Multi-Modal Video Topic Segmentation with Dual-Contrastive Domain Adaptation

Linzi Xing[⋆1], Quan Tran[2], Fabian Caba[2], Franck Dernoncourt[2], Seunghyun Yoon[2], Zhaowen Wang[2], Trung Bui[2], and Giuseppe Carenini[1]

[1] University of British Columbia, Vancouver, BC, Canada
[2] Adobe Research, San Francisco, CA, USA
{lzxing, carenini}@cs.ubc.ca
{qtran, caba, dernonco, syoon, zhawang, bui}@adobe.com

Abstract. Video topic segmentation unveils the coarse-grained semantic structure underlying videos and is essential for other video understanding tasks. Given the recent surge in multi-modal, relying solely on a single modality is arguably insufficient. On the other hand, prior solutions for similar tasks like video scene/shot segmentation cater to short videos with clear visual shifts but falter for long videos with subtle changes, such as livestreams. In this paper, we introduce a multi-modal video topic segmenter that utilizes both video transcripts and frames, bolstered by a cross-modal attention mechanism. Furthermore, we propose a dual-contrastive learning framework adhering to the unsupervised domain adaptation paradigm, enhancing our model's adaptability to longer, more semantically complex videos. Experiments on short and long video corpora demonstrate that our proposed solution, significantly surpasses baseline methods in terms of both accuracy and transferability, in both intra- and cross-domain settings.

Keywords: Topic Segmentation · Video Understanding · NLP

1 Introduction

Video Topic Segmentation aims to break stretches of videos into smaller segments consisting of video frames or clips consistently addressing a common topic. As an example given in Figure 1, video topic segmentation does the job of segmenting a creative livestream video (e.g., from livestream platform *Behance*[3]) into a sequence of topical-coherent pieces (Seg1 – Seg8), placing a boundary where a topic transition happens. This task can enhances both human-to-human and human-to-system interactions in modern social contexts, improving real-time engagement on live streaming platform [9]. In particular, the relatively coarse-grained temporal structure of the input video produced by video topic segmentation is shown to not only simplify video comprehension and helps viewers find content of interest easily. More importantly, it can substantially benefit other key

---

[⋆] Work done while the first author was an intern at Adobe Research.
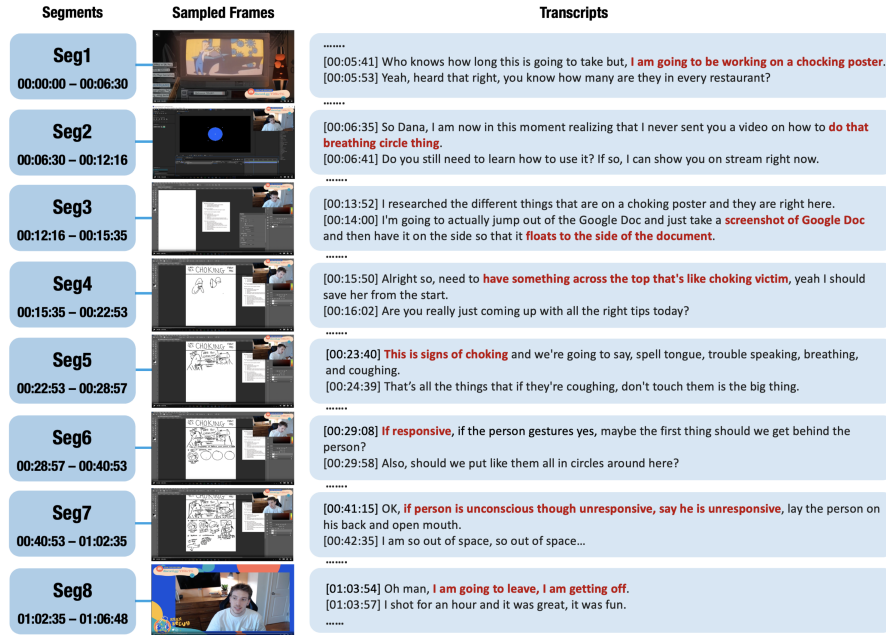[3] https://www.behance.net/live

Fig. 1: A Behance exemplar about making a choking victim poster. The left side of the figure illustrates the video's timeline after topic segmentation. The right side shows the transcript with words indicating a segment's topic highlighed.

video understanding tasks such as video summarization [28], and query-driven video localization [37].

Early computational models for video segmentation primarily targeted shot or scene detection by merely leveraging surface visual features like spatiotemporal aspects or frame colors [32,31,30,6]. These approaches typically measure the temporal similarity along a video's timeline to predict shot/scene boundaries. Despite the difference in definition from shot/scene segmentation (discussed in §2), the task of video topic segmentation focuses more on topic-related semantics in the video, which is not necessarily aligned with visual changes. As shown in Figure 1, the visual background remains similar in Seg3 – Seg7 for a considerable time, even though the streaming topic changes drastically. Moreover, prior video segmentation methods mostly focused on short videos with clear visual changes and simple patterns [13,18,16]. These distinctive features of short videos could be emphasized in model design or learned in supervised setups, making such models less adaptable to longer, more nuanced video content, like documentaries or instructional livestreams.

