arXiv:2510.05652v2 [cs.CV] 7 May 2026

SD-MVSum: Script-Driven Multimodal Video Summarization Method and Datasets *

Manolis Mylonas
CERTH-ITI
Thermi, Greece
emylonas@iti.gr

Charalampia Zerva
CERTH-ITI
Thermi, Greece
charazerva@iti.gr

Evlampios Apostolidis
CERTH-ITI
Thermi, Greece
apostolid@iti.gr

Vasileios Mezaris
CERTH-ITI
Thermi, Greece
bmezaris@iti.gr

Abstract

In this work, we present a method and two large-scale datasets for Script-Driven Multimodal Video Summarization. The proposed method, SD-MVSum, builds on our earlier SD-VSum method for script-driven video summarization, which considered just the visual content of the video. SD-MVSum takes into account, in addition to the visual modality, the relevance of the user-provided script with the spoken content (i.e., audio transcript) of the video. The dependence between each considered pair of data modalities, i.e., script-video and script-transcript, is modeled using a new weighted cross-modal attention mechanism. This mechanism explicitly exploits the semantic similarity between the paired modalities in order to promote the parts of the full-length video with the highest relevance to the user-provided script. Furthermore, we extend two large-scale datasets for script-driven (S-VideoXum) and generic (MrHiSum) video summarization, to make them suitable for training and evaluation of script-driven multimodal video summarization methods. Experimental comparisons document the competitiveness of the proposed SD-MVSum method against other SotA approaches for script-driven and generic video summarization. Our new method and extended datasets are available at: https://github.com/IDT-ITI/SD-MVSum

1. Introduction

Various methods for text/query-driven video summarization have been proposed over the last year in the literature, aiming to assist the generation of summarized versions of a full-length video that are customized to the user's needs. In most cases, these needs are expressed using one or more keywords (e.g., "changing tire") [1–3] or a short sentence (e.g., "a man is washing the car") [4, 5], since the relevant methods are not compatible with more extensive descriptions. Consequently, the generated video summaries exhibit limited visual and semantic diversity, as they mainly contain the video parts that match a short-form user query.

To tackle the aforementioned limitation of existing methods, in our previous work [6] we introduced the task of script-driven video summarization and released a relevant large-scale dataset (called S-VideoXum). Using this dataset, we trained a method (called SD-VSum) that gets as input a long-form script outlining the content of the desired video summary, and forms the summary by finding associations between the user script and the visual content of the video based on a cross-modal attention mechanism. However, the spoken content in the video is also a rich source of information for spotting such associations. Driven by this observation, in this work we extend SD-VSum to leverage also the video's spoken content, forming the SD-MVSum method for script-driven multimodal video summarization. Moreover, we introduce a weighted cross-modal attention mechanism that explicitly exploits the semantic similarity between a pair of data modalities when modeling their dependence, to promote the parts of the video with the highest relevance to the user's script. Finally, to assist future research, we extend the previously released S-VideoXum dataset for script-driven video summarization [6], as well as the large-scale MrHiSum dataset for video highlight detection and summarization [7], making them suitable for the task of script-driven multimodal video summarization. Our main contributions are as follows:

- We extend the SD-VSum method for script-driven video summarization, originally considering just the visual content of the video, to leverage also the video's spoken content (audio transcripts).
- We introduce a weighted cross-modal attention mechanism, which explicitly exploits the semantic similarity be-

*This work was supported by project MediaPot (TAEDK-06196), implemented in the framework of National Recovery and Resilience Plan Greece 2.0, funded by the European Union – NextGenerationEU, under the call RESEARCH-CREATE-INNOVATE; and, by project eXeLMM, implemented in the framework of H.F.R.I call "3rd Call for H.F.R.I.'s Research Projects to Support Faculty Members & Researchers" (H.F.R.I. Project Number: 25957).

tween a pair of data modalities when modeling their dependence, in order to promote the parts of the video with the highest relevance to the user-provided script.

- We extend two large-scale datasets for script-driven (S-VideoXum) and generic (MrHiSum) video summarization, in order to make them suitable for training and evaluation of script-driven multimodal video summarization methods.

