Query Twice: Dual Mixture Attention Meta Learning for Video Summarization

Junyan Wang[1], Yang Bai[2], Yang Long[3],
Bingzhang Hu[2], Zhenhua Chai[1], Yu Guan[2], Xiaolin Wei[1]
[1]Vision Intelligence Center, Meituan-Dianping Group, Beijing, China
[2]OpenLab, Newcastle University, Newcastle upon Tyne, UK
[3]Department of Computer Science, Durham University, Durham, UK
{wangjunyan04, chaizhenhua, weixiaolin02}@meituan.com,
{y.bai13, bingzhang.hu, yu.guan}@newcastle.ac.uk, yang.long@ieee.org

ABSTRACT

Video summarization aims to select representative frames to retain high-level information, which is usually solved by predicting the segment-wise importance score via a softmax function. However, softmax function suffers in retaining high-rank representations for complex visual or sequential information, which is known as the *Softmax Bottleneck* problem. In this paper, we propose a novel framework named Dual Mixture Attention (DMASum) model with Meta Learning for video summarization that tackles the softmax bottleneck problem, where the Mixture of Attention layer (MoA) effectively increases the model capacity by employing twice self-query attention that can capture the second-order changes in addition to the initial query-key attention, and a novel Single Frame Meta Learning rule is then introduced to achieve more generalization to small datasets with limited training sources. Furthermore, the DMASum significantly exploits both visual and sequential attention that connects local key-frame and global attention in an accumulative way. We adopt the new evaluation protocol on two public datasets, SumMe, and TVSum. Both qualitative and quantitative experiments manifest significant improvements over the state-of-the-art methods.

CCS CONCEPTS

• **Computing methodologies** → **Video summarization**.

KEYWORDS

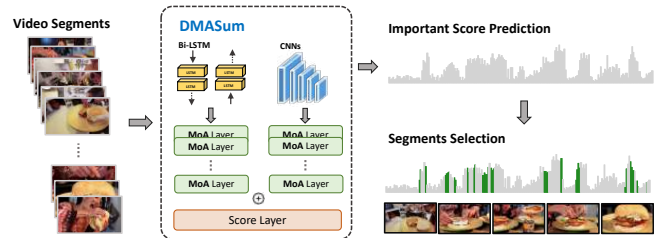video summarization, attention network, meta learning

Figure 1: An illustration of the video summarization task using our proposed DMASum. Each gray bar represents the predicted important score of a segment and green bars denote the key-segments in the summarized video. Highlights of DMASum include Visual-sequential Dual Channels, Stacked MoA modules.

1 INTRODUCTION

With the tremendous growth of video materials uploaded to various online video platforms like YouTube, automatic video summarization has received increasing attention in recent years. The summarized video can be used in many scenarios such as fast indexing and human-computer interaction in a light and convenient fashion. The main objective of video summarization is to shorten a whole video into summarized frames while preserving crucial plots. One of the mainstream directions focuses on key-frames summarization [8] is illustrated in Fig. 1 A video is first divided into 15-second segments, and the problem is modeled as an importance score prediction task to select the most informative segments.

The nature of video summarization task encourages a line of research [13, 21, 32, 35] focusing on unsupervised learning methods. Besides, [39] applied deep reinforcement learning with a diversity-representativeness reward function for the generated summary; Currently, the most popular benchmarks are SumMe [8] and TVSum [25]. Otani *et al.* [23] proposed to evaluate the methods by using the rank-order correlation between predicted and human-annotated importance scores. These key evaluation matrices measure agreements between generated summaries and reference summaries. Therefore, supervised methods [7, 32, 36, 38] are still very important for investigating essential technical questions because they can directly compare against human-annotated scores as ground truth. One of the mainstream directions focuses on key-frames summarization [8] is illustrated in Fig. 1.

The challenges for supervised key-frames summarization are two-fold. First, the importance scores are very subjective and highly related to human perception. Second, the annotations are expensive to be obtained; thus, the model should be able to cope with limited labeled data while retaining high generalization. These are not only unsolved questions for video summarization but also essential for many other research domains. To this end, this paper proposes a new framework, namely the Dual Mixture Attention model (DMA-Sum) that aims to achieve **1)** human-like attention by adopting cutting-edge self-attention architecture and takes both visual and sequential information into a unified process; and **2)** high-level semantic understanding of the whole content by incorporating a novel meta learning module to maximally exploit the training data and improve the model generalization.

The proposed framework manifested promising results in our early experiments. However, the early implementation reflected two major technical challenges. The first is known as the *Softmax Bottleneck* problem associated with the self-attention architecture. Both theoretical and empirical evidences in this paper show that traditional softmax function does not have the sufficient capability to retain high-rank representation for long and complex videos. To this end, we propose a *Query Twice* module by adding self-query attention to query-key attention. The Mixture of Attention layer can then compare the two attentions to capture the second-order changes and increase the model capability. The second problem is that the most common meta learning strategy does not naturally fit the video summarization task. We propose a Single-video Meta Learning rule to refrain the learner tasks so as to purify the meta learner updating processes. To summarize our contributions:

- To our best knowledge, this is the first paper that successfully introduces self-attention architecture and meta learning to jointly process dual representations of visual and sequential information for video summarization.
- We provide in-depth theoretical and empirical analyses of the Softmax Bottleneck problem when applying attention model to video summarization task. And a novel self-query module with Mixture-of-Attention is provided as the solution to overcome the problem effectively.
- We explore the meta learning strategy, and a Single-Video Meta Learning rule is particularly designed for video summarization tasks.
- Quantitatively and qualitatively experiments on two datasets: SumMe [8] and TVSum [25] demonstrate our superior performance over the state-of-the-art methods. More impressively, our model achieves human annotator level performance under new protocols of Kendall's τ correlation coefficients and Spearman's ρ correlation coefficients. The groundbreaking results suggested that our DMASum has effectively modeled human-like attention.