To address the aforementioned issues, in this paper, we first propose a simple yet effective multi-modal model for video topic segmentation, which can take both the aligned video transcript and visual frames as input. This ability considerably enhances the model's performance, as textual and visual features can work

together to more comprehensively reflect the input video's topic-related semantics [17,2]. Similar to the formulation used for text topic segmentation [21,39,24], we treat video topic segmentation as a sequence labeling task and introduce a neural model equipped with a cross-modal attention mechanism going beyond simple fusion to effectively integrate textual and visual signals in a complementary manner. We initially conduct intra-domain experiments by training and testing our proposed segmenter on a newly collected YouTube corpus equipped with high-quality human labeling. Empirical results show that our multi-modal approach outperforms a set of baseline video segmenters by a substantial margin.

To further adapt our proposed model trained on YouTube videos to longer videos with complex visuals and semantics, we propose an unsupervised domain adaptation strategy empowered by the sliding window inference and dual-contrastive learning scheme. Further experiments on two out-of-domain long video corpora demonstrate that the model's generality can be significantly improved through applying our dual-contrastive adaptation approach.

2   Related Work

Topic Segmentation seeks to uncover the semantic structure of a document (either monologue [39] or dialogue [38]) by dividing it, typically a sequence of sentences, into topical-coherent segments. Recently, a number of supervised neural solutions have been introduced owing to the availability of large-scale labeled corpora sampled from *Wikipedia* [21], with section marks as gold segment boundaries. Most of these neural segmentation approaches follow the same strategy to simply interpret text segmentation as a sequence labeling problem and further tackle it using a variety of hierarchical neural sequence labelers [21,23,39,40].

Inspired by above-mentioned neural text segmenters, our paper similarly frame video topic segmentation as a sequence labeling task due to the availability of textual input (video transcript) and utilize a hierarchical neural sequence labeling framework as the basic architecture of our proposal. Then we extend such framework into the multi-modal setting, with the injection of visual signals from video frames by adding a cross-modal attention network on top of it. More details will be presented in §3.

Shot and Scene Segmentation are closely related to video topic segmentation but more narrowly focused. A shot is a sequence of frames from a continuous camera capture. Hence, most shot segmentation techniques mainly use the visual modality to group video frames into shots. Conversely, a scene is more semantically intricate than a shot, representing a series of related shots that depict events defined by elements like actions, places, and characters. These elements are mostly found in narrative videos such as movies. Thus, past scene segmentation techniques are primarily developed for movies exclusively, aiming to group consecutive shots into scenes based on their visual consistency, spatiotemporal features, or shot color similarity. Similar but more difficult than shot or scene segmentation, our work focuses on video topic segmentation, which can be deemed as an extension of text topic segmentation with "topic" defined as a relatively self-contained collection of semantically close information. Notably, the topic's

definition relies heavily on the video's context and domain, and isn't strictly bound to the concepts of shots or scenes as previously defined. From a machine learning perspective, the more dynamic nature of shot/scene segmentation favours the visual signal, while in our task, the video topic segmentation relies more heavily on the semantic signals carried by natural language [17,2]. Therefore, our video topic segmenter is designed to extend from a topic segmentation framework for text, by integrating visual frames as the auxiliary signal.

Contrastive Learning algorithms aim to learn data representations by enlarging the distance between dissimilar samples and meanwhile minimizing the distance between similar samples with contrastive loss functions [27]. These techniques have been observed promising in domain adaptation for both uni-modal and cross-modal settings [19,20,5]. Recent works on using contrastive learning for video domain adaptation mainly focus on transferring models from source to target domain, where both domains consist of short videos [20]. In contrast, our work attempts to adapt the model trained on short videos to the low-resource domain containing long videos with subtler visual changes, by utilizing the semantic overlaps within a single modality or between two modalities to guide the contrastive learning process.

3   Neural Video Topic Segmentation

3.1   Problem Definition

Inspired by recent neural-based supervised approaches for text topic segmentation for text, we similarly frame video topic segmentation as a sequence labeling task, with sentences in the video transcript as the units for labeling. More precisely, given an input video with (1) the transcript as a sequence of timestamped sentences, and (2) a sequence of timestamped video frames, our model will predict a binary label for each transcript sentence to indicate whether or not the sentence indicates a topic segment boundary. Formally,

Given: A video v with its transcript T_v as a sequence of sentences $\{s_1, s_2, ..., s_n\}$ along with start time offsets $\{b_1, b_2, ..., b_n\}$ / end time offsets $\{e_1, e_2, ..., e_n\}$, and video frames $X_v = \{x_1, x_2, ... , x_m\}$, as a single frame x_i has timestamp t_i.

Predict: A sequence of labels $\{l_1, l_2, ..., l_{n-1}\}$ for the sequence of transcript sentences, where l is a binary label, 1 means the corresponding sentence overlaps a video topic segment boundary, 0 otherwise. We do not predict the label for the last sentence s_n, as it is by definition equal to 1, i.e., the end of the last segment.

