These methods observe and combine cues from multiple modalities of human expression, including faces, poses, vocal tones, eye movements hand and body gestures, and gaits. Existing methods commonly model observed modalities using points and graphs [33, 3, 38], making them suitable for learning action- and emotion-specific features. In our work, we utilize the fact that highlightable excerpts of human-centric videos can be determined based on the modalities. Following recent trends in multimodal action and emotion recognition [33, 38], we also model the modalities observed across the frames in videos as spatial temporal graphs, and leverage them to learn our highlight scores.

3. Multimodal Highlight Detection

Given human-centric videos, our goal is to detect the moments of interest or *highlights* from the videos. This section elaborates on how we detect such highlights by leveraging the human-centric modalities observed from the videos.

3.1. Human-Centric Modalities

In our work, we use the term *modalities* to imply the channels of human expression sensitive to human activities and interactions, *e.g.*, faces, eyes, body poses, hands, and gaits [5, 38, 39]. Activities constitute individual expressions and interactions occur with other humans, living beings, and

inanimate objects, pertinent to a variety of actions [56, 10] and emotions [2, 39]. We argue that the highlightable excerpts of human-centric videos preferred by human users focus on these activities and interactions. Therefore, we aim to learn from the observable human-centric modalities in our network. For each detected modality of each human, our network leverages the inter-relations at different time instances and the inter-relations between different humans to detect the most representative excerpts.

While we extract these modalities from the RGB image-space of the video frames, we note that the modalities better capture the rich semantics of the frames. Image-space representations build on variants of the intensity differences between different parts of images, without an underlying insight on how the different parts physically interact. Conversely, modalities provide insight on such interactions based on their structure, *e.g.*, the relative movements of arms and legs indicate certain actions, and the relative movements of various facial landmarks indicate certain expressions and emotions. We build our network to explicitly consider the structure of each modality and the evolution of those structures with activities and interactions over time.

We consider $M \geq 1$ observable human-centric modalities from an input video. We assume the modalities are extracted using standard detection and tracking techniques [29, 12], and are represented using a set of interconnected points in 2D or 3D, such as a set of 2D face landmarks for the face or a set of 3D body joints for the pose.

To represent each modality $m = 1, \dots, M$, we construct a spatial-temporal graph representation $\mathcal{G}_m = \{\mathcal{V}_m, \mathcal{E}_m\}$. The nodes in \mathcal{V}_m represent the points of the corresponding modality, and the edges in \mathcal{E}_m represent both the structure of the modality and how that structure evolves over time. To sufficiently capture this, we consider three edges types:

- **Intra-person edges** capturing the spatial relationships between the nodes of a single person, *e.g.*, bones between pose joints and connectors between face landmarks. These edges represent the baseline structure of the modality at every video frame.
- **Inter-person edges** connecting the identical nodes of different persons, *e.g.*, root to root, head to head, at every video frame. These edges capture how the nodes of different persons interact with each other. They form a bipartite graph for every pair of persons, and represent the inter-person interactions at every video frame.
- **Temporal edges** connecting the identical nodes of a person, *e.g.*, root to root, head to head, over multiple video frames. These edges capture how those nodes evolve with time for every person. They form a bipartite graph for every pair of video frames, and represent the evolution of activities and interactions over time.

The spatial positions of these nodes and the combination of all these edges allow our network to learn the activities and interactions of all humans in videos and learn highlight scores accordingly, without any prior knowledge on the video domains or user-provided preferences.

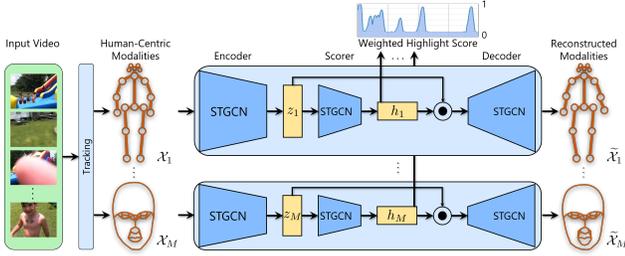


Figure 3: **Highlight detection with human-centric modalities:** Overview of our network for learning highlight scores from multiple human-centric modalities. We use standard techniques [29, 12] to detect the human-centric modalities. We represent the modalities as sets of connected points in either 2D or 3D. We train the networks for all the modalities in parallel. The only point of interaction between the networks is their predicted highlight scores, which we combine into our weighted highlight score for training.

3.2. Representativeness of the Video Frames

Since we aim to detect highlights from videos without requiring annotations or exemplars, our approach is aligned with detecting the representative frames from the videos, similar to video summarization [36, 14]. While detecting the representative frames in the image-space may or may not lead to detecting the moments of interest for highlight detection [47], our key observation is that detecting the representative frames in the space of human-centric modalities, in fact, leads to detecting the moments of interest for highlight detection in human-centric videos.

We define the *representativeness* of a video frame as the difference, in some metric space, between the video and the video without that frame. The larger the difference, the higher the representativeness of that frame. Intuitively, the representativeness of a frame measures the fraction of information it contains in relation to the entire video. Our goal in highlight detection is to detect a *minimal* set of frames from a video with a *maximal* representativeness.

In our work, we measure the representativeness in the metric space of the observable modalities. Figure 2 shows examples of frames with different values of representativeness in the space of poses and face landmarks. We consider each video to consist of a total T frames and P persons (zero-padding videos with fewer frames). Therefore, for each modality m , \mathcal{V}_m consists of $N \times T \times P$ nodes in total, where N is the number of nodes per person. We collate these nodes into a tensor $\mathcal{X}_m = [\mathcal{X}_m^{(1)}; \dots; \mathcal{X}_m^{(T)}]$, where $\mathcal{X}_m^{(t)} \in \mathbb{R}^{N \times P \times D}$ for each frame t , and D is the spatial dimension of each node, most commonly 2 or 3. We can then multiply a *highlight score* $h_m^{(t)}$ of 0 or 1 to each frame t to reflect their representativeness. Thus, we can write the net difference \mathcal{D} as a result of assigning the highlight scores as,

$$\mathcal{D} = \left\| \mathcal{X}_m - \left[h_m^{(1)}; \dots; h_m^{(T)} \right] \odot \mathcal{X}_m \right\|, \quad (1)$$

where \odot denotes the Hadamard product. We can now rewrite our goal as simultaneously minimizing \mathcal{D} in Eq. 1 and $\sum_t h_m^{(t)}$, for each modality m .

