

Multi-modal Segment Assemblage Network for Ad Video Editing with Importance-Coherence Reward

Yunlong Tang^{1,2}[0000–0003–2796–1787], Siting Xu¹[0000–0001–9934–7919],
Teng Wang^{1,2}, Qin Lin², Qinglin Lu², and Feng Zheng¹[0000–0002–1701–9141]*

¹ Southern University of Science and Technology, Shenzhen, China

² Tencent Inc., Shenzhen, China

{tangyl2019, xust2019, wangt2020}@mail.sustech.edu.cn

{angelqlin, qinglinlu}@tencent.com f.zheng@ieee.org

Abstract. Advertisement video editing aims to automatically edit advertising videos into shorter videos while retaining coherent content and crucial information conveyed by advertisers. It mainly contains two stages: video segmentation and segment assemblage. The existing method performs well at video segmentation stages but suffers from the problems of dependencies on extra cumbersome models and poor performance at the segment assemblage stage. To address these problems, we propose M-SAN (Multi-modal Segment Assemblage Network) which can perform efficient and coherent segment assemblage task end-to-end. It utilizes multi-modal representation extracted from the segments and follows the Encoder-Decoder Ptr-Net framework with the Attention mechanism. Importance-coherence reward is designed for training M-SAN. We experiment on the Ads-1k dataset with 1000+ videos under rich ad scenarios collected from advertisers. To evaluate the methods, we propose a unified metric, Imp-Coh@Time, which comprehensively assesses the importance, coherence, and duration of the outputs at the same time. Experimental results show that our method achieves better performance than random selection and the previous method on the metric. Ablation experiments further verify that multi-modal representation and importance-coherence reward significantly improve the performance. Ads-1k dataset is available at: <https://github.com/yunlong10/Ads-1k>.

Keywords: Ad Video Editing · Segment Assemblage · Advertisement Dataset · Multi-modal · Video Segmentation · Video Summarization.

1 Introduction

With the boom of the online video industry, video advertising has become popular with advertisers. However, different online video platforms have different requirements for the content and duration of ad videos. It is time-consuming and laborious for advertisers to edit their ad videos into a variety of duration

* Corresponding author

tailored to the diverse requirements, during which they have to consider which part is important and whether the result is coherent. Therefore, it is of great importance to automatically edit the ad videos to meet the requirements of duration, and the edited videos should be coherent and retain informative content.

Ad video editing is a task aiming to edit an ad video into its shorter version to meet the duration requirements, ensuring coherence and avoiding losing important ad-related information. Video segmentation and segment assemblage are the two main stages in ad video editing task [20], as Fig. 1 shows. An ad video will be cut into several segments with a small duration during the video segmentation stage. At the segment assemblage stage, the output will be produced by selecting and assembling a subset of the input segments of the source ad video. The key to video segmentation is to preserve the local semantic integrity of each video segment. For instance, a complete sentence of a speech or caption in the source video should not be split into two video segments. Existing method [20] has achieved this by aligning shots, subtitles, and sentences to form the segments. However, at the segment assemblage stage, the only pioneer work [20] suffers the following problems: (1) To calculate the individual importance of each segment and the coherence between segments, extra models are required to perform video classification [21, 14] and text coherence prediction [7], which is inefficient during inference. (2) Without globally modeling the context of videos, graph-based search adopted by [20] produces results with incoherent segments or irrelevant details.

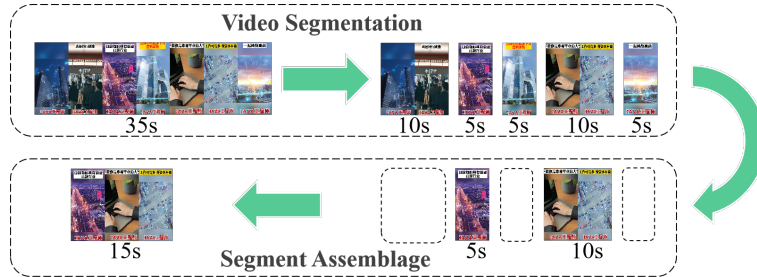


Fig. 1. The two stages of ad video editing: video segmentation and segment assemblage.

To tackle these problems, we propose an end-to-end Multi-modal Segment Assemblage Network (M-SAN) for accurate and efficient segment assemblage. It is free of extra cumbersome models during inference and strikes a better balance between importance and coherence. Specifically, we obtain segments at the video segmentation stage by boundary detection and alignment. Different from daily life videos, ad videos usually have sufficient multi-modal content like speech and caption, which contain abundant video semantics [13]. Therefore, pretrained unimodal models are applied to respectively extract the representation of shots, audios, and sentences, which are concatenated together yielding a multi-modal representation of segments. During the assemblage stage, we adopt a pointer

network with RNN-based decoder to improve the temporal dependency between selected segments. Importance-coherence reward is designed for training M-SAN with Policy Gradient [35]. Importance reward measures the amount of important ad-related information contained in the output. Coherence reward measures the text coherence between every two adjacent selected segments, which is computed as the mean of PPL (perplexity) [30] of sentences generated by concatenating the two texts extracted from adjacent selected segments.