3.2   Model Architecture

Figure 2 illustrates the detailed framework of our proposed video topic segmenter, which is similar in architecture to TextSeg [21]. It comprises two hierarchically linked encoding layers: one as text encoder for contextualized encoding within a sentence (orange in Figure 2) and the other for contextualized encoding between sequence units (green in Figure 2). To allow both textual and visual modalities contribute complementarily to the model's prediction, we add a frame
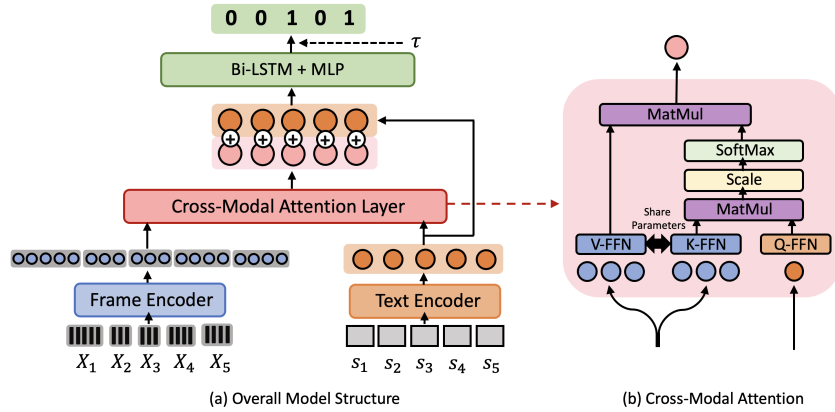
Fig. 2: Left: The overall architecture of our proposed multi-modal video topic segmentation model, with four sub-modules (in §3.2) coded in different colors. **Right**: The detailed illustration of cross-modal attention.

encoder module (blue in Figure 2) and a cross-modal attention mechanism (red in Figure 2). This design choice extends the framework to the multi-modal setting, accepting now both the video transcript and frames as input and making prediction of topic segment boundaries based on them.

The `text encoder` module E_t yields low-level features for sentences in the video transcript. Different from the proposal in [21] using the BiLSTM + attention as the backbone for text encoder, here we adopt the pre-trained vanilla BERT [7] in accordance to its achieved superiority on text segmentation observed in [39] and [26]. Parallel to the text encoder, a `frame encoder` E_f is introduced to extract features for visual signals with the standard pretrained ResNet-18 [14]. Formally, given a transcript sentence s_i with its time interval as $[b_i, e_i]$ and a set of video frames $X_i = \{x_1^i, ..., x_m^i\}$ associated with this sentence, where each frame (e.g., the k^{th} frame) in the set has timestamp $t_k^i \in [b_i, e_i]$, we can obtain the sentence representation $tr_i = E_t(s_i)$ and its corresponding set of frame representations $FR_i = \{fr_1^i, ..., fr_m^i\}$, where $fr_k^i = E_f(x_k^i)$.

Next, we propose to use a `cross-modal attention` mechanism to produce a text-aware visual representation for each sentence, rather than obtaining the visual representation by naively operating mean-pooling over the set of frame representations covered by the sentence interval. This design is motivated by the observation that transcript sentences and video frames are in an one-to-many relation, and frames sharing more semantics with the text should be given more attention weights for visual representation generation. Irrelevant frames sharing no or little semantics with the text may negatively affect the quality of the fused multi-modal representation passed to the subsequent module [36]. In practice, the cross-modal attention module adopts the standard scaled dot-product attention function proposed in [35]. With the transcript sentence representation tr_i and its corresponding frame representation set $FR_i = \{fr_1^i, ..., fr_m^i\}$ as query and key (value), we compute the text-aware visual representation vr_i as:

$$vr_i = A_i V_i, \tag{1}$$

$$A_i = softmax(\frac{q_i K_i^T}{\sqrt{d_k}}) \tag{2}$$

where $q_i \in \mathbb{R}^{1 \times d_k}$, $K_i \in \mathbb{R}^{m \times d_k}$, and $V_i \in \mathbb{R}^{m \times d_k}$ denote the query vector, key and value matrices generated by passing the sentence representation and frame representations through three parallel feedfoward layers, namely *Q-FFN*, *K-FFN* and *V-FFN* respectively. More formally, we have $q_i = Q\text{-}FFN(tr_i)$, $K_i = K\text{-}FFN(FR_i)$, $V_i = V\text{-}FFN(FR_i)$, where *K-FFN* and *V-FNN* share the same parameters and thus produce identical key and value matrices following [25,22].

Then all the obtained text-aware visual representations $\{vr_1, ..., vr_n\}$ are concatenated with their corresponding sentence representations $\{tr_1, ..., tr_n\}$ and fed into a BiLSTM layer which performs contextualization and returns hidden states. Next, a multilayer perceptron (MLP) followed by Softmax serves as a topic boundary predictor to make binary predictions regarding the input hidden states according to a threshold τ tuned on the validation set. More specifically, if a transcript sentence's output probability exceeds τ, it's marked as 1, indicating a segment boundary. The entire model is fine-tuned using cross-entropy loss.

We train and test this model on a newly collected YouTube corpus (in §5.1) and empirically verify its in-domain effectiveness (reported in Table 2). We leverage it as the source model in the next section to help deliver the domain adaptation strategy coupled with our proposed multi-modal framework.