2. Related Work

2.1. Text/Query-driven video summarization

Early attempts were based on the use of probabilistic and submodular optimization frameworks. Sharghi et al. [1, 2] used probabilistic models to select video shots that were both important to the video and relevant to the query, while Vasudevan et al. [8] employed a submodular optimization framework to ensure the selected frames were relevant to the textual query, but also visually diverse, representative and aesthetically pleasing. A significant shift was observed with the emergence of deep learning. Wei et al. [4] introduced a semantic-attended network that learns to select representative video parts by minimizing the distance between generated summaries and human-provided descriptions. On a different basis, Zhang et al. [9] trained a query-conditioned GAN with a three-player loss, where the generator aims to learn how to create a summary based on a joint representation of the query and video, and the discriminator tries to discriminate the real summary from a generated and a random one. More advanced approaches aimed to capture complex relationships. Jiang et al. [10] designed a hierarchical network with diverse cross-modal and self-attention mechanisms, to model query-related long-range temporal dependencies and take into account user-oriented diversity and stochastic factors. Xiao et al. [11] used local self-attention and query-aware global attention to rank shots according to their semantic relationship with the user query, while Narasimhan et al. [3] introduced CLIP-It, a method using a multi-head language-guided attention mechanism to estimate frames' importance based on their visual relevance and their correlation with the user query. Towards addressing data scarcity, Xiao et al. [12] pretrained a hierarchical self-attentive network for visual importance estimation on the ActivityNet Captions dataset [13], fine-tuned it using a reinforced caption generator, and developed a module that computes shot-level scores for a given query. Mujtaba et al. [14] presented a query-driven approach that employs 2D CNNs and is designed to run on client devices aiming to provide tailored summaries based on individual user preferences. Huang et al. [15] explored the use of self-supervision to generate pseudo-labels and model relationships between pseudo and human labels, and employed context-aware query representations to capture the relevance between visual and textual modalities. Guo et al. [16] presented a method for query-driven multimodal video summarization that gets as input a short-form text, the video frames and the audio transcripts, and fuses them using a transformer-based framework that applies both coarse-grained (taking all modalities into account) and fine-grained (combining text with video and audio individually) fusion, giving equal importance to each modality. Rodrigo et al. [17] described a text-guided framework for automatic sports video summarization that leverages contrastive language–image pretraining to classify video frames as highlight or non-highlight based on natural-language descriptions. Finally, in our previous work [6], we extended the VideoXum large-scale dataset for video summarization by producing textual descriptions of the ground-truth summaries, and trained the SD-VSum method that aligns and fuses visual and textual information using a cross-modal attention mechanism.

In most of the above methods, the users' preferences are expressed by a few keywords [1–3, 8–12, 14, 15] or a short sentence [3, 4, 16, 17]. Contrary to these methods, SD-MVSum gets as input a long-form textual description of the desired video summary, thus allowing the generation of visually and semantically diverse summaries. Moreover, differently from [6] that considers just the visual content of the video, SD-MVSum leverages also the video's spoken content to discover further associations between the user's script and the video, and produce summaries that are better tailored to the user's demands that are expressed in the script.

2.2. Multimodal video summarization

Several attempts have been made to advance the quality of automated video summarization using additional data modalities. Narasimhan et al. [3] examined the performance of CLIP-It when the textual input is formed as a set of auto-generated dense captions of the video content. Following, focusing on the summarization of instructional videos, Narasimhan et al. [20] developed a method that takes into account the video frames and the audio transcripts, and selects video fragments showing important steps of the procedure that are most relevant to the task, but also mentioned in the transcripts. Working also with instructional videos, Palaskar et al. [21] performed their summarization using a multi-source sequence-to-sequence model with hierarchical attention, while a similar approach was adopted by Sanabria et al. [22] for summarizing sports videos. Zhong et al. [23] built a method that creates semantically representative video summaries by minimizing the distance of learnable visual and text representations of the video content and its textual description, respectively, in a common embedding space. Argaw et al. [24] presented a method that employs the visual content and a long-form description of it or the audio transcripts, and trained it with

Table 1. Overview of large-scale datasets for generic (top three) and script-driven (bottom two) video summarization in the literature.

Dataset	Domains	Samples	Data modalities	Annotations per sample	Type of annotations	Task
VideoXum [18] (TMM'23)	open domain	14,001	video, text (video description)	10	ground-truth video summaries, text description of the video	video summarization with multimodal output
MrHiSum [7] (NeurIPS'23)	3,509	31,892	video	1	frame-level importance scores	video summarization & highlight detection
MMSum [19] (CVPR'24)	17	5,100	video, text, transcripts video metadata	1	ground-truth video and text summary	multimodal summarization with multimodal output
S-VideoXum [6] (ACM MM'25)	open domain	11,908	video, text (summary script)	10	ground-truth video summaries, text descriptions of the summaries	script-driven video summarization
SM-VideoXum (Ours)	open domain	11,908	video, text (summary script), transcripts	10	ground-truth video summaries, text descriptions of the summaries	script-driven multimodal video summarization
SM-MrHiSum (Ours)	3,509	29,917	video, text (summary script), transcripts	1	ground-truth video summaries, text descriptions of the summaries	script-driven multimodal video summarization

ground-truth pseudo-summaries obtained by prompting a Large Language Model (LLM) to extract the most informative moments from ASR transcripts. Jung Lee et al. [25] described a method that initially generates frame-level captions with the help of a frozen Multimodal LLM (LLaVA-1.5-7B [26]), and estimates the importance of each frame using a frozen LLM (Llama-2-13B-chat [27]). Then, taking into account the captions within a local temporal window around the frame, it refines the estimated frame-level importance scores using a global self-attention mechanism that considers the entire video context.