2 RELATED WORK

Video summarization. Video, as a media containing complex spatio-temporal relationship of visual contents, has a wide range of applications [4, 5, 28, 29, 33, 41]. However, because of its huge volume, video summarization is to compress such huge volume data into its light version while preserving its information. Early

works have presented various solutions to this problem, including storyboards [7, 9, 20, 20] and objects [16, 19, 30]. LSTM-based deep learning approaches are proposed for both supervised and unsupervised video summarization in recent years. Zhang *et al.* [36] proposed a bidirectional LSTM model to predict the importance score of each frame directly, and this model is also extended with determinantal point process [15]. Mahasseni *et al.* [21] specified a generative adversarial framework that consists of the summarizer and discriminator for unsupervised video summarization. The summarizer is an auto-encoder LSTM network for reconstructing the input video, and the discriminator is another LSTM network for distinguishing between the original video and its reconstruction. Meanwhile, based on the observation of Otani *et al.* [23], they propose another evaluation approach as well as a visualization of correlation between the estimated scoring and human annotations.
Attention-based Models. The attention mechanism was born to help memorize long source sentences in neural machine translation [2]. Rather than building a single context vector out of the translation encoder, the attention method is to create shortcuts between the context vector and the entire input sentence, then customize the weights of these shortcut connections for each element. The Transformer [27], without a doubt, is one of the most impressive works in the machine translation task. The model is mainly built on self-attention layers, also known as intra-attention, and the self-attention network is relating different positions of the same input sequence. Many recent works have applied self-attention to a wide range of video-related applications, such as video question answer [18] and video captioning [40]. Particularly for the video summarization task, Ji *et al.* [12] proposed an attention-based encoder-decoder network for selecting the key shots. He *et al.* [11] proposed an unsupervised video summarization method with attentive conditional Generative Adversarial Networks.
Meta Learning. Meta learning, also known as learning to learn, aims to design a model that can be learned rapidly with fewer training examples. Meta learning usually used in few-shot learning [6, 22] and transfer learning [31]. Finn *et al.* [6] propose a Model Agnostic Meta Learning (MAML) which is compatible with any model trained with gradient descent and applicable to a variety of different learning problems, including classification, regression, and reinforcement learning. Like MAML, the work of Nichol *et al.* [22] proposed a strategy which repeatedly sampling and training a single task, then moving the initialization towards the trained weights on that task. Recently, meta learning methods have been applied in a few video analysis tasks. Especially in video summarization, Li *et al.* [17] proposed a meta learning method that explores the video summarization mechanism among summarizing processes on different videos.

3 THE PROPOSED APPROACH

Video summarization is modeled as a sequence labeling (or sequence to sequence mapping) problem. Given a sequence of video frames, the task is to assign each frame an importance score based on which key-frames can be selected. Existing sequence labelling approaches include deep sequential models such as LSTM [36, 37], attention model [12]. However, the key difficulty is to learn the frame dependencies within the video and capturing the internal
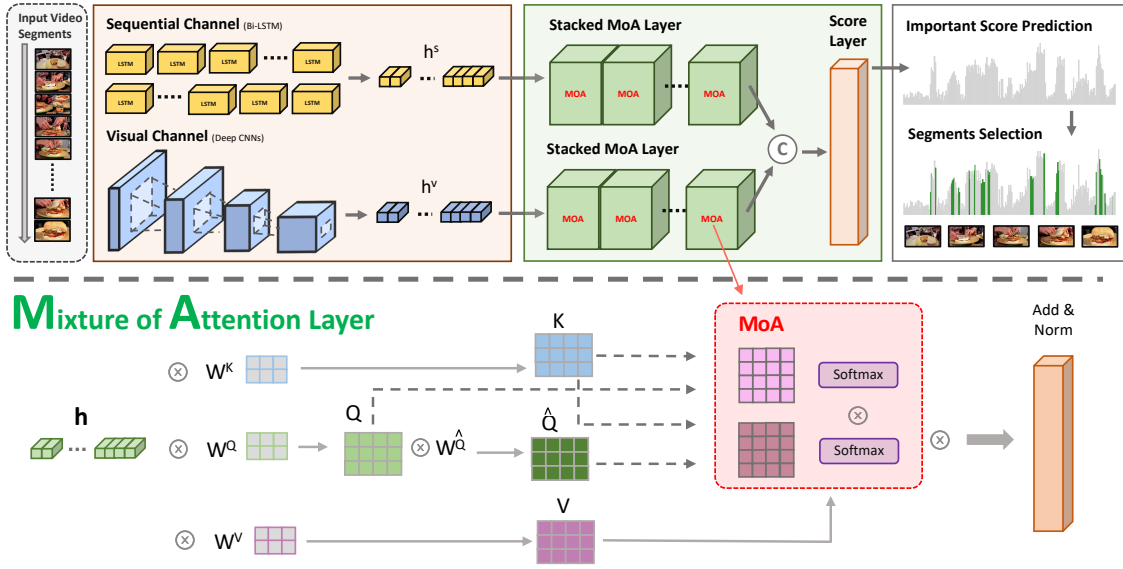
Figure 2: The overall architecture of our DMASum is shown as the top figure, which consists of a sequential channel and a visual channel and stacked MoA layers. The bottom part shows the structure of the Mixture of Attention layer.

contextual information of the video. Considering video is a highly context-dependent source that shares many similar properties in sentences. As the outstanding performance of the Transformer [27], we introduce the self-attention structure that has been widely used in natural language processing (NLP) as our architecture basis. Both visual and sequential representations are considered in order to model complex human-like attention and better match the subjective annotations. Also, the motivation of meta learning aims to improve the model generalization when training sources are insufficient due to expensive human annotations. An overview of the proposed video summarization architecture and the details of the Mixture of Attention layer that are illustrated in Figure 2.

3.1 Architecture Design

Dual-representation Learning: For the video summarization task, we introduce both visual and sequential channels as the input. The visual channel (deep CNNs) extracts visual features $H^v = \{h_t^v\}_{t=1}^T$ from each video frame image. Based on the extracted visual features, the sequential features $H^s = \{h_t^s\}_{t=1}^T$ is obtained by the sequential channel (bidirectional LSTM network) and consists of the dual-channel feature $H \in \{H^v, H^s\}$. The dual representation is critical to model complex human-like attention and can link frame-wise attention to the overall story line.