4   Long Video Adaptation

The rise of (live-)streaming platforms has increased the demand for topical segmentation of videos on these platforms. Unlike videos (e.g., on YouTube) with careful pre-editing and segment labeling provided by their creators, videos on (live-)streaming platforms (e.g., *BBC* documentaries and creative livestreams on *Behance*) are usually extremely long, with sparse visual changes, and more importantly, time-consuming to obtain segment annotations. Thus, it is impractical to learn a fully-supervised model on such videos, and the segmentation model described in §3.2 (source model) trained on short YouTube videos (source domain) might underperform when applied to these extensively long videos (target domain). To adapt the source model for the above-described target domain, we propose to equip it with two strategies, namely `Sliding Window Inference` (§4.1) and `Dual-Contrastive Adaptation` (§4.2).

4.1   Simple Sliding Window Inference

Due to the length discrepancy between videos from the source and target domains, directly applying the source model to the target input taking the full transcript sentence sequence as input is observed to cause extremely sparse boundary predictions [12]. Therefore, we propose to calibrate the input length by first breaking the long target input into snippets by a sliding window with the size consistent with the source input length (i.e., the average length of the YouTube

videos). After applying the source model on every snippet and aggregating outputs of all snippets, we eventually make binary prediction for each sentence if its aggregated probability exceeds the threshold τ pre-tuned on the source domain. Formally, given a long video input from the target domain with length $= n$ and a fixed window size $= k$ $(n \gg k)$, we can create $n - k + 1$ snippets $\{\mathbb{S}_1, ..., \mathbb{S}_{n-k+1}\}$ where each snippet consists of k consecutive sentences by sliding the window over the input with the stride of 1. As a result, each transcript sentence s_m can be covered by up to k snippets. Once we apply the source model to all snippets, we can obtain multiple probability predictions for s_m. We then aggregate these probability predictions associated with s_m by taking average:

$$\bar{p}_m = \frac{1}{k} \sum_{i=1}^{k} p_m^{\mathbb{S}_i} \tag{3}$$

Finally, we predict that s_m falls on a segment boundary if $\bar{p}_m > \tau$.

4.2   Dual-Contrastive Adaptation

To transfer the pre-trained source model to the target domain while still preserving its performance on the source domain, we follow a more sophisticated unsupervised domain adaptation paradigm, which leverages both labeled source data and unlabeled target data. Specifically, we fix the frame and text encoder (E_t and E_f) while updating the rest of the model in two steps, where the first step updates the model on unlabeled target data using two contrastive learning objectives namely *intra-modal contrastive loss* and *cross-modal contrastive loss*, while the second step updates the model on source data with the supervised training scheme described in §3.2.

As the overview of the first step shown in Figure 3, given a collection of the paired sentence representations and their corresponding frame representation sets $\{tr_i, FR_i\}_{i=1}^{b}$ in a training batch (size=b) produced by the text and frame encoder from the target domain, we need two distinct projection heads to map these sentence/frame representations into a shared space. Here we use *Q-FFN* and *K-FNN* (which shares parameters with *V-FNN*) in the cross-modal attention to serve as projection heads for textual and visual modality respectively. Thus we have:

$$q_i = Q\text{-}FFN(tr_i), K_i = K\text{-}FFN(FR_i) \tag{4}$$

where q_i and $K_i = \{k_1^i, ..., k_m^i\}$ denote the projected sentence embedding and the set of projected frame embeddings covered by the sentence. As frames covered by the same sentence are more likely to share similar semantics, we pull the frames attached to the same sentence closer and push the ones from different sentences far apart, by minimizing the intra-modal contrastive loss for visual modality defined as:

$$l^{intra} = - \sum_{i=1}^{b} log \frac{exp(\tilde{k}^i \cdot \tilde{k}^i)/\tau}{\sum_{j=1}^{b} exp(\tilde{k}^i \cdot \tilde{k}^j)/\tau} \tag{5}$$

where $\tilde{k}^i \in_R FR_i$ and τ is the hyper-parameter of temperature. To learn the semantic relation between modalities in the target domain, we first obtain the

visual representation matched with a sentence by averaging all frames covered by the sentence, and then bring semantically close sentence-visual pairs together and push away non-related pairs by minimizing the cross-modal contrastive loss:

$$l^{cross} = -\sum_{i=1}^{b} log \frac{exp(q_i \cdot MP(K_i))/\tau}{\sum_{j=1}^{b} exp(q_i \cdot MP(K_j))/\tau} \qquad (6)$$

where MP denotes *Mean-Pooling*. The total dual-contrastive loss is formed as:

$$l = l^{intra} + l^{cross} \qquad (7)$$

The above dual-contrastive learning phase is followed by the second step. In this step, the model is further trained on the labeled source data again to leverage the target-domain signals learned by the parameters in cross-modal attention for boundary prediction. This step ensures the model adapted to the target domain still preserves some level of effectiveness when applied to the source domain.
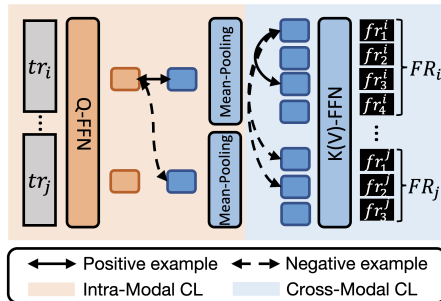

Fig. 3: An illustration of updating cross-modal attention module with dual-contrastive adaptation.