Furthermore, there are a few multimodal summarization approaches that aim to generate both visual and textual summaries. In this context, Fu et al. [28] presented a method that integrates a jump-attention mechanism to align features extracted from transcripts and video frames, and trained it using multi-task learning to simultaneously optimize text and video summarization. He et al. [29] built the A2Summ method, which aligns and attends multimodal inputs leveraging time correspondence using an alignment-guided self-attention mechanism; the latter learns how to form a keyframe-based and a text-based summary with the help of dual contrastive losses. Finally, Qui et al. [19] used a hierarchy of cross-modal attention mechanisms to fuse visual features from video frames/fragments with textual features from audio transcripts, and create a video and a text summary using a frame scorer and a text encoder, respectively.

The methods presented above produce generic summaries that aim to provide a complete synopsis of the entire video, thus not being tailored to specific needs about the summaries' content. Contrary to these methods, SD-MVSum takes into account such needs through the user-provided script, thus being capable to produce more personalized video summaries that are aligned with the users' demands.

2.3. Video summarization datasets

As discussed in [6], most of the existing datasets for text/query-driven video summarization are either very small and cover a restricted set of domains (UT Egocentric [1], TV Episodes [30], QFVS [2], SumMe [31], TVSum [32], ARS [33]), or contain annotations based on a small set of short-form (one/two-word) queries (RAD [8]). To tackle data scarcity, a few large-scale datasets for video summarization have been introduced in the literature over the last years. For example, the VideoXum dataset for cross-modal video summarization [18] comprises $14,001$ open-domain videos up to 12.5 min. long (2 min. avg. duration) with diverse visual content, derived from the ActivityNet Captions dataset [13]. Each video is accompanied by 10 ground-truth video summaries (in the form of binary frame-level annotations, indicating the inclusion (label 1) or not (label 0) of a frame in the summary), obtained by 40 different human annotators; and, a set of dense video captions that provide a high-level description of the full-length video. The MrHiSum dataset for video highlight detection and summarization [7] includes $31,892$ videos up to 5 min. long (3.3 min. avg. duration), derived from the YouTube-8M dataset [34]. Each video is associated with a series of frame-level importance scores (the so-called highlight labels in [7]) that have been computed after aggregating the viewing preferences of at least $50,000$ viewers of the video on YouTube, and used to formulate the ground-truth video summary based on the Knapsack algorithm and a predefined time-budget about the summary duration. The MMSum dataset for multimodal summarization and thumbnail generation [19], contains $5,100$ videos up to 115 min. long (14.5 min. avg. duration), showing various everyday activities from 17 main categories (e.g., cooking, sports, hobbies, travel). Each full-length video is accompanied by a ground-truth video and textual summary, as well as other metadata, such as title, author and category. Other large-scale datasets in the literature, such as MMS [35], MSMO [36], How2 [37], VMSMO [38] and MM-AVS [28], contain ground-

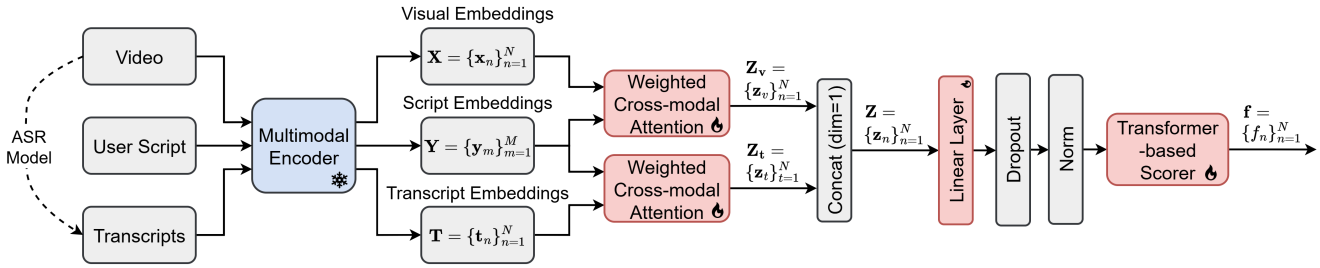

Figure 1. Overview of the SD-MVSum network architecture. Given an input video, a user script about the content of the summary, and a set of audio transcripts, SD-MVSum produces a video summary by finding associations of the user script with both the visual and the spoken content in the video, using two weighted cross-modal attention mechanisms. The outputs of these mechanisms are concatenated and forwarded to a trainable Transformer-based scorer which computes frame-level importance scores. These scores are used by a frame/fragment selection component that forms the video summary given a video fragmentation and a time-budget about the summary duration.

truth annotations that are suitable for training/evaluating methods that generate only textual summaries of the original video, and thus are out of the scope of this work. As shown in Table 1, none of the datasets mentioned above provides the necessary data for training and evaluating script-driven video summarization methods. The only existing large-scale dataset for this task is S-VideoXum [6], an extension of VideoXum which contains $11,908$ videos and 10 different ground-truth summaries (binary frame-level annotations) and summary descriptions (the so-called scripts in [6]) per video. The available triplets of "video, summary, and summary description" can be used for training methods to produce different summaries for a given video, driven by the descriptions of the desired content of each summary.