The Attention Module: Taking a feature sequence $H = \{h_t\}_{t=1}^T \in \mathbb{R}^{D \times T}$ extracted from the video as input, the attention network can re-express each h_t^* within input H by utilizing weighted combination of the entire neighborhood from h_1 to h_T, where D is the feature dimension and T is number of frames within a video. In concreteness, the attention network first linearly transforms H into $Q = W^Q H^*$, $K = W^K H^*$ and $V = W^V H^*$, where $Q = \{Q_t\}_{t=1}^T \in \mathbb{R}^{D_a \times T}$, $K = \{K_t\}_{t=1}^T \in \mathbb{R}^{D_a \times T}$ and $V = \{V_t\}_{t=1}^T \in \mathbb{R}^{D_a \times T}$ are known as Queries, Keys and Values vectors, respectively and D_a represents the attention feature size, and $W^Q, W^K, W^V \in \mathbb{R}^{D_a \times D}$

are the corresponding learnable parameters. K is employed to learn the distribution of attention matrix on condition of the query matrix Q, and V is used to exploit information representation. Thus the scaled dot-product attention A is defined as:

$$\mathcal{F}_{Scale}(K, Q) = \frac{K^T Q}{\sqrt{D_a}} , \tag{1}$$

$$A = \mathcal{F}_{Softmax}(K, Q) = \frac{\exp(\mathcal{F}_{Scale}(K, Q))}{\sum_{t=1}^T \exp \mathcal{F}_{Scale}(K, Q)} , \tag{2}$$

where $A \in \mathbb{R}^{T \times T}$ and we consider A as the distribution of attention matrix on condition of the query matrix Q. In Eq.1, due to the large degree of high dimensional $K^T Q$, scaling factor $\frac{1}{\sqrt{D_a}}$ is used to prevent the potential small gradient suffered by softmax. The output of attention network is:

$$Z = V A. \tag{3}$$

After applying the attention module to both channels, We concatenate their outputs and feed into a score layer, which consists of multiple fully-connected layers ended with a sigmoid function. The score layer predicts the importance score $\hat{s}$ is sampled as:

$$\hat{S} = \mathcal{F}_{Score}(\mathcal{F}_{Concat}(Z^v, Z^s)), \tag{4}$$

where $\mathcal{F}_{Score}$ denotes the score layer and $\mathcal{F}_{Concat}$ in this paper means concatenation operation on different channels.

Overall Objective Function. We intend to treat the outputs as the importance scores of the whole video frames in this work. Thus, we simply employ the mean square loss $\mathcal{L}$ between the ground truth importance scores and the predicted importance scores.

$$\mathcal{L} = \frac{1}{T} \sum_{i=1}^T (s_i - \hat{s_i})^2, \tag{5}$$

3.2 The Softmax Bottleneck

Almost all existing attention models follow the original pipeline from NLP tasks using the softmax function Eq. (2) to compute the attention. However, this section identifies the key limitation of softmax function for video summarization. It can be considered that the attention distribution is a finite set of pairs of a context and its conditional distribution $\mathcal{V} = \{(c_1, P^*(X|c_1)), \ldots, (c_1, P^*(X|c_T))\}$, where $X = \{x_1, x_2, \ldots, x_N\}$ denotes T compatible keys in the video $\mathcal{V}$ and $C = \{c_1, c_2, \ldots, c_N\}$ denotes the contexts. It is assumed $P^* > 0$ and A^* represents the true attention distribution. Thus the true attention distribution in (2) can be re-formulated as:

$$A^* = \begin{bmatrix} logP^*(x_1|c_1) & logP^*(x_2|c_1) & \cdots & logP^*(x_T|c_1) \\ logP^*(x_1|c_2) & logP^*(x_2|c_2) & \cdots & logP^*(x_T|c_2) \\ \vdots & \vdots & \ddots & \vdots \\ logP^*(x_1|c_T) & logP^*(x_2|c_T) & \cdots & logP^*(x_T|c_T). \end{bmatrix} \quad (6)$$

The objective of attention model is to learn the conditional attention distribution $P_\theta(X|C)$ parameterized by θ to match the true attention distribution $P^*(X|C)$. It can be seen that the attention distribution problem is now turned into a **matrix factorization problem**. Since A is a matrix with size $N \times N$, the rank of learned attention distribution A is upper bounded by the embedding size d. If $d < rank(A^*) - 1$, for any model parameter θ, there exists a context c in $\mathcal{V}$ such that $P_\theta(X|C) \neq P^*(X|C)$. This is so called **Softmax Bottleneck** which reflects the circumstance when softmax function does not have the capacity to express the true attention distribution when d is smaller than $rank(A^*) - 1$. In the contexts of video summarization, the log probability matrix A becomes a high-rank matrix when the visual contents are complex and the changes between frames are severe. For example, cooking may contain multiple repetitive actions than eating. While humans can intuitively assign equal importance to both of the actions, the former one actually results in a much higher rank in the representation matrix. The softmax function may compromise features from rich content to maintain consistency.

In Figure 3 we empirically verify such a Softmax Bottleneck problem can degrade the performance severely. We choose the TVSum dataset and calculate the difference $\mathcal{D} = T - rank(A)$, where T denotes the video length. This is because video lengths are not consistent so we only consider the difference between the actual rank and the full rank T. Lower difference values indicate the attention layer, after softmax, can retain high rank with minimum redundancy. On the other hand, Higher difference values mean the attention matrix of the whole video is low-rank. It can be due to the input video is not complex, e.g. no movement and the background is monotonous. But for most of the cases, the low-rank attention matrix is often resulted by key information missing due to long videos with high complexities. The statistics are collected from attention matrices of both visual and sequential channels. Our key observations are summarised as follows.

(1) From Figure 3 (a) and (b), higher rank representations tend to achieve higher F1 score. But due to the softmax capacity, significant performance drops can be seen in visual (after range 8-11) and sequential channels (after range 4-7), which confirms the existence of bottleneck. In other words, the
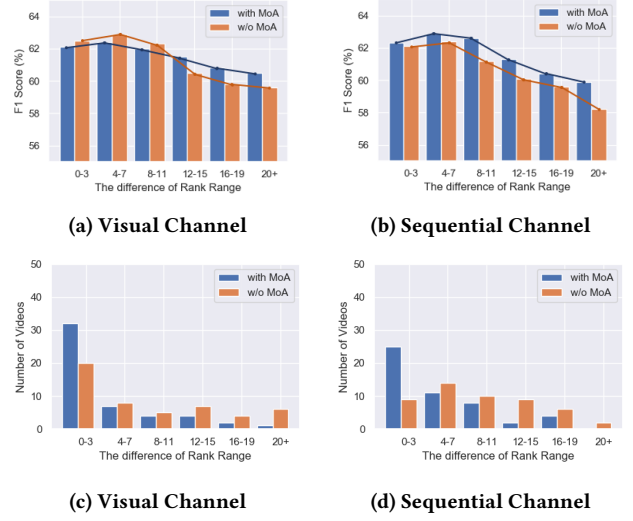

(a) Visual Channel   **(b) Sequential Channel**

(c) Visual Channel   **(d) Sequential Channel**

Figure 3: Averaged F1-score (%) and Number of videos with respect to the rank difference $\mathcal{D}$ in TVSum dataset. Blue and Orange bars compare our MoA against traditional softmax.

softmax function cannot retain high-rank information for long complex videos.

(2) From the distribution of video numbers in Figure 3 (c) and (d), many video representations fall out of high-rank range (0-7) after softmax. According to the last observation, these videos are prone to getting lower performances.