We note that a trivial solution to Eq. 1 is to pick a threshold $0 \leq \tau \leq T$, then assign a highlight score of 1 to the top τ most representative frames from \mathcal{X}_m , and a highlight score of 0 to all other frames. However, the choice τ is non-trivial and needs to be learned from the data in practice. Therefore, we train an autoencoder-based deep neural network to learn the highlight scores for a wide range of data. We also allow the highlight scores to be continuous in $[0, 1]$ to keep our network differentiable. Moreover, making the highlight scores continuous also helps us understand the relative representativeness of each frame, which is an inbuilt component of modern highlight detection systems [53, 42].

3.3. Network Architecture

Figure 3 shows our overall network architecture for predicting highlight excerpts from input videos. The goal of our network is to learn per-frame highlight scores to minimize an analogous form of Eq. 1. Our network achieves this by taking in the per-frame graph-based representations of the observable human-centric modalities. It attempts to reconstruct all the activities in the video using as few frames of the input modalities as possible, *i.e.*, a weighted reconstruction, where the weights are the highlight scores. In this training process, our network learns to assign higher highlight scores to the frames with higher representativeness. We now describe our network architecture in detail.

Our autoencoder consists of an encoder, a scorer, and a decoder. Our encoder takes in the spatial-temporal graph $\mathcal{G}_m = \{\mathcal{V}_m, \mathcal{E}_m\}$ for each observable modality m from an input video. It uses a separate spatial temporal graph convolutional network (STGCN) [56, 28, 9] to transform \mathcal{X}_m of each modality m into a latent activity-based feature $z_m \in \mathbb{R}^{N \times T \times P \times D_l}$, D_l being the dimension of each node in the latent feature. We thus have the operation,

$$z_m = \text{STGCN} \left(\mathcal{A}_m, \mathcal{X}_m; W_m^{(\text{enc})} \right), \quad (2)$$

where \mathcal{A}_m denotes the adjacency matrix obtained from \mathcal{E}_m , and $W_m^{(\text{enc})}$ consists of the set of trainable STGCN parameters in the encoder. We note here that the data \mathcal{X}_m forms a full-rank tensor, therefore the STGCN avoids the degenerate solution of assigning 0's to all z_m 's.

Our latent activity-based features $z_m \forall m$ connect to our scorer, which consists of a single layer of spatial temporal graph convolution followed by a sigmoid operation per modality. Our scorer transforms each z_m into a normalized highlight score $h_m \in [0, 1]^{N \times T \times P \times 1}$ for each node, *i.e.*,

$$h_m = \sigma \left(\text{STGCN} \left(\mathcal{A}_m, z_m; W_m^{(\text{hlt})} \right) \right), \quad (3)$$

where $\sigma(\cdot)$ represents the sigmoid function and $W_m^{(\text{hlt})}$ consists of the set of trainable STGCN parameters.

Our decoder takes in the feature z_m and the highlight

score h_m for each modality m , and produces a weighted latent feature $\tilde{z}_m \in \mathbb{R}^{N \times T \times P \times D_l}$ by performing a Hadamard product of h_m with each node dimension of z_m , *i.e.*,

$$\tilde{z}_m = \underbrace{[h_m; h_m; \dots]}_{D_l \text{ times}} \odot z_m. \quad (4)$$

In other words, we aim to pick the latent features in z_m that correspond to the most representative frames in \mathcal{X}_m . While training, our scorer successfully learns to assign higher h_m values to the z_m features representing the more representative frames, and favors them in the reconstruction process.

From the weighted latent feature \tilde{z}_m , our decoder produces a reconstruction $\tilde{\mathcal{X}}_m \in \mathbb{R}^{N \times T \times P \times D}$ of the input graph nodes using another STGCN, *i.e.*,

$$\tilde{\mathcal{X}}_m = \text{STGCN} \left(\mathcal{A}_m, \tilde{z}_m; W_m^{(\text{dec})} \right), \quad (5)$$

where $W_m^{(\text{dec})}$ consists of the set of trainable STGCN parameters in the decoder.

3.4. Loss Function for Training

Analogous to Eq. 1, we train our network architecture to maximally reconstruct the input graph nodes in all the modalities while minimizing the number of frames considered for reconstruction. Our approach is based on the assumption that the frames with higher representativeness, constitute the more highlightable excerpts of the video. Therefore, in effect, we aim to suppress as many frames as possible in the reconstruction of the input video while focusing on only the frames with high representativeness.

Given the highlight scores h_m for each modality m , we perform a max-pooling of the scores across all dimensions but the time to obtain $h_m^{(\max T)} \in [0, 1]^{T \times 1}$, the maximum highlight score per frame of the video for that modality, *i.e.*,

$$h_m^{(\max T)} = \max_{n \in N, p \in P} h_m. \quad (6)$$

We also consider a weighted contribution of $h_m^{(\max T)}$ for each modality m , such that the weight is proportional to the number of frames in which the modality was visible in the input video. We define a modality to be observable in a frame if more than half the constituent points of that modality are visible in the frame. By that definition, we construct a weight α_m for each modality m as

$$\alpha_m = \frac{\# \text{ frames where modality } m \text{ is observable}}{T}. \quad (7)$$

We have $0 \leq \alpha_m \leq 1 \forall m$ since each frame can contain between no and all modalities.