In this work, we extend the S-VideoXum and MrHiSum datasets by generating and publicly releasing textual descriptions of the human-annotated summaries (for MrHiSum; updated such scripts are also produced for S-VideoXum) and audio transcripts (for both S-VideoXum and MrHiSum). In this way, we make these extended datasets, called SM-VideoXum and SM-MrHiSum, suitable for training and evaluation of script-driven multimodal video summarization methods that take into account both the visual and the spoken content of the video.

3. Proposed SD-MVSum Method

3.1. Problem statement

Let us consider a full-length video and a user script (composed of a number of sentences) outlining the content of the desired video summary. Different sentences of the script may refer to different parts of the full-length video with varying visual and semantic content. The goal of script-driven multimodal video summarization is to assess the relevance of the user script with both the visual and the spoken content of the video, and select the video frames/fragments that are semantically associated to one or more sentences of the user script and necessary for providing a complete synopsis of the video. The selected frames/fragments must form a concise video summary with a duration that is typically set to 15% of the full-length video's duration [39].

3.2. Network architecture

An overview of the SD-MVSum network architecture is provided in Fig. 1. Let us assume a video of N frames (after sampling one frame per second), a user script outlining the content of the desired video summary formed by M sentences, and a set of automatically extracted audio transcripts containing K timestamped sentences. All these different input data pass through a pretrained multimodal encoder which produces three different sets of embeddings of the same size D; i.e., a set of visual embeddings ($\mathbf{X} = \{\mathbf{x}_n\}_{n=1}^N$), a set of script embeddings ($\mathbf{Y} = \{\mathbf{y}_m\}_{m=1}^M$), and a set of transcript embeddings ($\tilde{\mathbf{T}} = \{\mathbf{t}_k\}_{k=1}^K$). The transcript embeddings are expanded according to the timestamps of the associated transcripts, such that each embedding is repeated as many times as needed to match the number of video frames it spans, forming a new set $\mathbf{T} = \{\mathbf{t}_n\}_{n=1}^N$ that has the same number of embeddings with $\mathbf{X}$. This step is necessary for enabling, later on, the concatenation of the cross-modal embeddings that will come out the weighted cross-modal attention mechanisms.

The script embeddings $\mathbf{Y}$ are fused with the visual embeddings $\mathbf{X}$ and transcript embeddings $\mathbf{T}$, via two weighted cross-modal attention mechanisms. The latter explicitly exploit the semantic similarity between a pair of data modalities, when modeling their dependence and forming the cross-modal embeddings. The concatenation of these embeddings ($\mathbf{Z_v} = \{\mathbf{z}_v\}_{n=1}^N$ and $\mathbf{Z_t} = \{\mathbf{z}_t\}_{t=1}^N$) is then performed, along the feature dimension, resulting in an overall set of cross-modal embeddings $\mathbf{Z} = \{\mathbf{z}_n\}_{n=1}^N$ with size $2D$, which are subsequently reduced in size by half, using a linear layer.

The embeddings obtained after dimensionality reduction pass through dropout and normalization layers, and are then given as input to a trainable Transformer-based

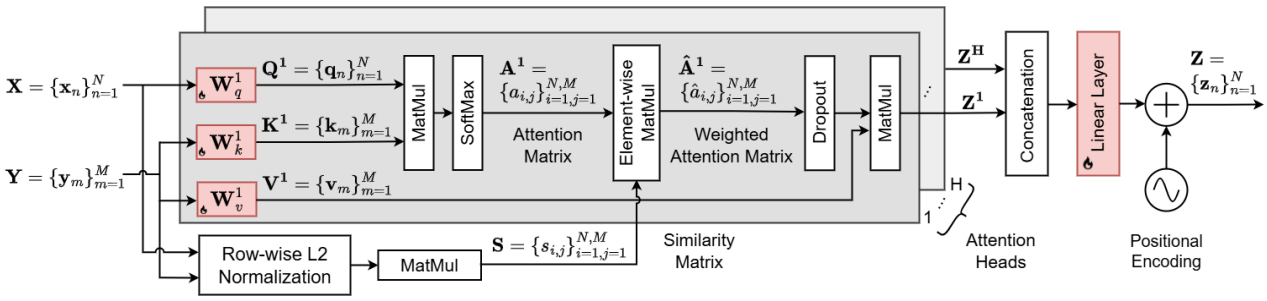

Figure 2. The processing pipeline in the weighted cross-modal attention mechanism when fusing the visual and the script embeddings. The dynamic scaling of the attention weights is performed based on the computed cosine similarity matrix of the input embeddings.

scorer, which computes frame-level importance scores $\mathbf{f} = \{f_n\}_{n=1}^{N}$. These scores are finally used by a frame/fragment selection component that assembles the final summary, given a predefined temporal fragmentation of the full-length video and a fixed time-budget about the summary duration. As a note, despite the fact that SD-MVSum currently leverages the aforementioned data modalities, its design makes it easily extensible for taking into account additional modalities (e.g., dense video captions). Data from new modalities can be incorporated by introducing extra cross-modal attention mechanisms and adjusting accordingly the input size of the employed linear layer for dimensionality reduction.