Dataset	YouTube	BBC	Behance
# of Vids.	5,422	11	575
# of Tokens/Vid.	1,411	2,868	11,554
# of Sents./Vid.	108	216	1,287
# of Segs./Vid.	6.7	29.1	5.2
# of Tokens/Seg.	209	99	2,229
# of Sents./Seg.	16	8	248
# of Frames/Sec.	10	4	4
Avg. Length	0:09:30	0:45:02	2:07:51

Table 1: Statistics of corpora used for training and evaluation in §5.

5   Experimental Setup

To evaluate the effectiveness and generality of our proposed video topic segmentation model (in §3) and long video adaptation strategy (in §4), we conduct experiments on two different settings, namely ***Intra-Domain Segment Inference*** and ***Cross-Domain Segment Inference***.

5.1   Intra-Domain Dataset – YouTube

For intra-domain segment inference, we train and test models with the data from the same domain (corpus). Due to the lack of large-scale human-annotated dataset in the field of video segmentation, we collect a novel corpus consisting of 5,422 user-generated videos randomly sampled from YouTube. During the video collecting process, we applied filtering criteria to eliminate unsuitable samples to ensure the quality of the dataset to construct. These criteria include constraints on video length (> 100s), word count (> 0.5 word/second on average), chapter durations (all chapters with length > 10s), and sentence length (all sentences

with length < 60 tokens) inspired by [4]. Each video in this corpus is associated with a series of user-defined chapter timestamps indicating the beginning of each topic chapter contained in this video. We thereby use these available chapter beginning timestamps as ground-truth topic segment boundaries since a topic chapter is by definition a main thematic division within a video. The video transcripts are generated by YouTube ASR with token-level offsets and we further exploit a top performing punctuation restoration model [1] to boost the quality of transcript sentence segmentation. Compared with other existing video understanding corpora such as *BBC* [2] (for scene segmentation, size = 11) and *TVSum* [33] (for video summarization, size = 50), our constructed YouTube corpus is (1) larger in size; (2) covering more diverse topics; (3) with reliable segmentation which has already been specified by video creators. We split this corpus into train/dev/test portions with size: 5148/134/140.

5.2 Datasets for Cross-Domain Inference

To evaluate our proposal's robustness in cases where a domain-shift is present, we conduct experiments for cross-domain segment inference, in which our proposed supervised segmenter is initially trained on the YouTube video corpus, and tested on two corpora with videos significantly longer than YouTube videos:

BBC Planet Earth [2] consists of 11 episodes from the BBC educational TV series Planet Earth. Topic segment boundaries of the dataset have been manually annotated by human experts and transcript sentences are obtained using *Whisper*[4] [29]. As the statistics in Table 1 indicates, this dataset has longer videos with more segments than YouTube videos, while each segment covers much fewer sentences and tokens. We split this dataset into 5 and 6 for validation and testing.

Behance Livestream consists of 575 videos sampled from the creative livestream platform *Behance*. Livestreams, in general, are very challenging to segment into coherent sections, since they contain mixes of multi-user dialogues, with visual features that change little over the video. They are also very long, with hours of content in each video. Behance livestreams, in particular, contain another layer of challenge, since they are highly specialized in creative tasks such as animation and image edit with unknown entities and intricate visual operations. We believe finding structures in such noisy environment would be not only a good benchmark for future topic segmentation research, but also of great practical value for consumers, who would find it hard to navigate contents in a back-and-forth conversation. Thus, we collect such a corpus and annotate its videos with topic segments given the help from domain experts in animation and image edits. For each video, we pay \$23.16 for human annotators to watch and create segments based on the video content. The total cost to annotate 575 videos is \$13,317. We split this dataset into validation and test sets with 57 (10%), 518 (90%).

5.3 Baselines
We consider the following representative baselines as comparisons:

---

[4] https://openai.com/research/whisper

- Random: Given a video with its transcript consisting of k sentences, we first randomly sample the segment number $b \in \{0, ..., k-1\}$ for this video, and then determine if a sentence overlaps a segment boundary with probability $\frac{b}{k}$.

- BayesSeg [8] originally proposed for text segmentation by predicting segment boundaries through modeling the lexical cohesion in a Bayesian context.

- GraphSeg [11] generates a semantic relatedness graph with sentences as nodes. Segments are then predicted as the maximal cliques in graph.

- Cross-BERT [26] is a supervised text topic segmenter representing candidate segment boundaries using their left and right contexts encoded with BERT. The model is trained on YouTube and applied to longer BBC and Behance videos.

- TransNet [34]: This model adopts stacked Dilated DCNN blocks as its basic framework and targets video segmentation on shot level. Here we use its publicly available version trained on synthesized video corpora off-the-shelf.

- LGSS [30] is a multi-modal movie scene segmenter pre-trained on the limited-scale *MovieScenes* corpus. We use its public version with visual frames as input.