We then construct weighted highlight scores $\bar{h}_m \in [0, 1]^{T \times 1}$ for all the frames of the video as

$$\bar{h}_m = \alpha_m h_m^{(\max T)}. \quad (8)$$

Finally, given the decoder reconstructions $\tilde{\mathcal{X}}_m$ and the weights per modality α_m , we construct our loss function \mathcal{L}

for training our network as

$$\mathcal{L} = \sum_m \left\| \mathcal{X}_m - \tilde{\mathcal{X}}_m \right\| + \left\| \bar{h}_m \right\| + \lambda_m \left\| W_m \right\|, \quad (9)$$

where W_m collates all the trainable parameters $W_m^{(\text{enc})}$, $W_m^{(\text{hit})}$, and $W_m^{(\text{dec})}$, λ_m are the regularization factors, and we use the smooth- ℓ_1 norm for $\|\cdot\|$. We note that \mathcal{L} consists of contrasting objectives that provide the competition needed to learn the highlight scores. The subtrahend $\tilde{\mathcal{X}}_m$ in the first term, $\left\| \mathcal{X}_m - \tilde{\mathcal{X}}_m \right\|$, obtained from Eqs. 4 and 5, is a stand-in for the subtrahend in Eq. 1. Minimizing this first term would require setting all highlight scores to 1 (so all frames are highlights). Conversely, minimizing the second term $\left\| \bar{h}_m \right\|$ would require setting all highlight scores to 0 (so no frames are highlights). Consequently, our network ends up assigning high highlight scores to only the set of frames with maximal representativeness.

4. Implementation and Testing

We train our network on the large-scale AVA-Kinetics dataset [31]. This dataset consists of 235 training videos and 64 validation videos, each 15 minutes long and annotated with action labels in 1-second clips. We ignore the action labels and use the original videos to train and validate our highlight detection network. The dataset consists of a wide variety of human activities but no supervision on highlightable excerpts. Thus, it is suitable for our task of learning to detect human-specific highlight excerpts. Owing to memory constraints, we process each video in non-overlapping excerpts of 30 seconds, leading to a total of 7,050 training excerpts and 1,920 excerpts for validation.

4.1. Implementation

We use $M = 2$ modalities, poses and faces, which are the two most observable modalities in all the datasets we tested our method on. Other modalities, such as hand gestures and eye movements, are either rarely visible or suffer from noisy detection. We build the pose graph following the CMU panoptic model [21, 37], and the face landmarks graph following the face landmarks model of Geitgey [12].

We use a multi-person tracker [29] to track the persons across all the frames. We use a pose detector [37] and a face landmark detector [12], to respectively detect the coordinates of their 3D poses and 2D face landmarks. We scale all the coordinates to lie in the range $[-1, 1]$. To build our graph for each modality, we consider up to $P = 20$ persons in each frame and temporal edges to $30f$ temporally adjacent frames combining the past and the future, f being the frame rate of processing the videos. When available, we use an equal number of frames in the past and the future for temporal adjacency. We have observed efficient performance in terms of both accuracy and memory requirements for frame rates between 2 and 5, use $f = 5$ for our experiments. For all z_m 's, we use a latent dimension of $D_l = 8$.

Table 1: Mean average precision on the DSH dataset [47]. Bold: **best**, underline: second-best. Our method performs second-best in the surfing domain, where not enough poses and faces were detected, and best in all the other domains.

| Domain | RRAE [57] | Video2 GIF [17] | LSVM [47] | Less is More [53] | Ours |
|------------|-----------|-----------------|-------------|-------------------|-------------|
| dog show | 0.49 | 0.31 | <u>0.60</u> | 0.58 | 0.63 |
| gymnastics | 0.35 | 0.34 | 0.41 | <u>0.44</u> | 0.73 |
| parkour | 0.50 | 0.54 | 0.61 | <u>0.67</u> | 0.72 |
| skating | 0.25 | 0.55 | <u>0.62</u> | 0.58 | 0.64 |
| skiing | 0.22 | 0.33 | 0.36 | <u>0.49</u> | 0.52 |
| surfing | 0.49 | 0.54 | 0.61 | 0.65 | <u>0.62</u> |
| Mean | 0.38 | 0.46 | 0.54 | <u>0.57</u> | 0.64 |

Table 2: Mean average precision on PHD² [11]. Bold: **best**, underline: second-best.

| Random | FCSN [44] | Video2 GIF [17] | Ad-FCSN [42] | Ours |
|--------|-----------|-----------------|--------------|-------------|
| 0.12 | 0.15 | 0.15 | <u>0.16</u> | 0.20 |

We train using the Adam optimizer [27] for 200 epochs with a batch size of 2, an initial learning rate of 10^{-3} , a momentum of 0.9, and a weight decay of 10^{-4} . We decrease our learning rate by a factor of 0.999 after every epoch. Our training took around 4.6 GPU days at around 40 minutes per epoch on an Nvidia GeForce GTX 1080Ti GPU.

4.2. Testing

At test time, we obtain weighted highlight scores $\sum_m \bar{h}_m$ following Eq. 8 for each frame of the input video. We combine all contiguous frames above a threshold h_{thres} to generate highlight excerpts for the video. Based on our experiments, we have observed that values of $h_{\text{thres}} \geq 0.5$ leads to the detection of representative highlight excerpts in the benchmark datasets. The difference between h_{thres} and τ (Section 3.2) is that h_{thres} is used for trained highlight scores that capture domain- and user-preference-agnostic representativeness. In practice, we assign the individual highlight excerpts a score that is the mean of the weighted highlight scores for each of its constituent frames. We rank the excerpts based on these scores so that users can select their own thresholds to obtain the excerpts above those thresholds. The higher the threshold they choose, the fewer excerpts that survive the thresholding, thus reducing their manual effort of sifting through less representative excerpts.