To evaluate our methods, we propose the new metric, Imp-Coh@Time, which takes the importance, coherence, and duration of the outputs into consideration at the same time instead of evaluating importance or coherence respectively. We experiment on the Ads-1k dataset with 1000+ ad videos collected from advertisers. Experimental results show that our method achieves better performance than random selection and the previous method [20] on the metric. Ablation experiments further verify that multi-modal representation and importance-coherence reward significantly improve the performance.

Our work mainly focuses on segment assemblage in ad video editing, and its main contributions can be summarized as follows:

- We propose M-SAN to perform segment assemblage efficiently and improve the result of ad video editing, without relying on an extra model when inference.
- We propose importance-coherence reward and train M-SAN with policy gradient to achieve a better trade-off between importance and coherence.
- We collect the dataset Ads-1k with 1000+ ad videos and propose Imp-Coh@Time metrics to evaluate the performance of ad video editing methods. Our M-SAN achieves state of the art on the metrics.

2 Related Work

2.1 Video Editing

There are three main categories of automated video editing [20]. They're video summarization, video highlight detection, and task-specific automated editing.

Video Summarization. The most relevant task to ad video editing is video summarization. It is a process that extracts meaningful shots or frames from video by analyzing structures of the video and time-space redundancy in an automatic or semi-automatic way. To perform video summarization, a load of work focuses on supervised learning based on frames [11, 34, 8, 23], shots [9, 39, 16], and clips [19]. Other than these works, DSN [40] is the first proposed unsupervised video summarization model training with diversity-representativeness reward by policy gradient. Without utilizing annotation, the method reached fully unsupervised. Our rewards design mainly refers to [40].

Video Highlight Detection. Learning how to extract an important segment from videos is the main motivation we focus on video highlight detection. [31] proposed frameworks that exploit users' previously created history. Edited videos

created by users are utilized in [38] to achieve highlight detection in an unsupervised way since human-edited videos tend to show more interesting or important scenes. [36] presents an idea that shorter videos tend to be more likely to be selected as highlights. Combining the above ideas, we exploit the ability of MSVM [20] which extracts potential selling points in segments that tend to be of higher importance when performing assemblage.

Task-specific Automated Editing. Videos can be presented in various forms such as movies, advertisements, sports videos, etc. To extract a short video from a long video also exists in specific scenarios of movies [15, 2, 37]. [15] pointed out that a movie has the trait that its computational cost is high. Also, sports videos have been explored to extract highlights [25, 26] for sports having the characteristic that they are high excitement. Advertisements are rich in content and they vary in duration among different platforms [20].

2.2 Text Coherence Prediction and Evaluation

Text information is extracted in the stage of video segmentation. When assembling segments, texts that are concatenated act as a reference to coherence. [4] proposed narrative incoherence detection, denoted semantic discrepancy exists causes incoherence. In [20], next sentence prediction (NSP) [7] is exploited to assess coherence. Although our work is inspired by their thoughts, we use perplexity as the coherence reward and metric to evaluate sentence coherence.

2.3 Neural Combinatorial Optimization

Combinatorial optimization problem is a problem that gets extremum in discrete states. Common problems like Knapsack Problem (KP), Travelling Salesman Problem (TSP), and Vehicular Routing Problem (VRP) belong to combinatorial optimization problem. Pointer Network (Ptr-Net) is proposed in [33] and performs better than heuristic algorithm in solving TSP. Later, Ptr-Net has been exploited with reinforcement learning [3, 29, 6, 18, 12]. In [10], Ptr-Net is used to solve the length-inconsistency problem in video summarization. In our network, Ptr-Net is used to model the video context and select the tokens from the input sequence as output.

3 Method

3.1 Problem Formulation

Given a set of N segments of ad video $S = \{s_i\}_{1 \leq i \leq M}$, our goal is to select a subset $A = \{a_i\}_{1 \leq i \leq N} \subseteq S$ which can be combined into the output video so that it can take the most chance to retain the important information and be coherent as well as meeting the requirements of duration. We denote the ad importance of the segment a_i as $imp(a_i)$, the coherence as $coh(a_i, a_j)$ and the duration as

$dur(a_i)$. Overall the task of segment assemblage can be regarded as a constrained combinatorial optimization problem and is defined formally as follows:

$$\begin{aligned} \max_{A \subseteq S} \quad & \sum_{a_i \in A} imp(a_i) + \sum_{a_i \prec a_j} coh(a_i, a_j) , \\ s.t. \quad & T_{min} \leq \tau(A) \leq T_{max}, \text{ and } \forall a_i, a_j \in A, a_i \neq a_j , \end{aligned} \quad (1)$$

where

- T_{min} and T_{max} are the lower and upper bound of requirement duration,
- $\tau(A) = \sum_{a_i \in A} dur(a_i)$,
- $a_i \prec a_j$ is defined as $(a_i, a_j \in A) \wedge (i < j) \wedge (\forall a_k \in A, k < i \vee k > j)$.

This is an NP-hard problem that can not be solved in polynomial time. Instead of utilizing graph modeling and optimization [20] to search for an optimal solution, we adopt a neural network with pointer [33] that follows the framework of neural combinatorial optimization to optimize the objective directly.

