

Sentence Guided Temporal Modulation for Dynamic Video Thumbnail Generation

Mrigank Rochan¹
mrochan@cs.umanitoba.ca

Mahesh Kumar Krishna Reddy¹
kumarm@cs.umanitoba.ca

Yang Wang^{1,2}
ywang@cs.umanitoba.ca

¹ University of Manitoba
Winnipeg, Canada

² Huawei Technologies
Canada

Abstract

We consider the problem of sentence specified dynamic video thumbnail generation. Given an input video and a user query sentence, the goal is to generate a video thumbnail that not only provides the preview of the video content, but also semantically corresponds to the sentence. In this paper, we propose a sentence guided temporal modulation (SGTM) mechanism that utilizes the sentence embedding to modulate the normalized temporal activations of the video thumbnail generation network. Unlike the existing state-of-the-art method that uses recurrent architectures, we propose a non-recurrent framework that is simple and allows much more parallelization. Extensive experiments and analysis on a large-scale dataset demonstrate the effectiveness of our framework.

1 Introduction

With the massive rise in videos available online, video thumbnails play a vital role in defining the browsing and searching experience of users. A video thumbnail shows viewers a quick and condensed preview of the entire content in the video [19, 28, 67]. Viewers often decide whether to watch or skip the video based on its thumbnail [8, 28]. Given its importance, there is increasing interest in how to create attractive and expressive thumbnails.

Most of the existing approaches [6, 11, 13, 20, 28, 65] focus on static thumbnail generation, where a thumbnail is generated solely based on the input video. These static thumbnail generation methods overlook rich semantic information such as user query sentence which is usually provided in searching a video. Static thumbnails are not tailored to each viewer's unique interest and may not provide the best online videos browsing experience. Some approaches [17, 19] consider the user search query for video thumbnail generation. But they either limit the thumbnail to a single keyframe or confine queries to a single word or a short phrase [67]. In this paper, we study the recently proposed challenging problem of sentence specified dynamic video thumbnail generation (DVTG) [67]. Given an input video and a user query expressed as a free form natural language sentence, the goal of DVTG is to generate a video thumbnail that semantically corresponds to the sentence while giving a concise preview of the video. Figure 1 shows the difference between static thumbnail generation and

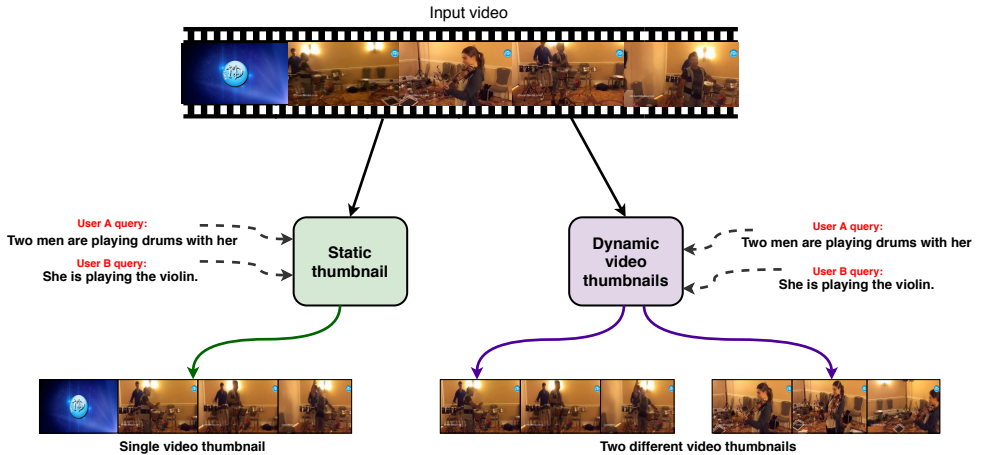


Figure 1: Illustration of the difference between static video thumbnails and sentence specified dynamic video thumbnails. The latter considers the user query sentence when creating the thumbnail.

DVTG. Despite many potential real world applications (e.g., video search), there is limited work on the DVTG problem.

The existing state-of-the-art method [57] for the DVTG task achieves promising results but has some limitations. First, it heavily relies on fine-grained modeling of semantic relationship within the user query sentence and clips or segments of the video [57]. It does not consider the overall guiding and modulating role that a user query sentence can play in temporally correlating video clips over time. Intuitively, global sentence semantics can serve as a reference in determining and associating sentence-specific video segments over time. Second, the method uses recurrent models (BiGRU [4] and pointer networks [63]) that are hard to parallelize as they inherently perform sequential computation.