(3) The softmax bottleneck problem is more severe on sequential attention, which indicates the changes between frames are the key missing information that results in the lower rank.

Motivated by the above insights and inspired by the work of Yang *et al.* [34], we come up with a **Mixture of Attention layer** (MoA) to alleviate the softmax bottleneck issue. We propose the Associated Query $\hat{Q} = tanh(W^{\hat{Q}}Q)$, where $W^{\hat{Q}}$ is the Associated Query parameter. The idea is to capture the second-order changes between queries so that both complex and simple contents can be represented in a more smoothed attention representation. The conditional attention distribution is defined as:

$$P(x|c) = \sum_{t=1}^{T} \frac{\exp(\mathcal{F}_{Scale}(K_{c,t}, Q_{c,t}))}{\sum_{t=1}^{T} \exp\mathcal{F}_{Scale}(K_{c,t}, Q_{c,t})} \hat{A}_{c,t} ,$$

$$s.t. \sum_{t=1}^{T} \hat{A}_{c,t} = 1 , \quad (7)$$

$$where \ \hat{A} = \mathcal{F}_{Softmax}(K, \hat{Q}) , \quad (8)$$

In Eq.8, $\hat{A} \in T \times T$ is the associated attention distribution. Thus, MoA formulates the conditional attention distribution as:

$$A_{moa} = A\hat{A}^T , \quad (9)$$

where $A_{moa} \in \mathbb{R}^{T \times T}$. In Eq.1, due to the large degree of high dimensional $K^T Q$, scaling factor $\frac{1}{\sqrt{D_a}}$ is used to prevent the potential
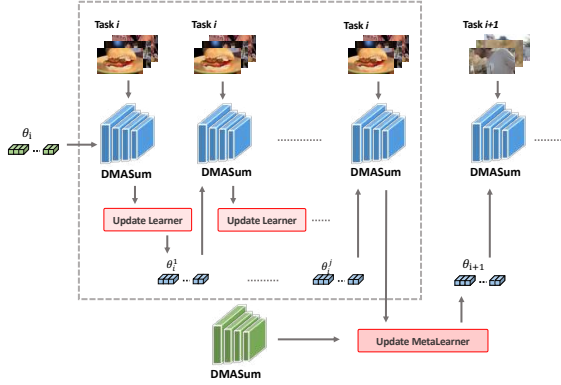
Figure 4: Overview of the i^{th} iteration for update θ_i to θ_{i+1}. There are two stages in this update process. The middle part shows the stage about how the Learner updates θ_i to θ_i^m by iterating m times. The outside parts shows the stage about how the Meta Learner updates θ_i to θ_{i+1}.

small gradient suffered by softmax. As A_{moa} is a non-linear function of the attention distribution, A_{moa} can be arbitrarily higher rank than standard self-attention structure A. Thus the output of the mixture of attention network $Z = VA_{moa}$ now can break the bottleneck problem. In Figure 3, after applying the MoA, we can see a large proportion of videos fall into the 0-3 high rank range compared that of traditional softmax. Also, videos especially with lower ranks ($\mathcal{D} > 11$) can be predicted with higher F1 scores. The performance of the sequential channel is boosted, which indicates that all of the previous softmax representations missed high rank information. The smoothed performance drop and increased number of high rank videos serve as strong evidence to manifest the Softmax Bottleneck has been resolved by proposed MoA.

Besides, the DMASum utilizes stacked mixture of attention networks, and in each layer we employ residual dropout connection [10] for allowing gradients to flow through a network directly and layer normalizaion [1] for normalizing the inputs across the features. Overall, the n^{th} layer output can be defined as:

$$Z_n = \mathcal{F}_{Normalize}(\mathcal{F}_{Attention}(Z_{n-1}) \oplus Z_{n-1}), \quad (10)$$

where $\mathcal{F}_{Normalize}$ denotes as layer normalization, $\mathcal{F}_{attention}$ represents the attention layer and $\oplus$ represents the residual connection.

3.3 Single-Video Meta Learning

The key motivation to introduce Meta Learning is to improve the model generalization when the dataset of video summarization is small. Different from gradient descent, the *MetaLearner* is updated by weighted parameters of *Learner* in subtasks, which can be formularized as:

$$Learner^* = MetaLearner(Learner(\tau_i)) \quad (11)$$

where τ_i denotes i^{th} video, *Learner* and *MetaLearner* means the DMASum model in meta learning. We first employ the MAML [6] due to its flexibility and superior performance but did not achieve expected results. Our observation is that in the video summarization

context each video has its own latent mechanism that is not shared by different videos. Therefore, we propose a *Single-Video Meta Learning* rule to refrain the learner by only one video at each task. The process is shown as Figure 4.

There are two stages of each epoch in this meta learning strategy. Firstly, to train the task τ_i, the *Learner* updates the parameter θ_i by traditional gradient descent. And, the *Learner* trains the task in a set number m recurrently to explore its latent summarizing context. The equation of updating parameter θ is:

$$\theta_i^j = \theta_i^{j-1} - \alpha \nabla \mathcal{L}_i^j(\mathcal{F}_{\theta_i^{j-1}}), \quad where\, j = 1 \ldots m \quad (12)$$

where α denotes learning rate and ∇ denoted as the gradient, and $\mathcal{F}_\theta$ is the loss function on i^{th} task. After j^{th} iteration, the *MetaLearner* updates the parameter θ_{i+1} by using the parameter θ_i^m of the *Learner* by:

$$\theta_i = \theta_{i-1} - \beta \nabla \mathcal{L}_i(\mathcal{F}_{\theta_i^m}), \quad (13)$$

where β is the learning rate of the *Learner*. θ_i updated state of *Learner* after the j^{th} iteration in *MetaLearner*. Overall, our meta learning is summarized in Algorithm 1. Note that in the last step of the algorithm, we treat $\theta_i^m - \theta_i$ as a gradient and plug it into Adam instead of simply updating θ_i in the direction $\theta_i^m - \theta_i$.