3.2 Architecture

The architecture of M-SAN is shown as Fig. 2. It incorporates a multi-modal video segmentation module [20] (MVSM), multi-modal representation extraction module (MREM) and assemblage module (AM). To preserve the local semantic integrity of each segment, we adopt MVSM to obtain video segments with reasonable boundaries. With ASR and OCR, MVSM also captures the texts from each segment. Given the segments and the corresponding texts, MREM extracts the segment-level representations of shots, audios, and texts, which are further jointed into the multi-modal representations. AM utilizes these representations to model the context of video and make decisions by a pointer network [33].

Video Segmentation Module. At the video segmentation stage, we first apply MVSM [20] to obtain the segments of each input video. MVSM splits a video into the video track and audio track and extracts shots $\{v_i\}$, audios $\{\alpha_i\}$, ASR and OCR results $\Omega_i = \{\omega_i\}$ (a sentence with words ω_i) from video to generate the boundaries of the content in each modality. The boundaries of audio space and textual space are first merged to form the joint space, followed by merging the boundaries of visual space and joint space to yield the final segments set $\{s_i\}$. The segment $s_i = (\{v_{p_0}, \dots, v_{q_0}\}, \{\alpha_{p_1}, \dots, \alpha_{q_1}\}, \Omega_i)$ preserves the integrity of local atomic semantic, where $p(\cdot) < q(\cdot)$.

Multi-modal Representation Extraction Module. MREM integrates three kinds of representation extractors: pre-trained Swin-Transformer [24], Vggish [14] and BERT [7] models. The Swin-Transformer extracts the visual representations $\{\tilde{v}_{p_0}, \dots, \tilde{v}_{q_0}\}$ from shots $\{v_{p_0}, \dots, v_{q_0}\}$ in each segment $\{s_i\}$, and a segment-level visual representation is computed as the mean of single shot-level representation. Vggish and BERT model extract audio representations $\{\tilde{\alpha}_{p_1}, \dots, \tilde{\alpha}_{q_1}\}$ and




Fig. 2. The Architecture of M-SAN

text representations $\tilde{\Omega}_i$, respectively. Similarly, a segment-level audio representation $\tilde{\alpha}_i$ is given by the mean of $\{\tilde{\alpha}_{p_1}, \dots, \tilde{\alpha}_{q_1}\}$. Since each segment contains at most one sentence, it is $\tilde{\Omega}_i$ that is the segment-level text representation. These three modalities will be jointed by concatenating directly to yield the final multi-modal representation of segment $\tilde{s}_i = [\tilde{v}_i, \tilde{\alpha}_i, \tilde{\Omega}_i]^T$.

Assemblage Module. The output of an Encoder-Decoder Ptr-Net [33] is produced by iteratively copying an input item that is chosen by the pointer [27], which is quite suitable for segment assemblage task. Therefore, our assemblage module (AM) follows this framework.

The encoder integrates a linear embedding layer and a bi-directional GRU. To enhance interactions between modalities, the linear embedding layer will perform preliminary a fusion of the three modalities and produce embedding token x_i corresponding to segment s_i . To enhance the interaction between segments and model the context of the whole video, a Bi-GRU is adopted to further embed the tokens:

$$H = GRU_e(X) , \quad (2)$$

where $X = [x_1, \dots, x_M]$, $H = [h_1, \dots, h_M]$, and the hidden state h_i is the context embedding for segment s_i .

Given the output of encoder H and X , the GRU [5] decoder with Attention mechanism [1] predicts the probability distribution of segment to be selected from S at every time-step t to get the result A :

$$p_\theta(A|S) = \prod_{t=1}^N p_\theta(a_t|a_{1:t}, S) = \prod_{t=1}^N p_\theta(a_t|a_{1:t-1}, H, X) = \prod_{t=1}^N p_\theta(a_t|a_{t-1}, d_t) , \quad (3)$$

where θ is the learnable parameter, and d_t is the hidden state computed by the decoder at time-step t . With d_t as the query vector, the decode will glimpse [32] the whole output of encoder H to compute the bilinear attention. Instead of utilizing the additive attention adopted in [10, 12], we compute bilinear attention μ_t with less computational cost:

$$\mu_t = \text{Softmax}(H^T W_{att1} d_t) . \quad (4)$$

Until now, the attention μ_t is used as the probability distribution to guide the selection in most of Ptr-Net framework [10, 12]. To dynamically integrate information over the whole video [28], we further calculate the context vector c_t of encoder output and update query vector d_t to \tilde{d}_t by concatenating it with c_t (Eq. 5). Then we compute attention a second time to obtain the probability distribution of segment selection at current time-step t (Eq. 6).

$$c_t = H\mu_t = \sum_{m=1}^M \mu_t^{(m)} h_m, \quad \tilde{d}_t = \begin{bmatrix} d_t \\ c_t \end{bmatrix} , \quad (5)$$

$$\tilde{\mu}_t = \text{Softmax}(M H^T W_{att2} \tilde{d}_t) , \quad (6)$$

$$p_\theta(a_t | a_{t-1}, d_t) = \tilde{\mu}_t , \quad (7)$$

where M can mask the position i corresponding to selected segment a_i to $-\infty$. The segment at this position will be not selected at later time-steps, since $\text{Softmax}(\cdot)$ modifies the corresponding probability to 0. Finally, the segment selected at time-step t will be sampled from the distribution:

$$a_t \sim p_\theta(a_t | a_{t-1}, d_t) . \quad (8)$$

If the sum of duration of selected segments $\tau(A)$ exceeds the tolerable duration limit $[T_{min}, T_{max}]$, the current and following segments selected will be replaced by [EOS] token [17].