- X-Tiling is our extension of *TextTiling* [15]. While TextTiling only allows textual embeddings as input, X-Tiling accepts both textual and visual embeddings and their concatenations as input to compute semantic coherence and then make segment boundary predictions for videos. For fair comparison, the input textual and visual embeddings are produced by pre-trained BERT and mean-pooling over ResNet-18 respectively.

All hyper-parameters required by the above baselines are tuned on the validation portion of the datasets included in this paper.

5.4   Evaluation Metrics

We apply three standard metrics in previous literature to evaluate the performances of our proposal and baselines. They are:

- F_1, with higher scores denoting better performance. It measures the exact match between ground truth and model's prediction.

- P_r error score [10], which fixes the inadequacies of P_k [3] and *WindowDiff* [?], previously considered as two standard evaluation metrics for text segmentation. Concretely, P_r is the mean of missing and false alarm probabilities, calculated based on the overlap between ground-truth segments and model's predictions within a certain size sliding window. Since it is a penalty metric, lower score indicates better performance.

- $mIoU$ score [41], shortened from the mean Intersection-over-Union. It is calculated by taking average over maximal IoUs of all ground-truth segments to predicted segments. Higher score indicates better performance.

5.5   Implementation Details

In our multi-modal video topic segmenter, we employ the `[CLS]` token representation from `bert-base-uncased` (dimension $d = 768$) for sentence representation and ResNet-18's avg-pooling layer output (dimension $d = 512$) for frame

Method	Modalities	Pr $\downarrow$	F1 $\uparrow$	mIoU $\uparrow$
Random	–	45.84	48.98	37.76
X-Tiling	Text	39.94	51.56	50.97
BayesSeg	Text	40.70	50.24	49.69
GraphSeg	Text	38.12	51.41	51.73
Cross-BERT*	Text	<u>32.89</u>	<u>60.48</u>	<u>60.00</u>
X-Tiling	Visual	37.53	52.08	53.24
TransNet	Visual	40.14	51.01	50.66
LGSS	Visual	39.77	50.87	51.13
X-Tiling	Text + Visual	38.78	52.30	51.45
NeuralSeg* (Ours)	Text	31.91	63.23	61.36
	Visual	50.18	47.59	16.15
	Text + Visual	**30.61**	**65.29**	**63.11**

Table 2: Results on YouTube for intra-domain evaluation. **Bold** results indicate the best performance across all comparisons. Underlined results indicate the best performance within their own sub-section. * indicates a fully supervised setting.

representation. The cross-modal attention's feedforward layers have an output dimension of 768. Our BiLSTM has 2 layers with a hidden size of 256. For YouTube training, we use the Adam optimizer with a learning rate of $1e^{-3}$ and a batch size of 16. In the long video adaptation with dual-contrastive learning, the SGD optimizer is applied with a mini-batch size of 256, learning rate of $3e^{-2}$, and softmax temperature of $1e^{-1}$. Training spans 10 epochs for both supervised learning and domain adaptation, with results averaged over 3 runs. The segmentation threshold, τ, is tuned on validation sets for both intra-domain and cross-domain evaluations.

6   Results and Discussion

Intra-Domain Segment Inference: Table 2 reports the performance of the chosen baselines and our proposal (NeuralSeg in the table) on the YouTube testing set, while NeuralSeg is trained on the YouTube training set with different input modality settings. Notably, our model significantly outperforms the best baseline, Cross-BERT, even when only using video transcripts (text) as input. But on the other hand, if we train the model by using only video frames (visual) as input, the model's performance is even considerably worse than the random baseline, possibly because that too diverse visual and topic presence in the corpus makes it difficult for the model to learn a meaningful visual input-to-prediction projection. Yet, by combining both modalities with cross-modal attention fusion, the model's performance can be further enhanced compared with only textual modality. These results confirm that the visual information itself may not be sufficient to capture a video's topic-related semantics under a supervised setting, but fusing it together with the textual information can provide a more clear picture of the video's underlying topics.

Cross-Domain Segment Inference: Table 3 compares the performance of the baselines and our proposed segmenter pre-trained on YouTube w/ or w/o applying our long video adaptation strategy to two challenging long video corpora.

Method	Modalities	BBC			Behance		
		Pr ↓	F1 ↑	mIoU ↑	Pr ↓	F1 ↑	mIoU ↑
Random	–	48.29	47.36	30.66	46.27	50.00	36.24
X-Tiling [15]	Text	<u>41.59</u>	50.72	33.27	45.93	49.95	<u>43.16</u>
BayesSeg [8]	Text	42.01	<u>51.01</u>	<u>33.79</u>	49.01	48.16	37.55
GraphSeg [11]	Text	45.11	49.28	30.23	46.76	50.12	40.38
Cross-BERT [26]	Text	44.63	49.70	31.51	46.15	49.88	41.02
X-Tiling [15]	Visual	44.48	49.89	33.26	<u>44.56</u>	<u>50.73</u>	41.81
TransNet [34]	Visual	42.54	49.82	31.12	46.34	49.88	40.54
LGSS [30]	Visual	42.88	50.02	32.66	45.67	50.12	39.90
X-Tiling [15]	Text + Visual	43.22	50.44	32.68	45.87	50.01	41.72
Ours (Plain)	Text + Visual	43.14	54.13	31.23	45.71	51.25	34.33
Ours (Window)	Text + Visual	40.66	54.92	37.50	43.03	51.83	47.66
Ours (Window + CL-Cross)	Text + Visual	36.45†	54.91	50.92†	42.72	51.68	48.65†
Ours (Window + CL-Intra)	Text + Visual	37.14†	55.10	49.22†	42.56	51.84	48.39
Ours (Window + CL-Dual)	Text + Visual	**<u>36.01</u>**†	**<u>55.75</u>**	**<u>51.56</u>**†	**<u>42.25</u>**†	**<u>51.91</u>**	**<u>49.35</u>**†