3.3. Weighted cross-modal attention mechanism

The processing pipeline of fusing the visual and script embeddings, within the introduced weighted cross-modal attention mechanism, is depicted in Fig. 2. The same process, after replacing $\mathbf{X}$ with $\mathbf{T}$, is applied when fusing the script and transcript embeddings. So, given the hth attention head of the attention mechanism, the visual embeddings $\boldsymbol{X}$ pass through a linear layer of size D/H, where H denotes the number of heads, forming the Query $\boldsymbol{Q}_h = \{\boldsymbol{q}_n\}_{n=1}^{N}$ matrix. The script embeddings $\boldsymbol{Y}$ pass through two different linear layers of size D/H, creating the Key $\boldsymbol{K}_h = \{\boldsymbol{k}_m\}_{m=1}^{M}$ and Value $\boldsymbol{V}_h = \{\boldsymbol{v}_m\}_{m=1}^{M}$ matrices. Then, the cross-modal embedding in the output of each attention head, is computed as follows:

$$\mathbf{A^h} = \mathbf{Q_h}\mathbf{K_h}^{\top}, \quad \hat{\mathbf{A}}^{\mathbf{h}} = \mathbf{A^h} \odot \mathbf{S},$$
$$\mathbf{Z_v^h} = Softmax(\hat{\mathbf{A}}^{\mathbf{h}})\mathbf{V^h}$$

where $\mathbf{A^h}$ is the initially computed attention matrix, and $\hat{\mathbf{A}}^{\mathbf{h}}$ is the weighted attention matrix after an element-wise multiplication (denoted by $\odot$) with $\mathbf{S}$, a cosine similarity matrix that is calculated by:

$$\mathbf{X_n} = L2(\mathbf{X}), \quad \mathbf{Y_n} = L2(\mathbf{Y})$$
$$\mathbf{S} = \mathbf{X_n}\mathbf{Y_n}^{T}$$

with $L2(\cdot)$ denoting L2 row-wise normalization. The output of the overall (multi-head) weighted cross-modal attention mechanism, is finally formulated as:

$$\mathbf{Z_v} = Concat(\mathbf{Z_v^1}, \mathbf{Z_v^2}, \ldots, \mathbf{Z_v^H})\mathbf{W}^o + pe,$$

where pe is the applied absolute positional encoding, that is implemented with the help of sine and cosine functions oscillating at different frequencies, as in [40].

So, instead of using a fixed scaling factor when computing the attention matrix (that is usually set equal to $\sqrt{D}$, following [40]), the proposed weighted cross-modal attention mechanism performs a dynamic scaling of the attention weights using the similarity matrix $\mathbf{S}$. Since the values in this matrix lie within $[-1, +1]$, our attention mechanism adaptively scales each entry of the attention matrix; values near ± 1 emphasize strongly-correlated elements in the common embedding space, while values near 0 suppress the weakly-related ones. Such an element-wise attention scaling approach provides finer control compared to uniform normalization, yielding more informative attention patterns.

4. Proposed SM-VideoXum and SM-MrHiSum Datasets

4.1. Extended datasets construction

The processing pipeline that was executed for constructing the extended SM-VideoXum and SM-MrHiSum datasets is presented in Fig. 3. As shown in the upper part of this figure, each ground-truth summary is submitted to a frame sampling process that keeps one frame per second; the set of sampled frames is then given as input to a video-to-text component. To exploit the visual content understanding and description capacity of modern Multimodal Large Language Models, we replaced the LLaVA-NeXT-Video-7B model [41] that was used to this end in [6] with the more powerful Qwen3-VL-8B-Instruct [42]. Moreover, based on experimentation with numerous prompts (5 in total) and through the qualitative analysis of the produced scripts, we replaced the prompt used to this end in [6] (*"describe the important scenes in this video"*) with the following (more specific) one: *"describe the scenery and the main persons*