3.3 Reward Design

Importance Reward. To extract important parts from original ad videos, rewards related to the importance of selected segments should be fed back to the network during training. We design importance reward R_{imp} :

$$R_{imp} = \frac{1}{|A|} \sum_{a_i \in A} imp(a_i) , \quad (9)$$

$$imp(a_i) = \frac{1}{|L_{a_i}|} \sum_{\ell \in L_{a_i}} w_\ell \cdot \ell , \quad (10)$$

where A is the set of selected segments, L_{a_i} stands for the total number of labels of the selected segment, $\ell \in \{1, 2, 3, 4\}$ is narrative techniques label hierarchy,

and w_l is the weight of one label. According to Eq. 9, the importance of a single segment is the weighted average of its annotated labels' weight. The labels listed in the supplementary are divided into four groups with four levels ranging from 1 to 4 according to their ad-relevance. We compute the importance reward of one output as the mean of the importance of single segments.

Coherence Reward. Importance reward mainly focuses on local visual dynamics within the single segment, neglecting temporal relationships between adjacent segments. We introduce a linguistic coherence reward to improve the fluency of caption descriptions of two segments.

Specifically, the extracted texts from adjacent segments are combined in pairs while retaining the original order, that is, preserving the same order of the original precedence. Then we compute the perplexity [30] (PPL) for each combined sentence by GPT-2 pre-trained on 5.4M advertising texts:

$$\begin{aligned} PPL(\Omega_1, \Omega_2) &= p(\omega_1^{(1)}, \omega_2^{(1)}, \dots, \omega_m^{(1)}, \omega_1^{(2)}, \omega_2^{(2)}, \dots, \omega_n^{(2)})^{-1/(m+n)} \\ &= \sqrt[m+n]{\prod_{i=1}^{m+n} \frac{1}{p(\omega_i | \omega_1, \omega_2, \dots, \omega_{i-1})}}, \end{aligned} \quad (11)$$

where $\Omega_i = (\omega_1^{(i)}, \dots, \omega_m^{(i)})$ is one sentence with m words. PPL reflects the *incoherence* of a sentence. We also maintain a PPL map to store the PPL of sentences in pairs as Fig. 3 shown. The grey parts represent invalid PPL values,

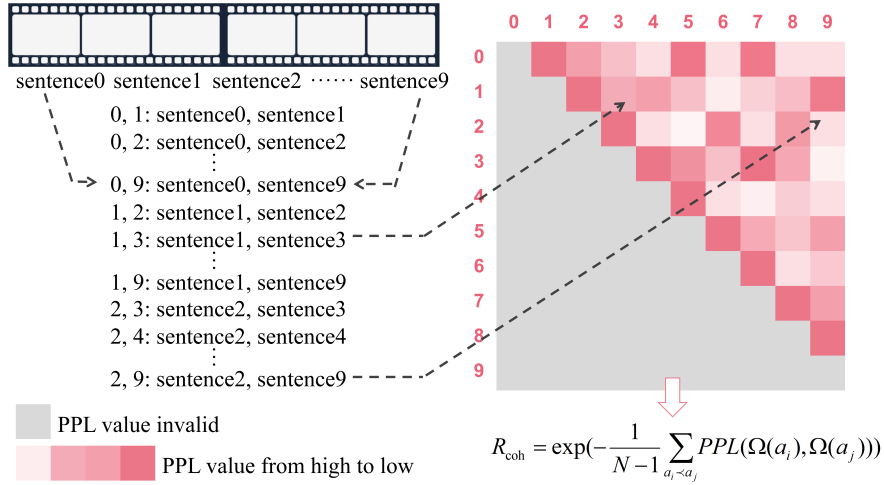


Fig. 3. PPL map and coherence reward

which means sentence pairs that violate the original orders have no valid PPL

value. The rest means valid PPL values vary from high to low as the color of red deepens. Since the smaller the PPL, the better the text coherence, the coherence reward is computed with a transfer function:

$$R_{coh} = \exp\left(-\frac{1}{N-1} \sum_{a_i \prec a_j} PPL(\Omega(a_i), \Omega(a_j))\right), \quad (12)$$

where $N = |A|$ is the total number of selected segments, $\Omega(a_i)$ is the sentence recognized from segment a_i . The $\exp(\cdot)$ ensures the R_{coh} at the same order of magnitude with R_{imp} .

Importance-Coherence Reward. To balance the importance and coherence of the selection, we make R_{imp} and R_{coh} complement each other and jointly guide the learning of M-SAN:

$$R = \beta \cdot R_{imp} + (1 - \beta) \cdot R_{coh}, \quad (13)$$

where the coefficient of reward β is a hyperparameter.