5. Experiments

We evaluate the comparative performance of our method and current state-of-the-art methods on two large-scale public benchmark datasets: the Domain-Specific Highlights (DSH) dataset [47] and the Personal Highlight Detection dataset (PHD²) [11]. We also evaluate on the smaller public datasets of TVSum [46] and SumMe [15]. Unlike any of the current approaches, however, we do not train or fine-

Table 3: Mean average precision on the TVSum dataset [46]. Full domain names are in Section 5.1. Bold: **best**, underline: second-best. Our method performs second-best in the domains that are not fully human-centric (BK, DS, GA, MS), and best in all the other domains.

| Domain | MBF [7] | KVS [41] | CVS [40] | Adv-LSTM [36] | Less is More [53] | Ours |
|--------|---------|----------|----------|---------------|-------------------|-------------|
| BK | 0.31 | 0.34 | 0.33 | 0.42 | 0.66 | <u>0.57</u> |
| BT | 0.37 | 0.42 | 0.40 | 0.48 | <u>0.69</u> | 0.93 |
| DS | 0.36 | 0.39 | 0.38 | 0.47 | 0.63 | <u>0.60</u> |
| FM | 0.37 | 0.40 | 0.37 | <u>0.46</u> | 0.43 | 0.88 |
| GA | 0.33 | 0.40 | 0.38 | 0.48 | 0.61 | <u>0.50</u> |
| MS | 0.41 | 0.42 | 0.40 | 0.49 | 0.54 | <u>0.50</u> |
| PR | 0.33 | 0.40 | 0.38 | 0.47 | <u>0.53</u> | 0.84 |
| PK | 0.32 | 0.38 | 0.35 | 0.46 | <u>0.60</u> | 0.76 |
| VT | 0.30 | 0.35 | 0.33 | 0.42 | <u>0.56</u> | 0.65 |
| VU | 0.36 | 0.44 | 0.41 | 0.47 | <u>0.50</u> | 0.77 |
| Mean | 0.35 | 0.40 | 0.37 | 0.46 | <u>0.58</u> | 0.70 |

tune our method on any of these datasets. We also test the performance of ablated versions of our network by removing individual modalities from training and evaluation.

5.1. Datasets

The DSH dataset [47] consists of YouTube videos across six domain-specific categories: dog show, gymnastics, parkour, skating, skiing, and surfing. There are roughly 100 videos in each domain, with a total duration of around 1,430 minutes. The PHD² dataset [11] consists of a total of around 10,000 YouTube videos in the test set, totaling about 55,800 minutes. It consists of highlights annotated by 850 users based on their preferences. The TVSum dataset [46] has 50 YouTube videos totaling about 210 minutes, collected across ten domains: beekeeping (BK), bike tricks (BT), dog show (DS), flash mob (FM), grooming animal (GA), making sandwich (MS), parade (PR), parkour (PK), vehicle tire (VT), and vehicle unstuck (VU). The SumMe dataset [15] has 25 personal videos, totaling about 66 minutes.

5.2. Evaluation Metrics

We compute the commonly used mean average precision (mAP) of the detected highlights matching the annotated highlights [47, 17, 11, 53, 42]. For evaluating highlights, we consider the precision for each video individually rather than across videos, because the highlights detected from one video need not necessarily have higher highlight scores than the non-highlighted segments of another video [47]. We also report the mean F-score (harmonic mean of the precision and the recall, calculated per video, and then averaged over all videos) of our method on all the datasets and for the provided baselines on the SumMe dataset [15].

5.3. Baselines

We compare with four baselines on the DSH dataset [47], four on PHD² [11], five on the TVSum dataset [46], and

Table 4: F-scores on the SumMe dataset [15]. Bold: **best**, underline: second-best.

| Int [15] | Sub [16] | DPP-LSTM [63] | GAN-S [36] | DRL-S [67] | S ² N [52] | Ad-FCSN [42] | Ours |
|----------|----------|---------------|------------|------------|-----------------------|--------------|-------------|
| 0.39 | 0.40 | 0.39 | 0.42 | 0.42 | 0.43 | <u>0.44</u> | 0.48 |

seven on the SumMe dataset [15]. We report the performances of the baselines as stated in the literature.

On the DSH dataset, we compare with the latent SVM-based highlight ranking (LSVM) method of Sun et al. [47], Video2GIF [17], which uses C3D features with fully connected layers to learn highlight ranking, the unsupervised robust recurrent autoencoder method (RRAE) of Yang et al. [57], and the method of Xiong et al. (Less is More) [53] that learns to rank highlights by using the duration of videos as weak supervision, with the insight that shorter videos are more likely to be edited and therefore more highlightable.

On PHD², we compare with Video2GIF [17] again, the fully convolutional sequence network (FCSN) that uses GoogLeNet to learn image-based features for highlight detection [44], and the adaptive FCSN method (Ad-FCSN) [42], which additionally consists of a history encoder to adapt to a user’s history of highlight preferences to detect personalized highlights. We also use a fully random highlight detector as the lowest baseline following [42].

On the TVSum dataset, we compare again with the duration-based highlight detection method (Less is More) [53], the visual correlation-based method of Chu et al. [7] that uses maximal biclique finding (MBF) to obtain co-occurring shots that are also relevant to the original video, the kernel-based video summarization method (KVS) of Potapov et al. [41] that trains an SVM on semantically consistent segments, the collaborative video summarization method (CVS) of Panda et al. [40] that uses a consensus regularizer to detect highlight segments satisfying sparsity, diversity, and representativeness, and the unsupervised video summarization method of Mahasseni et al. [36] using LSTMs with adversarial loss (Adv-LSTM).