Table 3: Results on BBC and Behance corpora for cross-domain evaluation. **Bolded** and <u>underlined</u> results indicate the best performance across all comparisons and within their own section. † indicates results applied contrastive domain adaptation which are significantly different ($p < 0.05$) from *Ours (Window)*.

To better investigate the effectiveness of the two contrastive objectives associated with long video adaptation, Table 3 also shows the results by adding each objective (CL-Cross/CL-Intra) individually. The primary takeaway is that our long video adaptation strategy (Window + CL-Dual) significantly and consistently improves the performance achieved by the segmenter initially trained on short YouTube videos (Plain). In more detail, we can observe that utilizing the sliding window inference strategy (Plain → Window) already yields noticeable performance gains on both corpora. Furthermore, adding contrastive learning objectives can further boost accuracy on two long video corpora. The greatest enhancement is seen when the model is tuned with dual-contrastive losses.

7   Conclusion and Future Work

We present a multi-modal video topic segmentation model accepting both video transcripts and frames as input. Further, we propose a novel unsupervised domain adaptation strategy empowered by a dual-contrastive learning framework to generalize our model pre-trained on short videos to longer videos with more complex content and subtler visual changes. Experiments on two settings (intra-domain and cross-domain segment inference) show that (1) our system achieves the SOTA performance on a newly collected YouTube corpus consisting of large-scale but short videos; (2) When we apply our long video adaption strategy, the model for short videos achieves better performance when transferred to two domains comprising long (live-)stream videos.

References

1. Alam, T., Khan, A., Alam, F.: Punctuation restoration using transformer models for high-and low-resource languages. In: Proceedings of the Sixth Workshop on Noisy User-generated Text (W-NUT 2020). pp. 132–142 (2020)

2. Baraldi, L., Grana, C., Cucchiara, R.: A deep siamese network for scene detection in broadcast videos. In: Proceedings of ACM MM 2015. pp. 1199–1202 (2015)
3. Beeferman, D., Berger, A., Lafferty, J.: Statistical models for text segmentation. Machine Learning **34**(1), 177–210 (1999)
4. Cao, X., Chen, Z., Le, C., Meng, L.: Multi-modal video chapter generation. ArXiv **abs/2209.12694** (2022)
5. Chen, D., Wang, D., Darrell, T., Ebrahimi, S.: Contrastive test-time adaptation. In: Proceedings of CVPR 2022. pp. 295–305 (2022)
6. Chen, S., Nie, X., Fan, D.D., Zhang, D., Bhat, V., Hamid, R.: Shot contrastive self-supervised learning for scene boundary detection. In: Proceedings of CVPR 2021. pp. 9791–9800 (2021)
7. Devlin, J., Chang, M.W., Lee, K., Toutanova, K.: BERT: Pre-training of deep bidirectional transformers for language understanding. In: Proceedings of NAACL 2019. pp. 4171–4186. Association for Computational Linguistics (2019)
8. Eisenstein, J., Barzilay, R.: Bayesian unsupervised topic segmentation. In: Proceedings of EMNLP 2008. pp. 334–343 (2008)
9. Fraser, C., Kim, J., Shin, H., Brandt, J., Dontcheva, M.: Temporal segmentation of creative live streams. In: Proceedings of CHI 2020. pp. 1–12 (2020)
10. Georgescul, M., Clark, A., Armstrong, S.: An analysis of quantitative aspects in the evaluation of thematic segmentation algorithms. In: Proceedings of SIGdial 2006. pp. 144–151 (2006)
11. Glavaš, G., Nanni, F., Ponzetto, S.P.: Unsupervised text segmentation using semantic relatedness graphs. In: Proceedings of the Fifth Joint Conference on Lexical and Computational Semantics. pp. 125–130 (2016)
12. Glavas, G., Somasundaran, S.: Two-level transformer and auxiliary coherence modeling for improved text segmentation. In: Proceeding of AAAI-20. pp. 2306–2315 (2020)
13. Gygli, M., Grabner, H., Riemenschneider, H., Van Gool, L.: Creating summaries from user videos. In: Proceedings of ECCV 2014. pp. 505–520 (2014)
14. He, K., Zhang, X., Ren, S., Sun, J.: Deep residual learning for image recognition. In: Proceedings of CVPR 2016. pp. 770–778 (2016)
15. Hearst, M.A.: Text tiling: Segmenting text into multi-paragraph subtopic passages. Computational Linguistics **23**(1), 33–64 (1997)
16. Jadon, S., Jasim, M.: Unsupervised video summarization framework using keyframe extraction and video skimming. In: Proceedings of ICCCA 2020. pp. 140–145 (2020)
17. James, N., Todorov, K., Hudelot, C.: Combining visual and textual modalities for multimedia ontology matching. In: Semantic Multimedia. pp. 95–110 (2011)
18. Jayaraman, D., Grauman, K.: Slow and steady feature analysis: Higher order temporal coherence in video. In: Proceeding of CVPR 2016. pp. 3852–3861 (2016)
19. Kang, G., Jiang, L., Yang, Y., Hauptmann, A.G.: Contrastive adaptation network for unsupervised domain adaptation. In: Proceedings of CVPR 2019. pp. 4893–4902 (2019)
20. Kim, D., Tsai, Y.H., Zhuang, B., Yu, X., Sclaroff, S., Saenko, K., Chandraker, M.: Learning cross-modal contrastive features for video domain adaptation. In: Proceedings of ICCV 2021. pp. 13598–13607 (2021)
21. Koshorek, O., Cohen, A., Mor, N., Rotman, M., Berant, J.: Text segmentation as a supervised learning task. In: Proceedings of NAACL 2018. pp. 469–473 (2018)
22. Kumar, A., Mittal, T., Manocha, D.: MCQA: multimodal co-attention based network for question answering. CoRR **abs/2004.12238** (2020)