3.4 Training

Policy gradient (a.k.a. REINFORCE algorithm [35]) is adopted during training:

$$\nabla_{\theta} J(\theta) \approx \frac{1}{K} \sum_{k=1}^K \sum_{t=1}^N (R(A^{(k)}) - b^{(k)}) \nabla_{\theta} \log p_{\theta}(a_t^{(k)} | a_{t-1}^{(k)}, d_t^{(k)}), \quad (14)$$

where K denotes the number of episodes, a_t stands for the actions (which segment to choose) and d_t is the hidden state estimated by the decoder. $R(A)$ is the reward calculated by Eq.13. A baseline value b so that the variance can be reduced. For the optimization, neural network parameter θ is updated as:

$$\theta = \theta - \eta \nabla_{\theta} (-J(\theta)), \quad (15)$$

where η is the learning rate.

4 Experiments

4.1 Dataset

In [20], there are only 50 ad videos used in experiments. To obtain better training and evaluation, we collect 1000+ ad videos from the advertisers to form the Ads-1k dataset. There are 942 ad videos for training and 99 for evaluation in total. However, the annotation methods of the training set and test set are somehow different. Instead of preparing the ground-truth for each data, we annotate each video with multi-labels shown in the supplementary.

Table 1. Dataset statistics. N_{seg} and N_{label} are respectively the average number of segments and labels of each video. D_{seg} and D_{video} are the average duration of a segment and a video, respectively.

Dataset	N_{seg}	$D_{seg}(s)$	$D_{video}(s)$	N_{label}
Training Set	13.90	2.77	34.60	30.18
Test Set	18.81	1.88	34.21	35.77
Overall	14.37	2.68	35.17	30.71

We counted the average segment number for videos, the average length of segments in seconds, the average duration of each video in seconds, and the average label number for each video for the training set, test set, and the whole dataset respectively. The results are shown as Tbl. 1. Besides, the number of annotated segment pairs and the proportion are counted. The number of *coherent*, *incoherent*, and *uncertain* pairs are 6988, 9551, and 2971, occupying 36%, 49%, and 15%, respectively.

4.2 Metric

Imp@T. We define the score of ad importance given the target duration T as follows:

$$Imp@T = \frac{1}{|A|} \sum_{a_i \in A} imp(a_i) \cdot \mathbb{I}[c_1 \cdot T \leq \tau(A) \leq c_2 \cdot T], \quad (16)$$

where A is the set of selected segments, $imp(a_i)$ is defined by Eq. 10. $\mathbb{I}(\cdot)$ is the indicator function. c_1 and c_2 are two constant that produced a interval based on given target duration T . We set $c_1 = 0.8$ and $c_2 = 1.2$, since a post-processing of $0.8\times$ slow down $1.2\times$ or fast forward can resize the result close to the target T without distortion in practice. Take $T = 10$ as example, the interval will be $[8, 12]$, which means if the duration of result $\tau(A) = \sum_{a_i \in A} dur(a_i) \in [8, 12]$ then this result is valid and gain the score.

Coh@T. The coherence score given the target duration T is defined as follows:

$$Coh@T = \frac{1}{|A| - 1} \sum_{a_i \prec a_j} coh(a_i, a_j) \cdot \mathbb{I}[c_1 \cdot T \leq \tau(A) \leq c_2 \cdot T], \quad (17)$$

where the $coh(a_i, a_j)$ is the coherence small score between the text of segment i and the text of segment j . With the annotation for coherence on test set, we can score the results produced by our models. When scoring, for each combination of consecutive two segments i and j in output, if it is in the *coherent* set, the $coh(a_i, a_j)$ will be 1. If it is in the *incoherent* set, the $coh(a_i, a_j)$ will be 0. Otherwise, it is in the *uncertain* set, the $coh(a_i, a_j)$ will be 0.5.

Imp-Coh@T. The overall score is defined as follows:

$$ImpCoh@T = \frac{Imp@T}{|A| - 1} \cdot \sum_{a_i \prec a_j} coh(a_i, a_j), \quad (18)$$

where Imp@T is defined in Eq. 16. The score reflects the ability of trade-off among importance, coherence and total duration.

4.3 Implementation Details

Baselines. Besides SAM (Segments Assemblage Module) proposed in [20], we also adopt two random methods to perform the segment assemblage task, given the segments produced by MVSM [20].

- Given the set of input segments $S = \{s_i\}_{1 \leq i \leq M}$ and the target time T , *Random* will first produce a random integer $1 \leq r \leq M$. Then it will randomly pick up r segments from S to get the result $A = \{a_i\}_{1 \leq i \leq r}$, regardless of the target time T .
- Given $S = \{s_i\}_{1 \leq i \leq M}$ and T , *Random-Cut* randomly picks up segments from S and add to A until $c_1 \cdot T \leq \tau(A) \leq c_2 \cdot T$, ensuring satisfying the requirement of duration.
- Given $S = \{s_i\}_{1 \leq i \leq M}$ and T , the *SAM* [20] will utilize an extra model to perform video classification or named entity recognition to obtain some labels for each segment and compute an importance score for each segment. It also utilizes an extra BERT [7] to perform next sentence prediction (NSP) to compute a coherence score for each pair of segments. Then the segments $\{s_i\}_{1 \leq i \leq M}$ will be modeled as a graph with $|S|$ nodes and $|S|(|S| - 1)$ edges, where the weight of nodes are the importance score and the weight of edges are the coherence score. DFS with pruning is then adopted to search on the graph to collect a set that maximizes the sum of importance scores and coherence scores.