In this paper, we propose a sentence guided temporal modulation (SGTM) mechanism for the DVTG task. We use a self-attention network [54] for encoding the user query sentence and adapt a fully convolutional temporal network [76] for the thumbnail generation of an input video. The SGTM mechanism leverages the semantic information from the user query sentence to modulate the normalized temporal activations in the video thumbnail generation network. We also introduce a small auxiliary network to further ensure that the generated video thumbnail semantically corresponds to the given user sentence. Our framework is computationally efficient as its computations are easily parallelizable on GPU architectures. It is also noteworthy that our framework is free from sophisticated multimodal feature fusion (unlike prior method [57]), making it computationally less complex. We also propose an unsupervised extension of our method that does not need ground-truth video thumbnails during training.

In summary, our contributions in this paper are as follows. (1) We propose a sentence guided temporal modulation (SGTM) mechanism for sentence specified dynamic video thumbnail generation (DVTG). Our method dynamically modulates the temporal activations in video thumbnail generation network using the semantic information of the query sentence. (2) We propose a computationally efficient and simple framework for the DVTG task that offers better parallelization and is free from complex multimodal feature fusion. (3) We pro-

pose both supervised and unsupervised version of our method. (4) We conduct extensive experiments and analysis on a large-scale dataset to evaluate our framework.

2 Related Work

Video thumbnail selection: Traditional methods [6, 9, 10, 13, 21, 28] of video thumbnail generation operate entirely on visual features and characteristics. These methods do not consider any other information related to the video such as textual queries from the user when generating the thumbnail. Some recent methods on automatic thumbnail selection [12, 19, 37] propose to leverage user textual queries to generate video thumbnails. Most of these methods, however, do not handle complex user queries and are inspired from multi-modal semantic matching model [1, 22] that are popular for image search and tagging. The work by Yuan *et al.* [37] is the most relevant one to ours. This method combines a variant of pointer network [33], graph convolution network [15] and BiGRU [1] to address thumbnail selection based on the user query. This method is computationally complex and prevents parallelization due to sequential computation within. In this paper, we propose an approach that is simple and much more parallelizable.

Video highlight and summarization: Thumbnail generation is related to video highlight detection and video summarization. The goal of video highlight detection is to find the most interesting events or segments in the video [5]. Video summarization aims to produce a short and compact overview of the video [5]. Although these two tasks are different from thumbnail generation, there is some overlap in techniques employed [28]. Vasudevan *et al.* [30] propose a query-adaptive video summarization technique. While promising, this method requires another image annotation dataset to learn the model. Gygli *et al.* [5] use user-generated GIFs to learn video highlight detection. This method can be used to produce thumbnails but they ignore the user query sentence information.

Temporal sentence localization in videos: Temporal sentence localization in video aims to determine the starting and ending of a continuous video segment that matches with the given natural language sentence [10, 8, 18, 36]. Different from this task, video thumbnail generation may contain several nonconsecutive video segments [37]. Additionally, temporal sentence localization mainly focuses on matching a sentence to a video segment. In contrast, a dynamic video thumbnail should also present a quick preview of the video content along with establishing the relationship with the user query sentence [37].

Conditional normalization methods: Our work is also related to conditional batch normalization [9] and adaptive instance normalization [12] methods. These methods spatially normalize the activation from a layer to zero mean and unit variance, and then apply an affine transformation whose parameters are computed using external data [12]. These methods have been successfully applied in image understanding tasks such as visual question answering and image-to-image translation.

3 Our Approach

The input to DVTG consists of a video V and a user query sentence S . We denote video as $V = \{v_c\}_{c=1}^C$, where v_c is the feature representation of c -th video clip and C is the total number of clips in the video. Similarly, we represent the user query sentence as $S = \{x_n\}_{n=1}^N$, where x_n denotes the word embedding of n -th word and N denotes the total number of words

in the sentence. The goal of DVTG is to identify a set of video clips (which may not be consecutive) from V that provide a good preview of the original video V and are semantically consistent with the sentence S .

Given V and S , our goal is to learn a mapping function $F(V, S) \in \mathbb{R}^{1 \times C \times 2}$, where the output of $F(V, S)$ indicate the scores of whether or not a video clip should be included in the video thumbnail. We call the function F as the sentence guided video thumbnail generation model.

3.1 Sentence Guided Video Thumbnail Generation Model

We propose the sentence guided video thumbnail generation model F which consists of three sub-networks, namely a video thumbnail generation network (T), a self-attention sentence encoder network (S_{enc}) and an auxiliary network (T_{aux}). A sentence guided temporal modulation (SGTM) mechanism is proposed to modulate certain activations in T using the output from S_{enc} . Figure 2 shows an overview of our proposed model. In the following, we discuss our model in detail.