23. Li, J., Sun, A., Joty, S.: Segbot: A generic neural text segmentation model with pointer network. In: Proceedings of IJCAI-18. pp. 4166–4172 (2018)
24. Lo, K., Jin, Y., Tan, W., Liu, M., Du, L., Buntine, W.: Transformer over pre-trained transformer for neural text segmentation with enhanced topic coherence. In: EMNLP 2021 (Findings). pp. 3334–3340 (2021)
25. Lu, J., Yang, J., Batra, D., Parikh, D.: Hierarchical question-image co-attention for visual question answering. In: Proceedings of NeurIPS 2016. pp. 289–297 (2016)
26. Lukasik, M., Dadachev, B., Papineni, K., Simões, G.: Text segmentation by cross segment attention. In: Proceedings of EMNLP 2020. pp. 4707–4716 (2020)
27. van den Oord, A., Li, Y., Vinyals, O.: Representation learning with contrastive predictive coding. ArXiv **abs/1807.03748** (2018)
28. Qiu, J., Zhu, J., Xu, M., Dernoncourt, F., Bui, T., Wang, Z., Li, B., Zhao, D., Jin, H.: Semantics-consistent cross-domain summarization via optimal transport alignment. ArXiv **abs/2210.04722** (2022)
29. Radford, A., Kim, J.W., Xu, T., Brockman, G., McLeavey, C., Sutskever, I.: Robust speech recognition via large-scale weak supervision. ArXiv **abs/2212.04356** (2022)
30. Rao, A., Xu, L., Xiong, Y., Xu, G., Huang, Q., Zhou, B., Lin, D.: A local-to-global approach to multi-modal movie scene segmentation. In: Proceedings of CVPR 2020. pp. 10146–10155 (2020)
31. Rasheed, Z., Shah, M.: Scene detection in hollywood movies and tv shows. In: Proceedings of CVPR 2003. pp. II–343 (2003)
32. Rui, Y., Huang, T.S., Mehrotra, S.: Exploring video structure beyond the shots. In: Proceedings of the IEEE International Conference on Multimedia Computing and Systems. pp. 237–240 (1998)
33. Song, Y., Vallmitjana, J., Stent, A., Jaimes, A.: Tvsum: Summarizing web videos using titles. In: Proceeding of CVPR 2015. pp. 5179–5187 (2015)
34. Souček, T., Lokoč, J.: Transnet v2: An effective deep network architecture for fast shot transition detection. ArXiv **abs/2008.04838** (2020)
35. Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A.N., Kaiser, L.u., Polosukhin, I.: Attention is all you need. In: Advances in Neural Information Processing Systems. vol. 30 (2017)
36. Wang, Z., Zhong, Y., Miao, Y., Ma, L., Specia, L.: Contrastive video-language learning with fine-grained frame sampling. In: Proceedings of AACL-IJCNLP 2022. pp. 694–705 (2022)
37. Xiao, S., Chen, L., Zhang, S., Ji, W., Shao, J., Ye, L., Xiao, J.: Boundary proposal network for two-stage natural language video localization. In: Proceedings of AAAI-21 (2021)
38. Xing, L., Carenini, G.: Improving unsupervised dialogue topic segmentation with utterance-pair coherence scoring. In: Proceedings of SIGdial 2021. pp. 167–177 (2021)
39. Xing, L., Hackinen, B., Carenini, G., Trebbi, F.: Improving context modeling in neural topic segmentation. In: Proceedings of AACL-IJCNLP 2020. pp. 626–636 (2020)
40. Xing, L., Huber, P., Carenini, G.: Improving topic segmentation by injecting discourse dependencies. In: Proceedings of the 3rd Workshop on Computational Approaches to Discourse. pp. 7–18 (2022)
41. Zhu, W., Pang, B., Thapliyal, A.V., Wang, W.Y., Soricut, R.: End-to-end dense video captioning as sequence generation. In: Proceedings of COLING 2022. pp. 5651–5665 (2022)