Parameters and Training Details. The Swin-Transformer [24] we used is the Large version with an output size of 1536. The 5 shots will be extracted to generate a visual representation every second of the video. The segment-level representation is computed as the mean of. The output size of BERT [7] and Vggish [14] are 768 and 128 respectively, and the sampling rate of Vggish is 5 every second to align with visual information. The dimension of \tilde{s}_i is 2432. The linear embedding layer performs a projection from 2432 to 768. We optimize the sum reward $R = 0.5 \cdot R_{imp} + 0.5 \cdot R_{coh}$, where the R_{imp} and R_{coh} are given in Eq. 9 and Eq. 12 with $\beta = 0.5$. The w_l in R_{imp} for all ℓ are 0.25. The optimizer we used is Adam [22]. More details can be found in supplementary.

4.4 Performance Comparison

We evaluated our M-SAN on the test set of Ads-1k with the three baselines mentioned above. The results are provided by Tbl. 2. It shows that our M-SAN

is state of the art on segment assemblage task given target duration $T = 10$ and $T = 15$. There is a significant improvement from Random to Random-Cut. Therefore, simply sticking to the time limit can improve performance by leaps and bounds. From Random-Cut to SAM, the improvement at $T = 10$ is also obvious, while the difference between them at $T = 15$ is not. This is probably because a longer duration budget forces Random-Cut to select more segments. The segment pairs as result have a greater chance of being coherent pairs.

Although SAM performs better than Random-Cut, its ability to trading-off between importance and coherence is still weak. M-SAN addresses the problem by trained with importance-coherence reward and achieves a better performance.

Table 2. Performance comparison results. Our M-SAN is state of the art on segment assemblage task.

	Imp-Coh@10			Imp-Coh@15		
	Imp	Coh	Overall	Imp	Coh	Overall
Random	10.68	12.57	7.92	17.04	20.49	12.71
Random-Cut	55.35	73.75	41.77	60.39	76.56	47.25
SAM [20]	72.54	86.55	65.97	63.97	78.58	58.09
M-SAN(ours)	80.29	92.16	74.19	77.00	90.83	70.28

4.5 Ablation Studies

Ablation of Modalities. We explore the effect of the representation from different modalities. Tbl. 3 shows that incorporating the text or audio representation can both improve the overall score while incorporating the former have much effect. After leveraging all representations from three modalities, the overall scores increase significantly, which has 8.7 and 3.75 gains compared with utilizing visual information only on Imp-Coh@10 and Imp-Coh@15, respectively. Even though adding text representation to video-audio dual modalities hurts the Imp@15, other scores increase obviously, which demonstrates the significance of multi-modal representation.

Ablation of Reward and Glimpse. To verify the effectiveness of reward and glimpse (two-stage attention calculation), we design the following ablation experiments with target duration $T = 10$ and $T = 15$. The results in the Tbl. 4 show that M-SAN gains a higher score on all metrics than the one without glimpse. Therefore, dynamically integrating information over the whole video by Glimpse can improve the performance. Similarly, we perform ablation on the rewards: importance-coherence reward (M-SAN), importance reward only, and coherence reward only. Results in Tbl. 4 show that M-SAN trained with

Table 3. Ablation study on modalities

Modalities			Imp-Coh@10			Imp-Coh@15		
V	T	A	Imp	Coh	Overall	Imp	Coh	Overall
✓			78.43	89.59	65.49	74.48	87.41	66.53
✓	✓		78.57	89.73	71.04	76.22	89.85	68.56
✓		✓	79.49	91.49	72.72	77.08	90.36	69.91
✓	✓	✓	80.29	92.16	74.19	77.00	90.83	70.28

importance-reward only gained relatively low scores comparing the other two kinds of rewards. The one trained with coherence-reward only gained a higher score than the one with importance-reward. M-SAN trained with importance-coherence reward prominently performs better than the other two.

Table 4. Ablation study on glimpse and rewards.

	Imp-Coh@10			Imp-Coh@15		
	Imp	Coh	Overall	Imp	Coh	Overall
M-SAN	80.29	92.16	74.19	77.00	90.83	70.28
w/o glimpse	79.80	91.50	72.97	75.98	90.58	70.12
coh-rwd only	79.37	91.02	72.28	76.90	90.29	69.75
imp-rwd only	67.82	71.62	49.87	67.42	79.89	54.10

Analysis of Reward Ratio. To further figure out which reward ratio brings the optimal results, we experiment on $T = 10$ and $T = 15$, assigning 0.0/0.3/0.5/0.7 to β . Importance score, coherence score, and overall score results are shown in the line chart Fig. 4. The results at the $T = 10$ and $T = 15$ are presented by lines painted in blue and orange respectively. Given target duration $T = 10s$, all three scores reach a peak at $\beta=0.5$. Given target duration $T = 15s$, the importance score and overall score reach a peak at $\beta=0.3$, which is slightly higher than the score at $\beta=0.5$. Coherence score reaches the highest point at $\beta=0.5$. On the whole, the performance is relatively good at the condition of $\beta=0.5$.