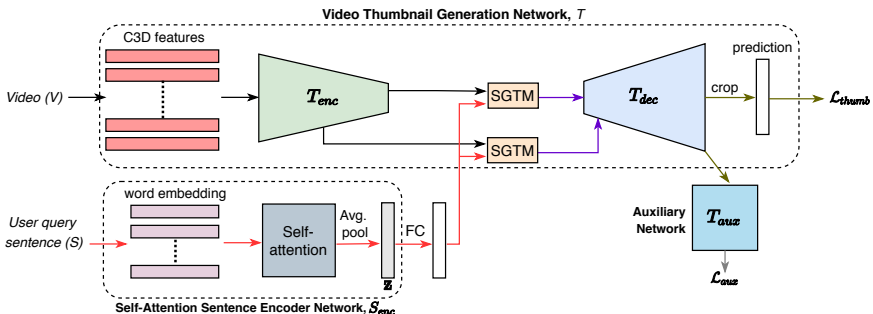


Figure 2: An overview of our sentence guided video thumbnail generation model. The model consists of a video thumbnail generation network (T), a self-attention sentence encoder network (S_{enc}), and an auxiliary network (T_{aux}). A sentence guided temporal modulation (SGTM) mechanism is introduced to allow interaction between T and S_{enc} . The network T consists of an encoder T_{enc} and a decoder T_{dec} . T takes features of video clips in an input video V and predicts whether or not each clip belongs to the video thumbnail. S_{enc} encodes the word-level embedding of user query sentence (S) to a vector \mathbf{z} which is then used by SGTM to modulate the temporal activations from the encoder of T (i.e., T_{enc}) to determine sentence-specific video content over time. The role of T_{aux} is to reconstruct \mathbf{z} to further ensure that the generated video thumbnail aligns well with S . We use two losses for learning: a thumbnail generation loss \mathcal{L}_{thumb} on the prediction of T_{dec} and an auxiliary loss \mathcal{L}_{aux} on the output of T_{aux} .

3.1.1 Video Thumbnail Generation Network

The video thumbnail generation network T is an encode-decoder style temporal convolution network. A temporal convolution network mainly performs 1D operations (e.g., 1D convolution, 1D pooling) over time. For instance, given a video with frame-level feature representations, this network can operate over frames enabling it to capture the temporal

dependencies among the frames. In this paper, we adapt FCSN [26] which is a form of encoder-decoder based temporal convolution network designed for video summarization.

The input to T is a sequence of clip-level feature representations of video V , i.e., $V \in \mathbb{R}^{1 \times C \times D_c}$ where C is the total number of clips and D_c is the dimension of the feature representation of each video clip. The output of T is of dimension $1 \times C \times 2$ which denotes the scores of each clip being a part of the thumbnail or not.

T consists of an encoder T_{enc} and a decoder T_{dec} . T_{enc} has seven convolutional blocks. Each of the first five convolution blocks ($conv1$ to $conv5$) consists of multiple temporal convolution and ReLU operations, with a max pooling at the end. The last two blocks ($fc6$ and $fc7$) has a temporal convolution, a ReLU and a dropout operation. T_{enc} produces two feature maps as outputs, one from the last layer and another one as a skip connection from $conv4$ block. The outputs of T_{enc} is fed to T_{dec} as inputs. The first input to T_{dec} goes through a 1×1 convolution and a temporally fractionally-strided convolution ($deconv1$) which is then combined with the second input after applying a 1×1 convolution to it. Lastly, it has another temporally fractionally-strided convolution ($deconv2$) and a crop operation so as to obtain the final prediction of dimension $1 \times C \times 2$.

3.1.2 Self-Attention Sentence Encoder Network

The self-attention sentence encoder network (S_{enc}) is responsible for encoding the user query sentence S to a fixed length vector \mathbf{z} .

Self-attention [23, 32] has been shown to be a powerful technique in natural language processing. We apply the non-local model [32, 33] with self-attention mechanism in S_{enc} . This enables S_{enc} to effectively model relationships between different words in the sentence.