On the SumMe dataset, we compare again with adaptive FCSN (Ad-FCSN) [42], the interestingness-based summarization method (Int.) of Gygli et al. [15], the submodularity-based summarization method (Sub.) of Gygli et al. [16], the LSTM network of Zhang et al. [63] employing a determinantal point process (DPP-LSTM), the GAN-based method of Lu and Grauman [35] with extra supervision (GAN-S), the deep reinforcement learning-based method of Zhou et al. [67] with extra supervision (DRL-S), and the sequence to segments detection method (S²N) [52] that uses an encoder-decoder architecture to detect segments with high relevance from sequence data.

5.4. Results

DSH [47] and TVSum [46]. We report the mAP across all the domains in these datasets in Tables 1 and 3 respectively. We outperform the baselines on all but a few do-

Table 5: Comparison of mean mAP and mean F-score for different ablated versions of our method on benchmark datasets. Bold: **best**, underline: second-best.

| Dataset | Using Modality | | | | | |
|-----------------------|----------------|-------------|-------------|-------------|-------------|-------------|
| | Face only | | Pose only | | Both | |
| | mAP | F | mAP | F | mAP | F |
| DSH [47] | 0.51 | 0.45 | <u>0.57</u> | <u>0.48</u> | 0.64 | 0.56 |
| TVSum [46] | 0.57 | 0.46 | <u>0.64</u> | <u>0.56</u> | 0.70 | 0.59 |
| PHD ² [11] | <u>0.16</u> | <u>0.20</u> | 0.15 | 0.18 | 0.20 | 0.22 |
| SumMe [15] | <u>0.48</u> | 0.39 | 0.45 | <u>0.41</u> | 0.52 | 0.48 |

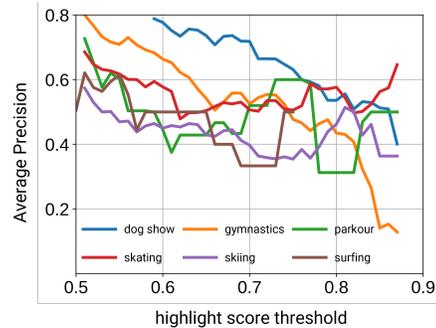


Figure 4: **Average precision by highlight score threshold h_{thres} .** On the domains in the DSH dataset [47].

main, which are either not fully human centric (beekeeping, dog show, grooming animals, and making sandwich in TVSum), or where sufficient poses and faces could not be detected (surfing in DSH). However, we come second-best on these domains, and on average, across the domains, we outperform the best baselines by an absolute 4% – 12%.

PHD² [11]. We report the mAP across the dataset in Table 2. Given the abundance of humans detected in the videos, we outperform the best baseline by an absolute 4%.

SumMe [15]. We report the mean F-scores across the dataset in Table 4. Following prior methods [52, 42], we randomly select 20% of the dataset for calculating the mean F-score, repeat this experiment five times, and report the mean performance. Based on these experiments, we outperform the best baseline by an absolute 4%.

These results demonstrate that our approach of using human-centric modalities to detect highlights leads to state-of-the-art performance on all these benchmark datasets.

5.5. Ablation Studies

In our experiments, we consider two modalities, poses and faces. We ablate each of these two modalities in turn and test the performance of our method by training it on the remaining modality. We report the mean mAP and the mean F-score of the ablated versions of our method on all four benchmark datasets in Table 5. Using only poses and no faces, we observe an absolute drop-off of 5% – 7% for the mean mAP and 3% – 8% for the mean F-score across the datasets, compared to using both modalities. Using only



Figure 5: **Sample highlight frames detected by our method.** We show sample frames across the range of highlight scores as detected by different ablated versions of our method. We show one sample video from the datasets SumMe [15], PHD² [11], DSH [47], and TVSum [46], in order from top to bottom. When using only faces or only poses, our method learns highlight scores based only on face- or pose-based representativeness. Combining both the modalities, our method learns highlight scores based on representativeness from both.

faces and no poses, we observe more severe absolute drop-offs of 4% – 13% for the mean mAP and 2% – 13% for the mean F-score across the datasets. This happens because poses are generally more abundant and more easily detected compared to face landmarks. For example, poses can be detected even when a human is partially occluded, in the dark, or not in clear focus, whereas detection of face landmarks requires the face to be well-lit and in focus. Therefore, not detecting poses leads to missing a significant number of highlightable excerpts. This trend is reversed only in PHD², where faces were more commonly detected than poses.

We also show the qualitative performance of our method

and all its ablated versions on one sample video from each of the four datasets, DSH, PHD², TVSum, and SumMe, in Figure 5. We see that when observing only poses and not faces, our method detects the representative highlight excerpts with pose-based expressions but fails to detect excerpts that primarily have facial expressions and emotions. Conversely, when observing only faces and not poses, our method can only detect the excerpts where the faces are prominent, and misses excerpts where the faces are too small, occluded, or in the dark. Using both modalities, our method can detect all the representative excerpts.

5.6. Effect of Highlight Score Threshold

We use a threshold h_{thres} on the highlight score predicted by our method to detect the highlightable excerpts (Section 4.2). To visualize the effect h_{thres} has on the average precision (AP), we show the plot of AP vs. h_{thres} on each domain in the DSH dataset [47] in Figure 4. We observe a general trend of the AP decreasing as we increase the threshold, as our method returns fewer and fewer highlights. However, this is not true for some domains like surfing, where the highlight scores of the representative excerpts are already sufficiently high. In practice, we consider the choice of threshold to be user-configurable for each video.

6. Conclusion, Limitations and Future Work

We have presented a novel method for detecting highlights from human-centric videos, which leverages the observable human-centric modalities, such as faces and poses, and uses these modalities to automatically detect the most representative highlights from the video. Extensive experimental results on the domain-specific highlights (DSH) dataset [47], the personal highlight detection dataset (PHD²) [11], the TVSum dataset [46], and the SumMe dataset [15] demonstrate the benefits of our proposed approach compared to several state-of-the-art baselines.