Algorithm 1: Meta learning in DMASum

```
/* θ : Parameter of Learner;                      */
/* α : Learning rate in Learner;                  */
/* β : Learning rate in MetaLearner;              */
/* n : The number of videos;                      */
/* m : Recurrent training Learner number;         */
/* F : the DMASum model;                          */
Initialize: θ
```
1 **for** *k = 1 to epoch number* **do**
2   **for** *i = 1 to n* **do**
3   Sample video i as task τ_i
4   **for** *j = 1 to m* **do**
5           $\theta_i^j = \theta_i^{j-1} - \alpha \nabla \mathcal{L}_i^j(\mathcal{F}_{\theta_i^{j-1}})$
6   Update $\theta_{i+1} \leftarrow \theta_i + \beta(\theta_i^m - \theta_i)$

---

4 EXPERIMENTS

4.1 Experiment Setup

Datasets. We evaluate our model on two datasets: SumMe [8] and TVSum [25]. SumMe consists of 25 videos covering a variety of events, such as sports and cooking. The duration of each video varies from 1 to 6.5 minutes. TVSum contains 50 videos downloaded from Youtube, which are selected from 10 categories. The video length varies from 1 to 10 minutes. Both datasets include ego-centric and third-person camera views, and the annotations were labeled by 25 human annotators. We also exploit two auxiliary datasets to augment the training data, where Open Video Project[1] (OVP) contains 50 videos and Youtube [3] contains 39 videos.

---

[1]Open video project: https://open-video.org

Evaluation Metrics. We follow the commonly used protocol from [36] and converted the importance scores to shot-based summaries for both datasets, and the user annotations are changed from frame-level scores to key-shots scores using the kernel temporal segmentation (KTS) [24] method, which can temporally segment a video into disjoint intervals. We then compute the harmonic mean F-score as the evaluation metric. In addition, according to the recent evaluation protocol [23], we apply Kendall's τ [14] and Spearman's ρ [42] correlation coefficients for comparing the ordinal association between generated summaries and the ground truth (i.e. the relationship between rankings). Also, they provided correlation curves to visualize the predicted importance score ranking with respect to the reference annotations, i.e., when the predicted importance scores are perfectly concordant with averaged human-annotated scores, the curve lies on the upper bound of the light-blue area. Otherwise, the curve coincides with the lower bound of the area when the ranking of the scores is in reverse order of the reference.
Evaluation Settings. Following [36], we conducted the experiments under three settings. (1) Canonical (C): we used the standard 5-fold cross-validation (5FCV) for SumMe and TVSum datasets. (2) Augmented (A): we used OVP and YouTube datasets to augment the training data in each fold under the 5FCV setting. (3) Transfer (T): we set a target testing dataset, e.g., SumMe or TVSum, and used the other three as the training data.
Implementation details. To be consistent with existing methods, the 1024 dimensional visual features extracted from the *pool*5 layer of the GoogLeNet [26] are used for training. To extract the temporal features, we design a Bi-LSTM model in the proposed network, as a two-layer LSTM with 512 hidden units per layer. For each attention layer, we set the attention dimension as 1024. We stack four attention layers for visual feature attention pipeline, and two layers for the sequential feature attention pipeline. The score layer consists of two fully-connected layers with 1024 hidden units. For Single-video Meta Learning, we set the learning rate of *Learner* as 3×10^{-5} and the learning rate of *MetaLearner* as 6×10^{-5}. Moreover, the recurrent training Learner number is set as 3 and 5 in SumMe and TVSum datasets respectively. During the test, we follow the strategy of prior work [21, 36, 39] to generate the summary. In addition, we employ the ADAM optimizer to train our network and the hyperparameters are optimized via cross-validation.

4.2 Quantitative Evaluation

We first compare our method with state-of-the-art supervised approaches in three evaluation settings. Then, we re-implement the VS-LSTM, SUM-GAN, and DR-DSN models, and quote results for other methods from [11–13, 17, 23, 32, 35]. An in-depth ablation study is then provided to better understand of our DMASum.
Comparison with State-of-the-art Methods. Our main comparison with state-of-the-art methods is summarized in Table 1. The compared methods can be mainly categorized into LSTM, GAN, Attention, and meta learning models. M-AVS [12] and ACGAN [11] are based-on attention models and MetaL-TDVS [17] is based on meta learning. It can be seen that DMASum outperforms other approaches on both datasets consistently. The F1-score results can reflect that our attention mechanism with meta learning can better predict importance scores.

Table 1: F1-score (%) of DMASum with state-of-the-art approaches on both SumMe and TVSum dataset.

Method	SumMe	TVSum
DPP-LSTM [36]	38.6	54.7
SASUM [32]	45.3	58.2
SUM-GAN [21]	41.7	54.3
Cycle-SUM [35]	41.9	57.6
DR-DSN [39]	42.1	58.1
MetaL-TDVS [17]	44.1	58.2
ACGAN [11]	46.0	58.5
CSNet [13]	51.3	58.8
M-AVS [12]	44.4	61.0
DMASum	**54.3**	**61.4**

Table 2: Rank-order correlation coefficients computed between predicted importance scores by different models and human-annotated scores on both SumMe and TVSum datasets using Kendall's τ and Spearman's ρ correlation coefficients.

Method	SumMe		TVSum	
	τ	ρ	τ	ρ
Random	0.000	0.000	0.000	0.000
DPP-LSTM [36]	-	-	0.042	0.055
SUM-GAN [21]	0.049	0.066	0.024	0.031
DR-DSN [39]	0.028	-0.027	0.020	0.026
Human	**0.227**	**0.239**	0.178	0.205
DMASum	0.063	0.089	**0.203**	**0.267**

We also evaluate our DMASum by using the most recent rank-order statistics [23]. The new evaluation matrix can also consider the frame dependencies and annotator consistency so as to reflect the true importance better. Because, F1 score can partially reflect the consistency between prediction and importance scores due to large variations in segment length (i.e. two-peak, KTS, and randomized KTS). The correlation coefficients (Kendall's τ and Spearman's ρ) can be used to measure the similarity between the implicit ranking provided by the frame-level importance score of the generated frame annotation and the human annotation. From Table 2, We can see the correlation coefficients given by DMASum are significantly higher than other state-of-the-art models. More importantly, the performance on the TVSum dataset (0.233 and 0.267) is even better than human annotators (0.205 and 0.267). We believe it is because the dual-channel attention mechanism itself is simulating human behavior and memorizing visual and sequential sources of information and the meta learning method could learn the latent mechanism of summarizing a video story. Also, different human annotators might pay different attention to the given video. Our model can summarize the information from multiple human annotators so that the learned attention-based model is moderated and can achieve better consistency. Figure 5 demonstrates two examples of the correlation coefficients. Curves above the random importance scores in the black dash are positive with better consistency. Our
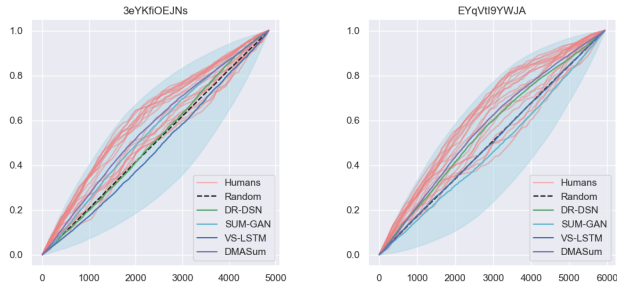
Figure 5: Example correlation curves produced for two videos from the TVSum dataset (3eYKfiOEJNs and EYqVtl9YWJA are video ids). The red lines represent correlation curves for 25 human annotators and the black dashed line is the expectation for a random importance score. The magenta curve shows the corresponding result.