4.6 Qualitative Analysis

One result is shown as Fig. 5. The source video is about a collectible game app with real mobile phones as the completion rewards. Its original duration is 36s, and the target duration is 15s. SAM [20] generates videos with too many foreshadowing parts, which exceeds the duration limitation. M-SAN produces a 15s result and tends to select latter segments in source videos, which is reasonable because key points usually appear in the latter part of the ad with the front part




Fig. 4. Comparison results with $\beta = 0.0/0.3/0.5/0.7$.

doing foreshadowing. The result of M-SAN first demonstrates using the app to get a new phone and then shows a scene where a new phone is packed, which emphasizes the rewards of completing the game. This verifies the M-SAN focuses on more informative segments and does better than SAM in duration control.



Fig. 5. Visualization. Source video and videos assembled by SAM and M-SAN given target duration $T = 15s$. τ is the actual duration of result.

5 Conclusion

The two main stages of ad video editing are video segmentation and segment assemblage. Existing methods perform poorly at the segment assemblage stage. To improve the performance of segment assemblage, we proposed M-SAN to perform segment assemblage end-to-end. We also proposed importance-coherence reward based on the characteristics of ad and train M-SAN with policy gradient. We collected an ad video dataset with 1000+ ad videos and proposed Imp-Coh@Time metrics. Experimental results show the effectiveness of M-SAN and verify that multi-modal representation and importance-coherence reward bring a significant performance boost.

Acknowledgment. This work is supported by the National Natural Science Foundation of China under Grant No. 61972188 and 62122035.

References

1. Bahdanau, D., Cho, K., Bengio, Y.: Neural machine translation by jointly learning to align and translate. arXiv preprint arXiv:1409.0473 (2014)
2. Bain, M., Nagrani, A., Brown, A., Zisserman, A.: Condensed movies: Story based retrieval with contextual embeddings. In: Proceedings of the Asian Conference on Computer Vision (2020)
3. Bello, I., Pham, H., Le, Q.V., Norouzi, M., Bengio, S.: Neural combinatorial optimization with reinforcement learning. arXiv preprint arXiv:1611.09940 (2016)
4. Cai, D., Zhang, Y., Huang, Y., Lam, W., Dolan, B.: Narrative incoherence detection. arXiv preprint arXiv:2012.11157 (2020)
5. Cho, K., Van Merriënboer, B., Gulcehre, C., Bahdanau, D., Bougares, F., Schwenk, H., Bengio, Y.: Learning phrase representations using rnn encoder-decoder for statistical machine translation. arXiv preprint arXiv:1406.1078 (2014)
6. Deudon, M., Cournut, P., Lacoste, A., Adulyasak, Y., Rousseau, L.M.: Learning heuristics for the tsp by policy gradient. In: International conference on the integration of constraint programming, artificial intelligence, and operations research. pp. 170–181. Springer (2018)
7. Devlin, J., Chang, M.W., Lee, K., Toutanova, K.: Bert: Pre-training of deep bidirectional transformers for language understanding. arXiv preprint arXiv:1810.04805 (2018)
8. Fajtl, J., Sokeh, H.S., Argyriou, V., Monekosso, D., Remagnino, P.: Summarizing videos with attention. In: Asian Conference on Computer Vision. pp. 39–54. Springer (2018)
9. Feng, L., Li, Z., Kuang, Z., Zhang, W.: Extractive video summarizer with memory augmented neural networks. In: Proceedings of the 26th ACM international conference on Multimedia. pp. 976–983 (2018)
10. Fu, T.J., Tai, S.H., Chen, H.T.: Attentive and adversarial learning for video summarization. In: 2019 IEEE Winter Conference on applications of computer vision (WACV). pp. 1579–1587. IEEE (2019)
11. Ghauri, J.A., Hakimov, S., Ewerth, R.: Supervised video summarization via multiple feature sets with parallel attention. In: 2021 IEEE International Conference on Multimedia and Expo (ICME). pp. 1–6s. IEEE (2021)
12. Gong, Y., Zhu, Y., Duan, L., Liu, Q., Guan, Z., Sun, F., Ou, W., Zhu, K.Q.: Exact-k recommendation via maximal clique optimization. In: Proceedings of the 25th ACM SIGKDD international conference on knowledge discovery & data mining. pp. 617–626 (2019)
13. Guo, D., Zeng, Z.: Multi-modal representation learning for video advertisement content structuring. In: Proceedings of the 29th ACM International Conference on Multimedia. pp. 4770–4774 (2021)
14. Hershey, S., Chaudhuri, S., Ellis, D.P., Gemmeke, J.F., Jansen, A., Moore, R.C., Plakal, M., Platt, D., Saurous, R.A., Seybold, B., et al.: Cnn architectures for large-scale audio classification. In: 2017 IEEE international conference on acoustics, speech and signal processing (icassp). pp. 131–135. IEEE (2017)
15. Huang, Q., Xiong, Y., Xiong, Y., Zhang, Y., Lin, D.: From trailers to storylines: An efficient way to learn from movies. arXiv preprint arXiv:1806.05341 (2018)
16. Ji, Z., Xiong, K., Pang, Y., Li, X.: Video summarization with attention-based encoder-decoder networks. IEEE Transactions on Circuits and Systems for Video Technology **30**(6), 1709–1717 (2019)