Our work has some limitations. Although our network design can accommodate any number of modalities, we have used only faces and poses as the modalities in our benchmarks. This leads to state-of-the-art results on average. However, many videos (*e.g.*, videos on grooming animals, making sandwich in TVSum) exhibit other modalities such as hands and fingers. Thus, we plan to incorporate more human-centric modalities into our experiments in the future. Our method may not offer much improvements in videos from domains that are not fully human-centric, *e.g.*, videos focusing on other living beings, inanimate objects, and natural scenes. We plan to explore these domains in the future using appropriate modalities. Our method can also be combined with domain-specific features and adapted to user preferences to detect more fine-tuned highlights.

Acknowledgment

This work was supported in part by ARO Grants W911NF1910069 and W911NF2110026.

References

- [1] Hillel Aviezer, Yaacov Trope, and Alexander Todorov. Body cues, not facial expressions, discriminate between intense positive and negative emotions. *Science*, 338(6111):1225–1229, 2012. [2](#)
- [2] Uttaran Bhattacharya, Trisha Mittal, Rohan Chandra, Tanmay Randhavane, Aniket Bera, and Dinesh Manocha. Step: Spatial temporal graph convolutional networks for emotion perception from gaits. In *Proceedings of the Thirty-Fourth AAAI Conference on Artificial Intelligence (AAAI-20)*, page 1342–1350. AAAI Press, 2020. [2](#), [3](#)
- [3] Carlos Busso, Murtaza Bulut, Chi-Chun Lee, Abe Kazemzadeh, Emily Mower, Samuel Kim, Jeannette N Chang, Sungbok Lee, and Shrikanth S Narayanan. Iemocap: Interactive emotional dyadic motion capture database. *Language resources and evaluation*, 42(4):335, 2008. [3](#)
- [4] Sijia Cai, Wangmeng Zuo, Larry S. Davis, and Lei Zhang. Weakly-supervised video summarization using variational encoder-decoder and web prior. In *Proceedings of the European Conference on Computer Vision (ECCV)*, September 2018. [2](#)
- [5] C. Chen, R. Jafari, and N. Kehtarnavaz. Utd-mhad: A multimodal dataset for human action recognition utilizing a depth camera and a wearable inertial sensor. In *2015 IEEE International Conference on Image Processing (ICIP)*, pages 168–172, 2015. [3](#)
- [6] Chong-Wah Ngo, Yu-Fei Ma, and Hong-Jiang Zhang. Automatic video summarization by graph modeling. In *Proceedings Ninth IEEE International Conference on Computer Vision*, pages 104–109 vol.1, 2003. [2](#)
- [7] Wen-Sheng Chu, Yale Song, and Alejandro Jaimes. Video co-summarization: Video summarization by visual co-occurrence. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2015. [2](#), [6](#), [7](#)
- [8] Cisco. Annual internet report (2018–2023): https://www.cisco.com/c/en/us/solutions/collateral/service-provider/visual-networking-index-vni/complete-white-paper-c11-481360.html#_toc484813989. *CISCO White paper*, 2020. [1](#)
- [9] Michaël Defferrard, Xavier Bresson, and Pierre Vandergheynst. Convolutional neural networks on graphs with fast localized spectral filtering. In D. D. Lee, M. Sugiyama, U. V. Luxburg, I. Guyon, and R. Garnett, editors, *Advances in Neural Information Processing Systems 29*, pages 3844–3852. Curran Associates, Inc., 2016. [4](#)
- [10] Annalisa Franco, Antonio Magnani, and Dario Maio. A multimodal approach for human activity recognition based on skeleton and rgb data. *Pattern Recognition Letters*, 131:293–299, 2020. [3](#)
- [11] Ana Garcia del Molino and Michael Gygli. Phd-gifs: Personalized highlight detection for automatic gif creation. In *Proceedings of the 26th ACM International Conference on Multimedia*, MM ’18, page 600–608, New York, NY, USA, 2018. Association for Computing Machinery. [1](#), [2](#), [6](#), [7](#), [8](#)
- [12] Adam Gettgey. *Face Recognition*, 2020. [3](#), [4](#), [5](#)
- [13] Boqing Gong, Wei-Lun Chao, Kristen Grauman, and Fei Sha. Diverse sequential subset selection for supervised video summarization. In Z. Ghahramani, M. Welling, C. Cortes, N. Lawrence, and K. Q. Weinberger, editors, *Advances in Neural Information Processing Systems*, volume 27, pages 2069–2077. Curran Associates, Inc., 2014. [3](#)
- [14] Hongxiang Gu and Viswanathan Swaminathan. From thumbnails to summaries—a single deep neural network to rule them all. In *2018 IEEE International Conference on Multimedia and Expo (ICME)*, pages 1–6. IEEE, 2018. [2](#), [4](#)
- [15] Michael Gygli, Helmut Grabner, Hayko Riemenschneider, and Luc Van Gool. Creating summaries from user videos. In David Fleet, Tomas Pajdla, Bernt Schiele, and Tinne Tuytelaars, editors, *Computer Vision – ECCV 2014*, pages 505–520, Cham, 2014. Springer International Publishing. [2](#), [3](#), [6](#), [7](#), [8](#)
- [16] Michael Gygli, Helmut Grabner, and Luc Van Gool. Video summarization by learning submodular mixtures of objectives. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2015. [3](#), [7](#)
- [17] Michael Gygli, Yale Song, and Liangliang Cao. Video2gif: Automatic generation of animated gifs from video. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2016. [2](#), [6](#), [7](#)
- [18] Yifan Jiao, Zhetao Li, Shucheng Huang, Xiaoshan Yang, Bin Liu, and Tianzhu Zhang. Three-dimensional attention-based deep ranking model for video highlight detection. *IEEE Transactions on Multimedia*, 20(10):2693–2705, 2018. [2](#)
- [19] Yifan Jiao, Tianzhu Zhang, Shucheng Huang, Bin Liu, and Changsheng Xu. Video highlight detection via region-based deep ranking model. *International Journal of Pattern Recognition and Artificial Intelligence*, 33(07):1940001, 2019. [2](#)
- [20] Hideo Joho, Jacopo Staiano, Nicu Sebe, and Joemon M Jose. Looking at the viewer: analysing facial activity to detect personal highlights of multimedia contents. *Multimedia Tools and Applications*, 51(2):505–523, 2011. [2](#)
- [21] Hanbyul Joo, Tomas Simon, Xulong Li, Hao Liu, Lei Tan, Lin Gui, Sean Banerjee, Timothy Scott Godisart, Bart Nabbe, Iain Matthews, Takeo Kanade, Shohei Nobuhara, and Yaser Sheikh. Panoptic studio: A massively multiview system for social interaction capture. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2017. [5](#)
- [22] Aditya Khosla, Raffay Hamid, Chih-Jen Lin, and Neel Sundaresan. Large-scale video summarization using web-image priors. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2013. [2](#)
- [23] Gunhee Kim, Leonid Sigal, and Eric 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 (CVPR)*, June 2014. [1](#), [2](#)
- [24] Gunhee Kim and Eric P. Xing. Reconstructing storyline graphs for image recommendation from web community photos. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2014. [2](#)
- [25] Hoseong Kim, Tao Mei, Hyeran Byun, and Ting Yao. Exploiting web images for video highlight detection with triplet deep ranking. *IEEE Transactions on Multimedia*, 20(9):2415–2426, 2018. [1](#), [2](#)
- [26] Yelin Kim, Honglak Lee, and Emily Mower Provost. Deep learning for robust feature generation in audiovisual emotion recognition. *ICASSP*, pages 3687–3691, 2013. [3](#)
- [27] Diederik P Kingma and Jimmy Ba. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*, 2014. [6](#)