We represent each word in the sentence using its word embedding vector. The sentence S can be written as $S \in \mathbb{R}^{D_w \times N}$, where N is the number of words in the sentence and D_w is the dimension of word embedding. The attention between two words can be computed as:

$$\Phi_{j,i} = \frac{\exp(A_{ij})}{\sum_{i=1}^N \exp(A_{ij})}, \text{ where } A_{ij} = f_1(x_i)^T f_2(x_j), \quad (1)$$

Here f_1 and f_2 represent two distinct feature spaces, i.e., $f_1(x_i) = W_{f_1} x_i$ and $f_2(x_j) = W_{f_2} x_j$. The attention score $\Phi_{j,i}$ denotes the extent with which i -th word is related to j -th word. The output of self-attention is $\Lambda = (\Lambda_1, \Lambda_2, \dots, \Lambda_N) \in \mathbb{R}^{D_w \times N}$, where

$$\Lambda_j = \sum_{i=1}^N \Phi_{j,i} h(x_i), \text{ where } h(x_i) = W_h x_i. \quad (2)$$

Note that $W_{f_1} \in \mathbb{R}^{d \times D_w}$, $W_{f_2} \in \mathbb{R}^{d \times D_w}$ and $W_h \in \mathbb{R}^{D_w \times D_w}$ are the learnable weights implemented using 1×1 convolutions. In our experiments, we set $d = D_w/8$. Lastly, we apply a average pooling on the resultant self-attended features Λ to obtain the vector $\mathbf{z} \in \mathbb{R}^{D_w}$ representing the whole sentence S .

3.1.3 Sentence Guided Temporal Modulation

In order to generate a video thumbnail that corresponds to a user sentence, it is necessary to establish the relationships between the clips of video and the sentence. The rich semantics in the sentence provides a strong signal to temporally associate it with the video clips. Motivated by this, we propose the sentence guided temporal modulation (SGTM) mechanism.

In practice, SGTm is similar to temporal-adaptive instance normalization [24] that extends adaptive instance normalization [14] (initially designed for images) to videos. SGTm plays a key role in our proposed model. It allows the interaction between the video thumbnail generation network T (Sec. 3.1.1) and the self-attention sentence encoder network S_{enc} (Sec. 3.1.2) so as to generate a video thumbnail that closely relates to the given user query sentence. SGTm uses the semantic information in the sentence to guide and temporally modulate the features of the input video fed to the video thumbnail generation network T .

Let \mathbf{z} be the sentence vector representation from S_{enc} and $A \in \mathbb{R}^{1 \times M \times C}$ be the activations from one specific layer in the temporal convolution network T , where M is the temporal length of the activation and C is the number of channels. We firstly forward \mathbf{z} to a fully-connected layer (FC) to generate a vector of length $2C$. We use this generated vector to obtain two modulation vectors $\alpha \in \mathbb{R}^C$ and $\beta \in \mathbb{R}^C$. Next, we use these modulation vectors to modulate the each channel of activation map A . We can express the modulated activation map at $m \in M$ and $c \in C$ as $\hat{A}_{m,c}$ which is computed as:

$$\hat{A}_{m,c} = \alpha_c \cdot \frac{A_{m,c} - \mu_c(A)}{\sigma_c(A)} + \beta_c, \quad (3)$$

where μ_c and σ_c are the channel-wise mean and standard deviation calculated independently for each input video along the temporal length of the activation map A .

From Eq. 3, we can infer that a normalized temporal activation map with zero mean and unit variance in each channel is subjected to an affine transformation (scale and shift) whose values are predicted using the sentence representation \mathbf{z} . Note that the affine transformation is temporally invariant and there are no learnable parameters involved in this mechanism. By using this translation-based design, our goal is to allow the sentence semantics information to modulate each temporal feature map of the input video.

In our model, we apply SGTm to the two output activations of the encoder T_{enc} (see Fig. 2). We use the sentence representation \mathbf{z} from S_{enc} to produce two sets of affine parameter vectors $(\alpha_i, \beta_i$ where $i = 1, 2$) using a FC layer. The parameters α_i and β_i correspond to the i -th SGTm.

3.1.4 Auxiliary Network

To further ensure that the generated thumbnail aligns well with the query sentence, we introduce an auxiliary network (T_{aux}) next to the decoder T_{dec} of T (Sec. 3.1.1). It is a small network whose input is the output of *deconv2* layer in T_{dec} which is forwarded to a 1×1 convolution and an average pooling operation so as to reconstruct the user sentence vector representation \mathbf{z} learned by S_{enc} (Sec. 3.1.2).

3.2 Learning and Optimization

Our learning objective includes a thumbnail loss and an auxiliary loss.

Thumbnail loss: For an input video V with C clips and the ground-truth binary indicator label vector denoting whether a clip belongs to thumbnail or not, we define a cross-entropy loss \mathcal{L}_{thumb} on the prediction of the video thumbnail generation network T as:

$$\mathcal{L}_{thumb} = -\frac{1}{C} \sum_{c=1}^C \log \left(\frac{\exp(\delta_{c,t_c})}{\sum_{j=1}^2 \exp(\delta_{c,j})} \right), \quad (4)$$

where $\delta_{c,j}$ is the predicted score of c -th video clip to be labeled as j -th class (non-thumbnailed or thumbnailed) and δ_{c,l_c} is the score for the ground-truth class l_c for the c -th video clip.