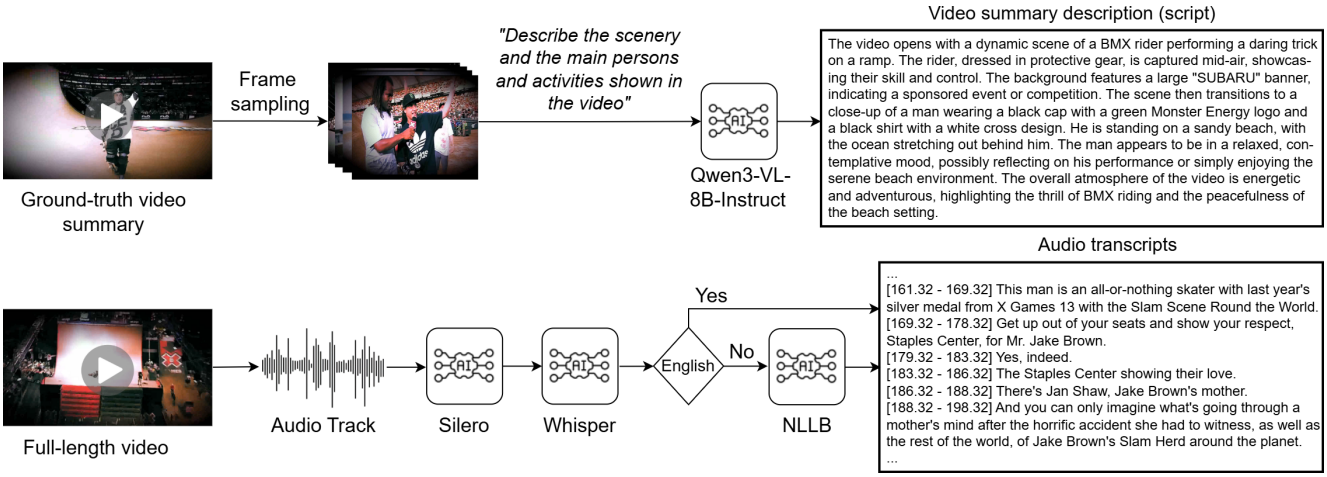

Figure 3. Overview of the processing pipeline for creating the SM-VideoXum and SM-MrHiSum datasets for script-driven multimodal video summarization.

and activities shown in the video." As before, the newly used Qwen3-VL-8B-Instruct model is prompted to generate a textual description of the ground-truth summary that is up to 200 tokens long.

Focusing on the lower part of Fig. 3, each full-length video undergoes an audio transcript extraction process. For this, the audio stream of the video is submitted to a pretrained model of Silero VAD for voice activity detection [43], which identifies the speech segments. The identified segments are then forwarded to a pretrained model of Whisper Turbo for speech-to-text transcription [44], which outputs a set of timestamped transcripts. Finally, given that the employed multimodal encoder for obtaining embeddings from the input data has been trained on English textual data, any transcript in a different language is translated in English using the NLLB-200 model for machine translation [45]. All the generated data and the full-length videos of the SM-MrHiSum and SM-VideoXum datasets, along with the extracted embeddings from visual and textual (script, transcript) data and the used data splits in our experiments, are publicly-available at: `https://github.com/IDT-ITI/SD-MVSum`

4.2. Dataset quality assessment

To ensure the high quality of the constructed datasets, we started by carefully selecting the models that would be leveraged for this task, as described in the previous section. Qwen3-VL-8B-Instruct [42] was selected based on its substantially improved video understanding capabilities, compared to its ancestors (Qwen2.5-VL-7B and Qwen2-VL-7B), and its competitive performance in comparison to other SotA models (e.g. OpenAI GPT-5 nano [46]) on several video question answer (e.g., RealWorldQA [47], MMStar [48], SimpleVQA [49]) and video understanding benchmarks (e.g., Video-MME [50], MLVU [51], VideoMMMU [52], MMVU [53]). Silero VAD [43] offers very high accuracy and is robust to environmental noise, surpassing other popular open-source solutions, such as Google's WebRTC[1] [54]. OpenAI's Whisper Turbo [44] is a SotA open model for multilingual automatic speech recognition, designed to balance high accuracy with significantly improved speed for speech-to-text transcription, that surpasses the performance of its ancestor WhisperX [55] and performs comparably with licensed solutions, such as Deepgram's Nova-2[2]. Meta AI's NLLB-200 model for multilingual machine translation [45] is a SotA model with superior performance on the FLORES-200 dataset[3], and performs well even for low-resource languages that were previously under-served by LLMs (e.g., GPT-4 [56], LLaMA-3 [57], Qwen-2.5 [58]).