- [28] Thomas N Kipf and Max Welling. Semi-supervised classification with graph convolutional networks. *arXiv preprint arXiv:1609.02907*, 2016. 4
- [29] Muhammed Kocabas. *Simple Multi Person Tracker*, 2019. 3, 4, 5
- [30] 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, 2012. 2, 3
- [31] Ang Li, Meghana Thotakuri, David A Ross, João Carreira, Alexander Vostrikov, and Andrew Zisserman. The ava-kinetics localized human actions video dataset. *arXiv preprint arXiv:2005.00214*, 2020. 5
- [32] G. Li, Y. Zhao, M. Ji, X. Yuan, and L. Fang. Zoom in to the details of human-centric videos. In *2020 IEEE International Conference on Image Processing (ICIP)*, pages 3089–3093, 2020. 2
- [33] Jianan Li, Xuemei Xie, Qingzhe Pan, Yuhan Cao, Zhifu Zhao, and Guangming Shi. Sgm-net: Skeleton-guided multimodal network for action recognition. *Pattern Recognition*, 104:107356, 2020. 3
- [34] S. Li and W. Deng. Deep facial expression recognition: A survey. *IEEE Transactions on Affective Computing*, pages 1–1, 2020. 2
- [35] Zheng Lu and Kristen Grauman. Story-driven summarization for egocentric video. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2013. 2, 7
- [36] Behrooz Mahasseni, Michael Lam, and Sinisa Todorovic. Unsupervised video summarization with adversarial lstm networks. In *Proceedings of the IEEE conference on Computer Vision and Pattern Recognition*, pages 202–211, 2017. 1, 2, 4, 6, 7
- [37] D. Mehta, O. Sotnychenko, F. Mueller, W. Xu, S. Sridhar, G. Pons-Moll, and C. Theobalt. Single-shot multi-person 3d pose estimation from monocular rgb. In *2018 International Conference on 3D Vision (3DV)*, pages 120–130, 2018. 5
- [38] Trisha Mittal, Uttaran Bhattacharya, Rohan Chandra, Aniket Bera, and Dinesh Manocha. M3er: Multiplicative multimodal emotion recognition using facial, textual, and speech cues. In *Proceedings of the Thirty-Fourth AAAI Conference on Artificial Intelligence, AAAI’20*, pages 1359–1367. AAAI Press, 2020. 2, 3
- [39] Trisha Mittal, Pooja Guhan, Uttaran Bhattacharya, Rohan Chandra, Aniket Bera, and Dinesh Manocha. Emoticon: Context-aware multimodal emotion recognition using frege’s principle. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2020. 3
- [40] Rameswar Panda and Amit K. Roy-Chowdhury. Collaborative summarization of topic-related videos. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, July 2017. 2, 6, 7
- [41] Danila Potapov, Matthijs Douze, Zaid Harchaoui, and Cordelia Schmid. Category-specific video summarization. In David Fleet, Tomas Pajdla, Bernt Schiele, and Tinne Tuytelaars, editors, *Computer Vision – ECCV 2014*, pages 540–555, Cham, 2014. Springer International Publishing. 2, 6, 7
- [42] Mrigank Rochan, Mahesh Kumar Krishna Reddy, Linwei Ye, and Yang Wang. Adaptive video highlight detection by learning from user history. In *Proceedings of the European Conference on Computer Vision (ECCV)*, August 2020. 1, 2, 4, 6, 7
- [43] Mrigank Rochan and Yang Wang. Video summarization by learning from unpaired data. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2019. 2
- [44] Mrigank Rochan, Linwei Ye, and Yang Wang. Video summarization using fully convolutional sequence networks. In *Proceedings of the European Conference on Computer Vision (ECCV)*, September 2018. 1, 2, 6, 7
- [45] A. Shahroudy, T. Ng, Y. Gong, and G. Wang. Deep multimodal feature analysis for action recognition in rgb+d videos. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 40(5):1045–1058, 2018. 3
- [46] Yale Song, Jordi Vallmitjana, Amanda Stent, and Alejandro Jaimes. Tvsum: Summarizing web videos using titles. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2015. 2, 6, 7, 8
- [47] Min Sun, Ali Farhadi, and Steve Seitz. Ranking domain-specific highlights by analyzing edited videos. In *European conference on computer vision*, pages 787–802. Springer, 2014. 1, 2, 4, 6, 7, 8
- [48] Zongheng Tang, Yue Liao, Si Liu, Guanbin Li, Xiaojie Jin, Hongxu Jiang, Qian Yu, and Dong Xu. Human-centric spatio-temporal video grounding with visual transformers. *arXiv preprint arXiv:2011.05049*, 2020. 2
- [49] Ba Tu Truong and Svetha Venkatesh. Video abstraction: A systematic review and classification. *ACM Trans. Multimedia Comput. Commun. Appl.*, 3(1):3–es, Feb. 2007. 2
- [50] Paul Vicol, Makarand Tapaswi, Lluís Castrejón, and Sanja Fidler. Moviegraphs: Towards understanding human-centric situations from videos. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2018. 1, 2
- [51] Xueyang Wang, Xiya Zhang, Yinheng Zhu, Yuchen Guo, Xiaoyun Yuan, Liuyu Xiang, Zerun Wang, Guiguang Ding, David Brady, Qionghai Dai, and Lu Fang. Panda: A gigapixel-level human-centric video dataset. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2020. 2
- [52] Zijun Wei, Boyu Wang, Minh Hoai Nguyen, Jianming Zhang, Zhe Lin, Xiaohui Shen, Radomir Mech, and Dimitris Samaras. Sequence-to-segment networks for segment detection. In S. Bengio, H. Wallach, H. Larochelle, K. Grauman, N. Cesa-Bianchi, and R. Garnett, editors, *Advances in Neural Information Processing Systems*, volume 31, pages 3507–3516. Curran Associates, Inc., 2018. 2, 7
- [53] Bo Xiong, Yannis Kalantidis, Deepti Ghadiyaram, and Kristen Grauman. Less is more: Learning highlight detection from video duration. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 1258–1267, 2019. 1, 2, 4, 6, 7
- [54] Bo Xiong, Gunhee Kim, and Leonid Sigal. Storyline representation of egocentric videos with an applications to story-based search. In *Proceedings of the IEEE International Conference on Computer Vision (ICCV)*, December 2015. 2
- [55] Jia Xu, Lopamudra Mukherjee, Yin Li, Jamieson Warner, James M. Rehg, and Vikas Singh. Gaze-enabled egocentric