Auxiliary loss: This loss aims to minimize the difference between the sentence representation \mathbf{z} from S_{enc} and the reconstructed sentence representation $\hat{\mathbf{z}}$ from T_{aux} . We define the reconstruction loss \mathcal{L}_{aux} as:

$$\mathcal{L}_{aux} = \|\mathbf{z} - \hat{\mathbf{z}}\|^2, \quad (5)$$

where $\mathbf{z}, \hat{\mathbf{z}} \in \mathbb{R}^{D_w}$ and $\|\cdot\|$ denotes the L_2 norm.

Final loss: We define the final loss \mathcal{L}_{final} as:

$$\mathcal{L}_{final} = \mathcal{L}_{thumb} + \mathcal{L}_{aux}. \quad (6)$$

The aim of the learning is to find the optimal parameters Θ_T^* , $\Theta_{S_{enc}}^*$ and $\Theta_{T_{aux}}^*$ of the networks T , S_{enc} and T_{aux} , respectively. We can express the learning objective as follows:

$$\Theta_T^*, \Theta_{S_{enc}}^*, \Theta_{T_{aux}}^* = \arg \min_{\Theta_T, \Theta_{S_{enc}}, \Theta_{T_{aux}}} \mathcal{L}_{final}(T, S_{enc}, T_{aux}). \quad (7)$$

For simplicity, we denote our proposed sentence guided dynamic video thumbnail generation model (Sec. 3.1) learned from Eq. 7 as Guided-DVTG.

In summary, Guided-DVTG captures the global information of the video through its video thumbnail generation network (Sec. 3.1.1) which is crucial to generate thumbnails that provide overall content preview. It also has the ability to dynamically modulate its prediction using the user sentence via the sentence guided temporal modulation (Sec. 3.1.3). Moreover, it further ensures semantic correspondence with the sentence using the auxiliary network (Sec. 3.1.4). As a result, Guided-DVTG can produce dynamic video thumbnails that provide a quick preview of the video while satisfying the user query sentence.

4 Experiments

4.1 Setup

Dataset: We conduct experiments on the sentence specified dynamic video thumbnail generation dataset by Yuan *et al.* [57]. This dataset is based on the ActivityNet Captions dataset [16]. It has 10,204 video-sentence pairs where each pair is labeled with 4 video thumbnail annotations. The thumbnail annotation for each video is at clip-level where each clip is of 2 seconds and with no more than 5 clips from the video included in the final video thumbnail. 70% of the dataset is used for training, 15% for validation and the rest 15% for testing.

Feature representation: We follow prior work [57] for video and word embedding representation. We evenly split every video in the dataset into 2 second clips and represent each clip with the C3D features [29] provided by the ActivityNet Challenge 2016. For each word, we obtain a 300 dimensional word embedding using Glove [25].

Training details: We train all our models from scratch with a constant learning rate 0.0001 using the Adam optimizer [14]. During training, for a video-sentence pair, we find the most consistent video thumbnail annotation among the 4 thumbnail annotations and treat it as the ground-truth annotation. However, in testing, we evaluate the predicted video thumbnail by comparing it against all the 4 thumbnail annotations. Note that similar process is followed by previous work [37].

Evaluation metrics: Following prior work [57], we use the F1 and IoU scores to measure the performance of our methods. These metrics measure the agreement between the generated thumbnail and the ground-truth thumbnail annotations. A higher value is desirable on these metrics.

4.2 Main Results and Comparisons

In addition to comparing with prior state-of-the-art video thumbnail generation methods, we also define several strong baselines that are as follows:

FCSN [26]: This is a state-of-the-art model in video summarization task. We extend it to create our video thumbnail generation network T (see Sec. 3.1.1). We directly train and evaluate FCSN on the dataset in this paper. Note that this is a generic video thumbnail model that does not consider the user query sentence.

IN-FCSN: This baseline is a variant of video thumbnail generation network T where we replace the proposed SGTM mechanism (see Sec. 3.1.3) with the temporal instance normalization layer [30] with learnable affine transformation parameters. Note that we do not have the networks S_{enc} and T_{aux} in this model. This results in another generic thumbnail model.