Table 3: F1-score (%) of ablation study on SumMe and TV-Sum datasets. There are five ablation models: $DMASum_{wom}$ (without meta learning strategy), $DMASum_{softmax}$ (with standard softmax function in self-attention network), $DMASum_v$ (without sequential channel), $DMASum_s$ (without visual channel), $DMASum_b$ (with multiple videos in a batch), and $DMASum_{maml}$ (with MAML)

Method	SumMe	TVSum
$DMASum_{wom}$	51.6	60.6
$DMASum_{softmax}$	50.6	60.1
$DMASum_v$	53.2	60.5
$DMASum_s$	53.3	61.0
$DMASum_b$	51.3	60.0
$DMASum_{maml}$	49.3	59.2
DMASum	**54.3**	**61.4**

model achieves averaged performance among all human annotators and outperforms the other compared methods.

4.3 Ablation study.

The success of our DMASum ascribes to both the framework design and technical improvement in each module. To analyze the effect of each component in DMASum, we conduct six ablation study models including DMASum without single video meta learning ($DMASum_{wom}$), DMASum with standard softmax function in self-attention network ($DMASum_{softmax}$), DMASum without sequential channel ($DMASum_v$), DMASum without visual channel ($DMASum_s$), $DMASum_b$ is developed with the batch version of Reptile, and DMASum is designed with MAML ($DMASum_{maml}$). Results are summarised in Table 3, from which we can understand the following questions.

The Basis of Self-attention Architecture provided the initial performance boost. By removing our meta learning module, we can make a straight comparison with state-of-the-art DPP-LSTM [36] and M-AVS [12] which are using no attention and normal attention
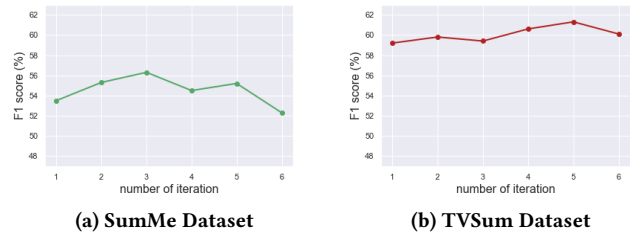

(a) SumMe Dataset   (b) TVSum Dataset

Figure 6: Different recurrent training Learner number with respect to the F1-score (%) in DMASum on both SumMe and TVSum datasets.

modules. Averaged improvement is around 5% to 10%. Note that M-AVS is slightly better than our method on the TVSum dataset due to their extra autoencoder architecture.

The Softmax Bottleneck problem results in severe performance gaps. By replace the MoA back to traditional softmax function, the performance drops 3.7% and 1.3% respectively on the two datasets. A more detailed analysis has been discussed in Section 3, from where we can see the problem is more critical when video contents are long and complex, involving rich sequential information.

Visual vs Sequential Representation. By comparing the performance of $DMASum_v$ or $DMASum_s$, we can observe that: 1) In TV-Sum dataset, the $DMASum_v$ gained a slightly better performance than $DMASum_s$. 2) The performance in SumMe dataset benefits more from the sequential channel. The self-attention network can effectively connect visual features from frames and the sequential information for the whole story line and thus our combined DMASum achieves better results.

The Necessity of Meta Learning. Removing the meta learning can heavily affect the performance by 2.7% on SumMe dataset. The key reason is that SumMe is a relatively small dataset. This observation serves as strong evidence to validate the motivation and necessity of our meta learning module.

MAML, Batch, and Single Video Meta Learning. The Single Video rule is the key finding that distinguish it from meta learning in other applications, e.g. few-shot learning. This is due to a video itself is rich and complex. By increasing each meta learning task from one video to three in a batch, the performance of $DMASum_b$ drops 3% and 1.4% with the clearly slowed training process. In addition, we can see that the performance of our proposed meta learning strategy is better than the batch version of the Reptile strategy, and the batch version of the Reptile strategy is time-consuming during the training process. The efficiency of Single-Video rule is also validated by comparing it to $DMASum_{maml}$.

Number of Recurrent Learning. In a controlled experiment, we observe that when the recurrent training Learner number is 3 for SumMe Dataset and 5 for the TVSum dataset, the F-score reaches the highest shown from Figure 6. Which means, the Learner might not learn the summarizing mechanism when the number is too low, and when the number is too high, the Learner might overfit the current video. In this paper, the number of recurrent training is automatically chosen by using the standard 5-fold cross validation.

(a) Example frames from video 37 in TVSum (indexed by [25])

(b) DMASum

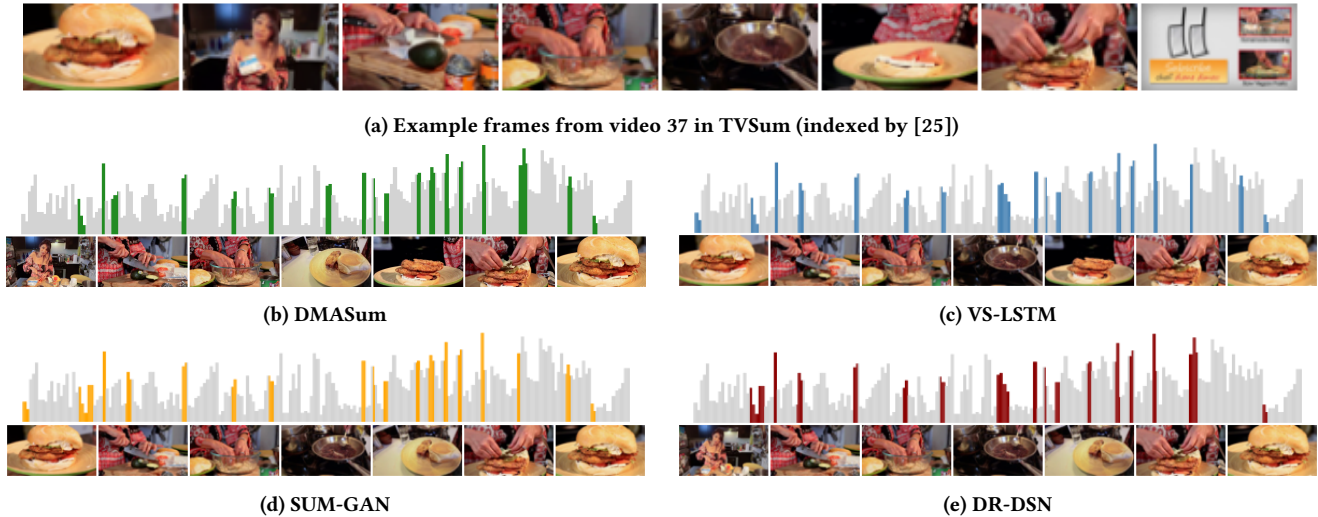(c) VS-LSTM

(d) SUM-GAN

(e) DR-DSN

Figure 7: Quantitative results of different approaches for video 16 in TVSum. In (b) to (e), the light-gray bars represent the ground truth importance scores, and the colored bars correspond to the selected frames by different methods.