- video summarization via constrained submodular maximization. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2015. 3
- [56] Sijie Yan, Yuanjun Xiong, and Dahua Lin. Spatial temporal graph convolutional networks for skeleton-based action recognition. In Sheila A. McIlraith and Kilian Q. Weinberger, editors, *Proceedings of the Thirty-Second AAAI Conference on Artificial Intelligence, (AAAI-18)*, pages 7444–7452. AAAI Press, 2018. 2, 3, 4
- [57] Huan Yang, Baoyuan Wang, Stephen Lin, David Wipf, Minyi Guo, and Baining Guo. Unsupervised extraction of video highlights via robust recurrent auto-encoders. In *Proceedings of the IEEE International Conference on Computer Vision (ICCV)*, December 2015. 2, 6, 7
- [58] Ting Yao, Tao Mei, and Yong 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. 2
- [59] Minerva Yeung, Boon-Lock Yeo, and Bede Liu. Segmentation of video by clustering and graph analysis. *Computer Vision and Image Understanding*, 71(1):94 – 109, 1998. 2
- [60] Youngjae Yu, Sangho Lee, Joonil Na, Jaeyun Kang, and Gunhee Kim. A deep ranking model for spatio-temporal highlight detection from a 360 video. pages 7525–7533, 2018. 2
- [61] AmirAli Bagher Zadeh, Paul Pu Liang, Soujanya Poria, Erik Cambria, and Louis-Philippe Morency. Multimodal language analysis in the wild: Cmu-mosei dataset and interpretable dynamic fusion graph. In *ACL (Volume 1: Long Papers)*, pages 2236–2246, 2018. 3
- [62] Wenjun Zeng. Toward human-centric deep video understanding. *APSIPA Transactions on Signal and Information Processing*, 9:e1, 2020. 1
- [63] Ke Zhang, Wei-Lun Chao, Fei Sha, and Kristen Grauman. Video summarization with long short-term memory. In Bastian Leibe, Jiri Matas, Nicu Sebe, and Max Welling, editors, *Computer Vision – ECCV 2016*, pages 766–782, Cham, 2016. Springer International Publishing. 1, 3, 7
- [64] Ke Zhang, Kristen Grauman, and Fei Sha. Retrospective encoders for video summarization. In *Proceedings of the European Conference on Computer Vision (ECCV)*, September 2018. 1, 2, 3
- [65] Bin Zhao, Xuelong Li, and Xiaoqiang Lu. Hierarchical recurrent neural network for video summarization. In *Proceedings of the 25th ACM International Conference on Multimedia, MM '17*, page 863–871, New York, NY, USA, 2017. Association for Computing Machinery. 3
- [66] Bin Zhao and Eric P. Xing. Quasi real-time summarization for consumer videos. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2014. 2
- [67] Kaiyang Zhou, Yu Qiao, and Tao Xiang. Deep reinforcement learning for unsupervised video summarization with diversity-representativeness reward. pages 7582–7589, 2018. 1, 2, 7