IN-FCSN-concat: We obtain this baseline when we train video thumbnail generation network T by concatenating the user sentence vector representation (obtained by averaging the words embedding) with the video-clip features of the input video. We again replace SGTM mechanism with temporal instance normalization layer with learnable affine transformation parameters. This results in a dynamic video thumbnail generation model as it incorporates user sentence information in the model.

In Table 1, we compare our final model Guided-DVTG with the prior and baseline methods. We outperform the baselines and other alternative methods except GTP [57]. Unlike GTP, we do not use sophisticated multimodal feature fusion. Moreover, GTP uses recurrent models that perform sequential computation within training samples which prevents parallelization, whereas our model is completely non-recurrent that allows much more parallelization. Figure 3 shows example video thumbnails generated by our Guided-DVTG model.

Method	F1	IoU
Random	0.3604	0.2379
RankNet [10]	0.4013	0.2770
VSEM [19]	0.4386	0.3098
QARE [51]	0.4285	0.2986
CTRL [8]	0.4303	0.3084
ACRN [18]	0.4456	0.3271
GTP [57]	0.5285	0.3933
FCSN [26] (ours)	0.4295	0.3101
IN-FCSN (ours)	0.4426	0.3140
IN-FCSN-concat (ours)	0.4286	0.3084
Guided-DVTG (ours)	0.4758	0.3405

Table 1: Performance comparison (in terms of F1 and IoU) between Guided-DVTG and other alternative methods. Results of previous methods are taken from [57]. Best and second best methods are highlighted in gray and cyan, respectively.




Figure 3: Example qualitative results produced by Guided-DVTG. The gray, green and orange bars indicate the video length, ground-truth and thumbnail predictions, respectively.

4.3 Analysis

Role of Modulation Parameters: We analyze the importance of modulation parameters α_c and β_c (Eq. 3) in the SGTm mechanism (Sec. 3.1.3). Table 2 compares the performance for different possible solutions of these parameters. We find that when these parameters are predicted ($\alpha_c=\alpha_c^s$, $\beta_c=\beta_c^s$) from another network using the user query sentence (i.e., in Guided-DVTG), the model achieves much better performance as compared to cases when they are set to fixed values ($\alpha_c=1$, $\beta_c=0$) or directly learned ($\alpha_c=\alpha_c^*$, $\beta_c=\beta_c^*$) in the main video thumbnail generation network (Sec. 3.1.1). Therefore, the proposed SGTm is key to dynamic video thumbnail generation.

Method	$\alpha_c = 1, \beta_c = 0$	$\alpha_c = \alpha_c^*, \beta_c = \beta_c^*$	$\alpha_c = \alpha_c^s, \beta_c = \beta_c^s$
IN-FCSN	0.4062 (0.2904)	0.4426 (0.3140)	-
IN-FCSN-concat	0.4480 (0.3192)	0.4286 (0.3084)	-
Guided-DVTG	-	-	0.4758 (0.3405)

Table 2: Impact of modulation parameters on video thumbnail generation. Here we indicate F1 and IoU (in bracket) for different solutions of parameters α_c and β_c in Eq. 3.

Impact of Auxiliary Network: We study the impact of auxiliary network (Sec. 3.1.4) and the loss \mathcal{L}_{aux} (Sec. 3.2) associated with it in learning the thumbnail model. In order to verify their contribution, we remove them from our final model Guided-DVTG and perform learning. We call the learned model Guided-DVTG-NA and compare the performance in Table 3(a). We notice a drop in performance which highlights the importance of the auxiliary network and its loss in our Guided-DVTG model.

Unsupervised Guided-DVTG: We develop an unsupervised variant of our Guided-DVTG model. When we remove the supervised loss \mathcal{L}_{thumb} (Sec. 3.2 and Eq. 6) and perform learning only using the auxiliary loss \mathcal{L}_{aux} which is completely unsupervised, we obtain the unsupervised version of our model to which we call Guided-DVTG_{unsup}. In Table 3(b), we compare its performance with the state-of-the-art unsupervised method, BeautThumb [28], when evaluated on the dataset in this paper. Our model Guided-DVTG_{unsup} significantly outperforms BeautThumb [28]. This result is very appealing since gathering labeled video thumbnail data is extremely expensive.

Method	F1	IoU
Guided-DVTG	0.4758	0.3405
Guided-DVTG-NA	0.4644	0.3296

(a)

Method	F1	IoU
Random	0.3604	0.2379
BeautThumb [28]	0.3837	0.2544
Guided-DVTG _{unsup}	0.4222	0.2835

(b)

Table 3: (a) Impact of auxiliary network and its loss. (b) Performance comparison of unsupervised methods. Result of BeautThumb [28] is taken from [57].