Table 4: F-score (%) of approaches in canonical, augmented and transfer settings on SumMe and TVSum datasets.

Method	SumMe			TVSum		
	C	A	T	C	A	T
DPP-LSTM [36]	38.6	42.9	40.7	54.7	59.6	58.7
SUM-GAN [21]	41.7	43.6	-	54.3	**61.2**	-
DR-DSN [39]	42.1	43.9	42.6	58.1	59.8	58.9
CSNet [13]	51.3	52.1	45.1	58.8	59.0	59.2
DMASum	**54.3**	**54.1**	**52.2**	**61.4**	**61.2**	**60.5**

Comparison under Different Settings. Another approach to examining the model generalization is to investigate its performance under different task settings. Table 4 shows the experimental results of the comparison between the DMASum and cited results of state-of-the-art approaches in canonical, augmented and transfer settings. Note that even though the performance of our model in augmented and transfer settings are partially better than the best results. We observe that the given importance scores in Youtube and OVP datasets are either 0 or 1. However, the DMASum is learning by the importance scores within the range of zero to one from SumMe and TVSum datasets. Such discrepancy of importance score format in both Youtube and OVP datasets would cause the meta learning strategy to be ineffective or even counterproductive because our model is not tailored to handle the discrepancy in labels. Thus in the future, we can improve our framework to adapt to this situation. But on the positive side, our DMASum is still capable in both augmented and transfer settings and achieves comparable results to that of state-of-the-art models despite the above difficulties.

4.4 Qualitative Evaluation

To better illustrate the important frames selection of different approaches, we provide qualitative results for an exemplary video in

Figure 7, which tells a story of how to cook a burger. Overall, we can observe that all summaries generated by the different models can cover the intervals with high importance scores. Moreover, according to the figure, the summaries produced by both our DMA-Sum and SUM-GAN contain more peaks, which proves that our proposed model can effectively capture key-frames from the original video. Also, the summary of our model is more sparse and much closer to the entire storyline, i.e., the different cooking stages, which means our meta learning strategy can learn the latent mechanism of summarizing a video.

5 CONCLUSION

We have presented the first work to introduce self-attention meta learning architecture to estimate the visual and sequential attentions jointly for video summarization. The self-attention formula was derived into a matrix factorization problem and key technical Softmax Bottleneck has been identified with both theoretical and empirical evidences. Our work also confirmed the importance of high-rank representation for video summarization tasks. A novel MoA module was proposed to replace the softmax, which can compare twice by query-key and self-query attentions. The Single-Video Meta Learning rule was designed and particularly tailored for video summarization tasks and significantly improved off-the-shelf Meta Learning, e.g. MAML. On two public datasets, our DMASum outperforms other methods in terms of both F1-score and achieved human-level performance using rank-order correlation coefficients. Future work could focus on further improve the generalisation for cross-dataset settings using an integrated framework.

REFERENCES

[1] Jimmy Lei Ba, Jamie Ryan Kiros, and Geoffrey E Hinton. 2016. Layer normalization. *arXiv preprint arXiv:1607.06450* (2016).

[2] Dzmitry Bahdanau, Kyunghyun Cho, and Yoshua Bengio. 2014. Neural machine translation by jointly learning to align and translate. *arXiv preprint arXiv:1409.0473* (2014).

[3] Sandra Eliza Fontes De Avila, Ana Paula Brandão Lopes, Antonio da Luz Jr, and Arnaldo de Albuquerque Araújo. 2011. VSUMM: A mechanism designed to produce static video summaries and a novel evaluation method. *Pattern Recognition Letters* 32, 1 (2011), 56–68.

[4] Debidatta Dwibedi, Yusuf Aytar, Jonathan Tompson, Pierre Sermanet, and Andrew Zisserman. 2019. Temporal cycle-consistency learning. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*. 1801–1810.

[5] Christoph Feichtenhofer, Haoqi Fan, Jitendra Malik, and Kaiming He. 2019. Slow-fast networks for video recognition. In *Proceedings of the IEEE international conference on computer vision*. 6202–6211.

[6] Chelsea Finn, Pieter Abbeel, and Sergey Levine. 2017. Model-agnostic meta-learning for fast adaptation of deep networks. In *Proceedings of the 34th International Conference on Machine Learning-Volume 70*. JMLR. org, 1126–1135.

[7] Boqing Gong, Wei-Lun Chao, Kristen Grauman, and Fei Sha. 2014. Diverse sequential subset selection for supervised video summarization. In *Advances in Neural Information Processing Systems*. 2069–2077.

[8] Michael Gygli, Helmut Grabner, Hayko Riemenschneider, and Luc Van Gool. 2014. Creating summaries from user videos. In *European conference on computer vision*. Springer, 505–520.

[9] Michael Gygli, Helmut Grabner, and Luc Van Gool. 2015. Video summarization by learning submodular mixtures of objectives. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*. 3090–3098.

[10] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. 2016. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*. 770–778.

[11] Xufeng He, Yang Hua, Tao Song, Zongpu Zhang, Zhengui Xue, Ruhui Ma, Neil Robertson, and Haibing Guan. 2019. Unsupervised Video Summarization with Attentive Conditional Generative Adversarial Networks. In *Proceedings of the 27th ACM International Conference on Multimedia*. 2296–2304.

[12] Zhong Ji, Kailin Xiong, Yanwei Pang, and Xuelong Li. 2019. Video summarization with attention-based encoder-decoder networks. *IEEE Transactions on Circuits and Systems for Video Technology* (2019).

[13] Yunjae Jung, Donghyeon Cho, Dahun Kim, Sanghyun Woo, and In So Kweon. 2019. Discriminative Feature Learning for Unsupervised Video Summarization. In *Proceedings of the AAAI Conference on Artificial Intelligence*, Vol. 33. 8537–8544.