To evaluate the quality of the generated scripts, we performed a visual inspection of the obtained scripts for a set of sampled videos from both the SM-VideoXum and SM-MrHiSum datasets. Our evaluation indicated that the scripts are sensible, highly descriptive, can refer to multiple related or unrelated actions, may contain abstract cues and narrative elements, and correspond to the content of the relevant video summaries. These characteristics are illustrated in the example in Fig. 4, where the first script relates to a video summary that focuses on the group of people kayaking, the second script relates to a video summary focusing on one of the rock climbers, and the third script relates to a video summary showing the efforts of two climbers and parts of the scenery. Moreover, the scripts in the extended SM-VideoXum and SM-MrHiSum datasets are diverse, varying from concise to extensive ones, as shown by the statistics

[1] https://github.com/reedom/VoiceActivityDetector
[2] https://deepgram.com/learn/nova-2-speech-to-text-api
[3] https://huggingface.co/datasets/Muennighoff/flores200

SM-VideoXum summary script #4: *"The video opens with a breathtaking view of a natural rock archway, framed by towering cliffs and lush greenery. The archway is illuminated by a soft, golden light, creating a serene and mystical atmosphere. The scene transitions to a group of people kayaking through a narrow, rocky channel. The kayakers are paddling in unison, their red paddles cutting through the calm, reflective water. The surrounding cliffs are steep and rugged, with patches of green vegetation clinging to their surfaces. The water below is a deep, tranquil blue, mirroring the sky above. The kayakers appear to be enjoying the peaceful journey, their movements synchronized as they navigate the winding path. The overall scene is one of natural beauty and adventure, capturing the essence of exploration and the awe-inspiring power of nature."*

SM-VideoXum summary script #7: *"The video opens with a blue screen displaying the "Lonely Planet" logo. The scene then transitions to a man climbing a steep, rocky cliff face. He is shirtless, wearing gray shorts, and is using his hands and feet to grip the rock as he ascends. The cliff is rugged and dark, with patches of green vegetation growing in the crevices. The man appears focused and determined as he continues his climb, his body moving with strength and precision. The camera follows his progress, capturing the height and difficulty of the climb. The overall atmosphere is one of adventure and challenge, as the climber pushes himself to reach new heights."*

SM-VideoXum summary script #9: *"The video opens with a breathtaking view of a towering cliff face, its rugged surface adorned with patches of greenery. A lone climber, dressed in a red shirt and dark pants, is seen scaling the cliff with determination. The camera pans to reveal another climber, this one in a blue shirt and dark pants, also making their way up the steep rock face. Both climbers are focused and skilled, their movements deliberate as they navigate the challenging terrain. The scene is set against a backdrop of a serene, misty landscape, with the sound of distant water adding to the tranquil yet adventurous atmosphere. The video captures the essence of outdoor adventure and the thrill of conquering nature's most formidable obstacles."*

Figure 4. A keyframe-based representation of the video "v_d3crFny-e3E", and examples of the generated scripts for three ground-truth summaries of this video from the SM-VideoXum dataset.

Table 2. Statistics about the scripts in the extended SM-VideoXum and SM-MrHiSum datasets.

	Min	Max	Avg	Std
SM-VideoXum				
# words per script	41	187	147.0	26.4
# sentences per script	3	25	8.5	1.8
SM-MrHiSum				
# words per script	32	191	148.0	25.4
# sentences per script	1	23	8.6	1.8

Table 3. Statistics about the audio transcripts in the extended SM-VideoXum and SM-MrHiSum datasets.

Spoken content statistics				
Dataset	# videos total	# videos w/o speech	% videos w/o speech	
SM-VideoXum	11,908	3,893	32.7%	
SM-MrHiSum	29,917	5,638	18.8%	
For videos with speech, # sentences per transcript				
Dataset	Min	Max	Avg	Std
SM-VideoXum	1	226	21.6	20.1
SM-MrHiSum	4	310	45.4	31.1

(number of words and sentences per script) reported in Table 2. In the case of SM-VideoXum, we additionally calculated such statistics for each video across its ten available ground-truth summaries, and then averaged over the entire dataset. The averaged std of script words and sentences are 17.8 and 1.3, respectively, indicating that within-video diversity is also quite high.

To assess the quality of the ASR transcripts, we made a thorough listening test, taking into account the transcripts of 25 randomly sampled videos from the SM-VideoXum dataset and computing the Word Error Rate (WER). The sampled videos included varying visual content from different categories, such as instructional videos (e.g., changing a tyre, playing the violin, waxing a ski), TV shows and movies, music videos, and sports videos (skiing, climbing). As a note, in most videos there was some background noise or music that made the task of automated transcription challenging. Nevertheless, our qualitative analysis documented the efficiency of the employed Whisper Turbo model. The computed WER score for the set of 25 sampled videos was equal to 0.082, indicating the reasonably high faithfulness and quality of the ASR transcripts [59]. Table 3 provides details about the number of videos without speech in the

extended SM-VideoXum and SM-MrHiSum datasets, and statistics that document the diversity of the obtained transcripts. As explained in Section 5.2, for videos without spoken content we use transcript embeddings with zero values.

5. Experiments

5.1. Evaluation protocol

We follow a slightly different evaluation approach on each dataset, to account for the differences in the available ground-truth annotations for each of them.