5 Conclusion

In this paper, we have proposed a simple yet effective framework for the DVTG task. At the core of our framework is the proposed SGTm mechanism that modulates the normalized temporal activations in the video thumbnail generation network to effectively correlate sentence-specific video clips over time. Instead of applying recurrent neural architectures, we propose a non-recurrent solution that offers much more parallelization on GPU hardware. Our proposed framework does not involve complex multimodal feature fusion commonly used in vision-language tasks such as DVTG. The experimental results and analysis on a large-scale dataset demonstrate that our proposed method achieves superior or competitive performance against the state-of-the-art methods.

6 Acknowledgements

The authors acknowledge financial support from NSERC and UMGF funding. We thank NVIDIA for donating some of the GPUs used in this work.

References

- [1] Lisa Anne Hendricks, Oliver Wang, Eli Shechtman, Josef Sivic, Trevor Darrell, and Bryan Russell. Localizing moments in video with natural language. In *IEEE International Conference on Computer Vision*, pages 5803–5812, 2017.
- [2] Kyunghyun Cho, Bart van Merriënboer, Caglar Gulcehre, Dzmitry Bahdanau, Fethi Bougares, Holger Schwenk, and Yoshua Bengio. Learning phrase representations using rnn encoder–decoder for statistical machine translation. In *Conference on Empirical Methods in Natural Language Processing*, pages 1724–1734, 2014.
- [3] Sally Jo Cunningham and David M Nichols. How people find videos. In *ACM/IEEE Joint Conference on Digital Libraries*, pages 201–210, 2008.
- [4] Harm De Vries, Florian Strub, Jérémie Mary, Hugo Larochelle, Olivier Pietquin, and Aaron C Courville. Modulating early visual processing by language. In *Advances in Neural Information Processing Systems*, 2017.
- [5] Ana García del Molino and Michael Gygli. Phd-gifs: Personalized highlight detection for automatic gif creation. In *ACM International Conference on Multimedia*, 2018.
- [6] F Dirfaux. Key frame selection to represent a video. In *IEEE International Conference on Image Processing*, pages 275–278, 2000.
- [7] Andrea Frome, Greg S Corrado, Jon Shlens, Samy Bengio, Jeff Dean, Marc’Aurelio Ranzato, and Tomas Mikolov. Devise: A deep visual-semantic embedding model. In *Advances in Neural Information Processing Systems*, pages 2121–2129, 2013.
- [8] Jiyang Gao, Chen Sun, Zhenheng Yang, and Ram Nevatia. Tall: Temporal activity localization via language query. In *IEEE International Conference on Computer Vision*, pages 5267–5275, 2017.

- [9] Yuli Gao, Tong Zhang, and Jun Xiao. Thematic video thumbnail selection. In *IEEE International Conference on Image Processing*, pages 4333–4336, 2009.
- [10] Michael Gygli, Yale Song, and Liangliang Cao. Video2gif: Automatic generation of animated gifs from video. In *IEEE Conference on Computer Vision and Pattern Recognition*, 2016.
- [11] Satoshi Hasebe, Makoto Nagumo, Shogo Muramatsu, and Hisakazu Kikuchi. Video key frame selection by clustering wavelet coefficients. In *European Signal Processing Conference*, pages 2303–2306, 2004.
- [12] Xun Huang and Serge Belongie. Arbitrary style transfer in real-time with adaptive instance normalization. In *IEEE International Conference on Computer Vision*, 2017.
- [13] Hong-Wen Kang and Xian-Sheng Hua. To learn representativeness of video frames. In *ACM International Conference on Multimedia*, pages 423–426, 2005.
- [14] Diederik Kingma and Jimmy Ba. Adam: A method for stochastic optimization. In *International Conference on Learning Representations*, 2015.
- [15] Thomas N. Kipf and Max Welling. Semi-supervised classification with graph convolutional networks. In *International Conference on Learning Representations*, 2017.
- [16] Ranjay Krishna, Kenji Hata, Frederic Ren, Li Fei-Fei, and Juan Carlos Nieves. Dense-captioning events in videos. In *International Conference on Computer Vision*, pages 706–715, 2017.
- [17] Chunxi Liu, Qingming Huang, and Shuqiang Jiang. Query sensitive dynamic web video thumbnail generation. In *IEEE International Conference on Image Processing*, pages 2449–2452, 2011.
- [18] Meng Liu, Xiang Wang, Liqiang Nie, Xiangnan He, Baoquan Chen, and Tat-Seng Chua. Attentive moment retrieval in videos. In *ACM SIGIR Conference on Research & Development in Information Retrieval*, pages 15–24, 2018.
- [19] Wu Liu, Tao Mei, Yongdong Zhang, Cherry Che, and Jiebo Luo. Multi-task deep visual-semantic embedding for video thumbnail selection. In *IEEE Conference on Computer Vision and Pattern Recognition*, pages 3707–3715, 2015.
- [20] Jiebo Luo, Christophe Papin, and Kathleen Costello. Towards extracting semantically meaningful key frames from personal video clips: from humans to computers. *IEEE Transactions on Circuits and Systems for Video Technology*, 19(2):289–301, 2008.
- [21] Tao Mei, Bo Yang, Shi-Qiang Yang, and Xian-Sheng Hua. Video collage: presenting a video sequence using a single image. *The Visual Computer*, 25(1):39–51, 2009.
- [22] Yingwei Pan, Ting Yao, Tao Mei, Houqiang Li, Chong-Wah Ngo, and Yong Rui. Click-through-based cross-view learning for image search. In *International ACM SIGIR Conference on Research and Development in Information Retrieval*, pages 717–726, 2014.
- [23] Ankur Parikh, Oscar Täckström, Dipanjan Das, and Jakob Uszkoreit. A decomposable attention model for natural language inference. In *Conference on Empirical Methods in Natural Language Processing*, pages 2249–2255, 2016.