[14] Maurice G Kendall. 1945. The treatment of ties in ranking problems. *Biometrika* 33, 3 (1945), 239–251.

[15] Alex Kulesza, Ben Taskar, et al. 2012. Determinantal point processes for machine learning. *Foundations and Trends® in Machine Learning* 5, 2–3 (2012), 123–286.

[16] Yong Jae Lee, Joydeep Ghosh, and Kristen Grauman. 2012. Discovering important people and objects for egocentric video summarization. In *2012 IEEE Conference on Computer Vision and Pattern Recognition*. IEEE, 1346–1353.

[17] Xuelong Li, Hongli Li, and Yongsheng Dong. 2019. Meta Learning for Task-Driven Video Summarization. *IEEE Transactions on Industrial Electronics* (2019).

[18] Xiangpeng Li, Jingkuan Song, Lianli Gao, Xianglong Liu, Wenbing Huang, Xiangnan He, and Chuang Gan. 2019. Beyond RNNs: Positional Self-Attention with Co-Attention for Video Question Answering. In *The 33rd AAAI Conference on Artificial Intelligence*, Vol. 8.

[19] David Liu, Gang Hua, and Tsuhan Chen. 2010. A hierarchical visual model for video object summarization. *IEEE transactions on pattern analysis and machine intelligence* 32, 12 (2010), 2178–2190.

[20] Zheng Lu and Kristen Grauman. 2013. Story-driven summarization for egocentric video. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*. 2714–2721.

[21] Behrooz Mahasseni, Michael Lam, and Sinisa Todorovic. 2017. Unsupervised video summarization with adversarial lstm networks. In *Proceedings of the IEEE conference on Computer Vision and Pattern Recognition*. 202–211.

[22] Alex Nichol, Joshua Achiam, and John Schulman. 2018. On first-order meta-learning algorithms. *arXiv preprint arXiv:1803.02999* (2018).

[23] Mayu Otani, Yuta Nakashima, Esa Rahtu, and Janne Heikkila. 2019. Rethinking the Evaluation of Video Summaries. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*. 7596–7604.

[24] Danila Potapov, Matthijs Douze, Zaid Harchaoui, and Cordelia Schmid. 2014. Category-specific video summarization. In *European conference on computer vision*. Springer, 540–555.

[25] Yale Song, Jordi Vallmitjana, Amanda Stent, and Alejandro Jaimes. 2015. Tvsum: Summarizing web videos using titles. In *Proceedings of the IEEE conference on computer vision and pattern recognition*. 5179–5187.

[26] Christian Szegedy, Wei Liu, Yangqing Jia, Pierre Sermanet, Scott Reed, Dragomir Anguelov, Dumitru Erhan, Vincent Vanhoucke, and Andrew Rabinovich. 2015. Going deeper with convolutions. In *Proceedings of the IEEE conference on computer vision and pattern recognition*. 1–9.

[27] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. 2017. Attention is all you need. In *Advances in neural information processing systems*. 5998–6008.

[28] Junyan Wang, Bingzhang Hu, Yang Long, and Yu Guan. 2019. Order Matters: Shuffling Sequence Generation for Video Prediction. In *Proc. BMVA British Mach. Vis. Conf.* 275.1–275.14.

[29] Qiang Wang, Li Zhang, Luca Bertinetto, Weiming Hu, and Philip HS Torr. 2019. Fast online object tracking and segmentation: A unifying approach. In *Proceedings of the IEEE conference on computer vision and pattern recognition*. 1328–1338.

[30] Xi Wang, Yu-Gang Jiang, Zhenhua Chai, Zichen Gu, Xinyu Du, and Dong Wang. 2014. Real-Time Summarization of User-Generated Videos Based on Semantic Recognition. In *Proceedings of the 22nd ACM International Conference on Multimedia* (Orlando, Florida, USA) *(MM âĂŽ14)*. Association for Computing Machinery, New York, NY, USA, 849âĂŞ852. https://doi.org/10.1145/2647868.2655013

[31] Yu-Xiong Wang and Martial Hebert. 2016. Learning to learn: Model regression networks for easy small sample learning. In *European Conference on Computer Vision*. Springer, 616–634.

[32] Huawei Wei, Bingbing Ni, Yichao Yan, Huanyu Yu, Xiaokang Yang, and Chen Yao. 2018. Video summarization via semantic attended networks. In *Thirty-Second AAAI Conference on Artificial Intelligence*.

[33] Chao-Yuan Wu, Christoph Feichtenhofer, Haoqi Fan, Kaiming He, Philipp Krahenbuhl, and Ross Girshick. 2019. Long-term feature banks for detailed video understanding. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*. 284–293.

[34] Zhilin Yang, Zihang Dai, Ruslan Salakhutdinov, and William W Cohen. 2017. Breaking the softmax bottleneck: A high-rank RNN language model. *arXiv preprint arXiv:1711.03953* (2017).

[35] Li Yuan, Francis EH Tay, Ping Li, Li Zhou, and Jiashi Feng. 2019. Cycle-SUM: Cycle-consistent Adversarial LSTM Networks for Unsupervised Video Summarization. *arXiv preprint arXiv:1904.08265* (2019).

[36] Ke Zhang, Wei-Lun Chao, Fei Sha, and Kristen Grauman. 2016. Video summarization with long short-term memory. In *European conference on computer vision*. Springer, 766–782.

[37] Ke Zhang, Kristen Grauman, and Fei Sha. 2018. Retrospective encoders for video summarization. In *Proceedings of the European Conference on Computer Vision (ECCV)*. 383–399.

[38] Bin Zhao, Xuelong Li, and Xiaoqiang Lu. 2018. Hsa-rnn: Hierarchical structure-adaptive rnn for video summarization. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*. 7405–7414.

[39] Kaiyang Zhou, Yu Qiao, and Tao Xiang. 2018. Deep reinforcement learning for unsupervised video summarization with diversity-representativeness reward. In *Thirty-Second AAAI Conference on Artificial Intelligence*.

[40] Luowei Zhou, Yingbo Zhou, Jason J Corso, Richard Socher, and Caiming Xiong. 2018. End-to-end dense video captioning with masked transformer. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*. 8739–8748.

[41] Yi Zhu, Yang Long, Yu Guan, Shawn Newsam, and Ling Shao. 2018. Towards universal representation for unseen action recognition. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*. 9436–9445.

[42] Daniel Zwillinger and Stephen Kokoska. 1999. *CRC standard probability and statistics tables and formulae*. Crc Press.