17. Kikuchi, Y., Neubig, G., Sasano, R., Takamura, H., Okumura, M.: Controlling output length in neural encoder-decoders. arXiv preprint arXiv:1609.09552 (2016)
18. Kool, W., Van Hoof, H., Welling, M.: Attention, learn to solve routing problems! arXiv preprint arXiv:1803.08475 (2018)
19. Koutras, P., Maragos, P.: Susinet: See, understand and summarize it. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops. pp. 0–0 (2019)
20. Lin, Q., Pang, N., Hong, Z.: Automated multi-modal video editing for ads video. In: Proceedings of the 29th ACM International Conference on Multimedia. pp. 4823–4827 (2021)
21. Lin, R., Xiao, J., Fan, J.: Nextvlad: An efficient neural network to aggregate frame-level features for large-scale video classification. In: Proceedings of the European Conference on Computer Vision (ECCV) Workshops. pp. 0–0 (2018)
22. Liu, M., Zhang, W., Orabona, F., Yang, T.: Adam: A stochastic method with adaptive variance reduction. arXiv preprint arXiv:2011.11985 (2020)
23. Liu, Y.T., Li, Y.J., Wang, Y.C.F.: Transforming multi-concept attention into video summarization. In: Proceedings of the Asian Conference on Computer Vision (2020)
24. Liu, Z., Lin, Y., Cao, Y., Hu, H., Wei, Y., Zhang, Z., Lin, S., Guo, B.: Swin transformer: Hierarchical vision transformer using shifted windows. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. pp. 10012–10022 (2021)
25. Merler, M., Joshi, D., Mac, K.N.C., Nguyen, Q.B., Hammer, S., Kent, J., Xiong, J., Do, M.N., Smith, J.R., Feris, R.S.: The excitement of sports: Automatic highlights using audio/visual cues. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition Workshops. pp. 2520–2523 (2018)
26. Merler, M., Joshi, D., Nguyen, Q.B., Hammer, S., Kent, J., Smith, J.R., Feris, R.S.: Automatic curation of golf highlights using multimodal excitement features. In: 2017 IEEE Conference on Computer Vision and Pattern Recognition Workshops (CVPRW). pp. 57–65. IEEE (2017)
27. Messaoud, S., Lourentzou, I., Boughoula, A., Zehni, M., Zhao, Z., Zhai, C., Schwing, A.G.: Deepqamvs: Query-aware hierarchical pointer networks for multi-video summarization. In: Proceedings of the 44th International ACM SIGIR Conference on Research and Development in Information Retrieval. pp. 1389–1399 (2021)
28. Mnih, V., Heess, N., Graves, A., et al.: Recurrent models of visual attention. *Advances in neural information processing systems* **27** (2014)
29. Nazari, M., Oroojlooy, A., Snyder, L.V., Takác, M.: Deep reinforcement learning for solving the vehicle routing problem. arXiv preprint arXiv:1802.04240 (2018)
30. Radford, A., Wu, J., Child, R., Luan, D., Amodei, D., Sutskever, I., et al.: Language models are unsupervised multitask learners. *OpenAI blog* **1**(8), 9 (2019)
31. Rochan, M., Krishna Reddy, M.K., Ye, L., Wang, Y.: Adaptive video highlight detection by learning from user history. In: European conference on computer vision. pp. 261–278. Springer (2020)
32. Vinyals, O., Bengio, S., Kudlur, M.: Order matters: Sequence to sequence for sets. arXiv preprint arXiv:1511.06391 (2015)
33. Vinyals, O., Fortunato, M., Jaitly, N.: Pointer networks. *Advances in neural information processing systems* **28** (2015)
34. Wang, J., Wang, W., Wang, Z., Wang, L., Feng, D., Tan, T.: Stacked memory network for video summarization. In: Proceedings of the 27th ACM International Conference on Multimedia. pp. 836–844 (2019)

35. Williams, R.J.: Simple statistical gradient-following algorithms for connectionist reinforcement learning. *Machine learning* **8**(3), 229–256 (1992)
36. Xiong, B., Kalantidis, Y., Ghadiyaram, D., Grauman, K.: Less is more: Learning highlight detection from video duration. In: *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*. pp. 1258–1267 (2019)
37. Xiong, Y., Huang, Q., Guo, L., Zhou, H., Zhou, B., Lin, D.: A graph-based framework to bridge movies and synopses. In: *Proceedings of the IEEE/CVF International Conference on Computer Vision*. pp. 4592–4601 (2019)
38. Yang, H., Wang, B., Lin, S., Wipf, D., Guo, M., Guo, B.: Unsupervised extraction of video highlights via robust recurrent auto-encoders. In: *Proceedings of the IEEE international conference on computer vision*. pp. 4633–4641 (2015)
39. Zhang, K., Grauman, K., Sha, F.: Retrospective encoders for video summarization. In: *Proceedings of the European conference on computer vision (ECCV)*. pp. 383–399 (2018)
40. Zhou, K., Qiao, Y., Xiang, T.: Deep reinforcement learning for unsupervised video summarization with diversity-representativeness reward. In: *Proceedings of the AAAI Conference on Artificial Intelligence*. vol. 32 (2018)