- [24] Taesung Park, Ming-Yu Liu, Ting-Chun Wang, and Jun-Yan Zhu. Semantic image synthesis with spatially-adaptive normalization. In *IEEE Conference on Computer Vision and Pattern Recognition*, 2019.
- [25] Jeffrey Pennington, Richard Socher, and Christopher D Manning. Glove: Global vectors for word representation. In *Conference on Empirical Methods in Natural Language Processing*, pages 1532–1543, 2014.
- [26] Mrigank Rochan, Linwei Ye, and Yang Wang. Video summarization using fully convolutional sequence networks. In *European Conference on Computer Vision*, 2018.
- [27] Mrigank Rochan, Mahesh Kumar Krishna Reddy, Linwei Ye, and Yang Wang. Adaptive video highlight detection by learning from user history. In *European Conference on Computer Vision*, 2020.
- [28] Yale Song, Miriam Redi, Jordi Vallmitjana, and Alejandro Jaimes. To click or not to click: Automatic selection of beautiful thumbnails from videos. In *ACM International Conference on Information and Knowledge Management*, pages 659–668, 2016.
- [29] Du Tran, Lubomir Bourdev, Rob Fergus, Lorenzo Torresani, and Manohar Paluri. Learning spatiotemporal features with 3d convolutional networks. In *IEEE International Conference on Computer Vision*, pages 4489–4497, 2015.
- [30] Dmitry Ulyanov, Andrea Vedaldi, and Victor Lempitsky. Instance normalization: The missing ingredient for fast stylization. *arXiv preprint arXiv:1607.08022*, 2016.
- [31] Arun Balajee Vasudevan, Michael Gygli, Anna Volokitin, and Luc Van Gool. Query-adaptive video summarization via quality-aware relevance estimation. In *ACM International Conference on Multimedia*, pages 582–590, 2017.
- [32] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. Attention is all you need. In *Advances in Neural Information Processing Systems*, pages 5998–6008, 2017.
- [33] Oriol Vinyals, Meire Fortunato, and Navdeep Jaitly. Pointer networks. In *Advances in Neural Information Processing Systems*, pages 2692–2700, 2015.
- [34] Xiaolong Wang, Ross Girshick, Abhinav Gupta, and Kaiming He. Non-local neural networks. In *IEEE Conference on Computer Vision and Pattern Recognition*, pages 7794–7803, 2018.
- [35] Zheshen Wang, Mrityunjay Kumar, Jiebo Luo, and Baoxin Li. Extracting key frames from consumer videos using bi-layer group sparsity. In *ACM International Conference on Multimedia*, pages 1505–1508, 2011.
- [36] Yitian Yuan, Lin Ma, Jingwen Wang, Wei Liu, and Wenwu Zhu. Semantic conditioned dynamic modulation for temporal sentence grounding in videos. In *Advances in Neural Information Processing Systems*, pages 534–544, 2019.
- [37] Yitian Yuan, Lin Ma, and Wenwu Zhu. Sentence specified dynamic video thumbnail generation. In *ACM International Conference on Multimedia*, pages 2332–2340, 2019.

-
- [38] Han Zhang, Ian Goodfellow, Dimitris Metaxas, and Augustus Odena. Self-attention generative adversarial networks. In *International Conference on Machine Learning*, 2019.