For SM-VideoXum, based on the methodology in [18], we form the video summary by selecting the top-15% scoring frames by the model, and quantify the similarity between the machine-generated and the ground-truth summary using the F-Score (%). So, a given test video is matched with each one of the multiple available user scripts for it, and each one of the generated summaries is compared with the corresponding ground-truth summary. Through this process, we compute an F-Score for each pair of compared summaries and we average these scores to form the final F-Score for this video. After performing this for all test videos of SM-VideoXum, we calculate the mean of the obtained F-Score values, forming a score that indicates the model's performance on the test set.

For SM-MrHiSum, we follow the evaluation strategy in [7] and formulate the video summary based on the machine-computed frame-level importance scores, a predefined temporal segmentation of the video, and a time-budget about the summary duration ($\leq 15\%$ of the video's length), by solving the Knapsack problem. Then, we quantify its similarity with the ground-truth summary using F-Score (%) only once, since there is only one ground-truth summary per video. After performing this for all test videos of SM-MrHiSum, we average the obtained F-Score values, resulting in a score that indicates the model's performance on the test set. Moreover, since SM-MrHiSum contains ground-truth annotations also in the form of frame-level scores, we additionally apply the evaluation protocol of [62]. Specifically, we quantify the alignment between the machine-computed and the ground-truth frame-level importance scores for a given video using the Kendall's τ [63] and Spearman's ρ [64] rank correlation coefficients. The computed τ and ρ values for all test videos are then averaged, defining the performance of the summarization model on the test set. When assessing the performance on generic video summarization, the same evaluation protocol of [62] is applied also in the case of SM-VideoXum, after forming a single ground-truth summary per video by averaging its multiple binary ground-truth annotations at the frame-level.

Both SM-MrHiSum and SM-VideoXum are divided into training, validation and test sets. During training, model selection is performed using the validation set, i.e., by measuring the model's performance on the validation set after each training epoch. When training is completed, we keep the model with the highest validation-set performance, and assess it on the test set using the evaluation protocols described above.

5.2. Implementation details

Similarly to [18] and [7], videos are sampled at one frame per second, and embeddings (of size $D = 512$) are obtained from the video frames, the user script and the audio transcripts, using the CLIP vision-language model. In the case of videos without spoken content, we use transcript embeddings with zero values. For the samples of SM-VideoXum, we employ a fine-tuned CLIP model on the data of VideoXum, that has been released by the authors of [18][4], while for the samples of SM-MrHiSum we use the CLIP ViT-B/32 model from HuggingFace[5].

Each cross-modal attention mechanism of SD-MVSum contains 8 heads. The frame scorer consists of a Transformer encoder, followed by a linear layer with 512 neurons and a sigmoid activation to compute frame-level importance scores. The network's weights are initialized based on the Xavier uniform initialization approach (gain $= \sqrt{2}$, bias $= 0.1$). Training on SM-VideoXum is based on the optimization of the BCE (Binary Cross-Entropy) loss between the predicted frame-level scores and the binary ground-truth labels, since this dataset does not include frame-level importance scores. Training on SM-MrHiSum is performed using the MSE (Mean Squared Error) loss, and the ground-truth frame-level importance scores. Training takes place for 50 epochs in a batch mode with a batch size equal to 4 and 64 for SM-VideoXum and SM-MrHiSum respectively, using the Adam optimizer and setting the learning rate, dropout rate and L2 regularization factor equal to $5 \cdot 10^{-5}$, 0.5 and 10^{-4}, respectively. All experiments were conducted on a workstation equipped with an Intel Core i5-11600K CPU and an NVIDIA RTX 3090 GPU. To allow the reproduction of our experiments, any used data and the PyTorch implementation of SD-MVSum have been made publicly-available at: `https://github.com/IDT-ITI/SD-MVSum`

5.3. Experimental comparisons and ablations

We compared the proposed SD-MVSum method against a number of SotA methods for query/script-driven and generic (multimodal- or visual-based) video summarization. For the first class, we considered the SD-VSum [6] and CLIP-It [3] methods that were discussed in Section 2. For the second class, we took into account the A2Summ [29] method for multimodal video summarization that also

[4] https://videoxum.github.io/

[5] https://huggingface.co/sentence-transformers/clip-ViT-B-32

Table 4. Performance comparisons with SotA methods for script-driven (upper part) and generic (lower part) video summarization on SM-VideoXum and SM-MrHiSum, in terms of F-Score (%, denoted "F1") and Kendall's τ and Spearman's ρ rank correlation coefficients. Best scores in bold, second-best scores underlined. Scores whose difference from those of SD-MVSum is not statistically significant are shown in italics.

		Data modalities			SM-VideoXum			SM-MrHiSum		
Task	Model	Script	Visual	Transcript	F1	τ	ρ	F1	τ	ρ
Script-driven summ.	SD-MVSum (proposed)	✓	✓	✓	**27.3**	N/A	N/A	**59.3**	**0.204**	0.273
	SD-VSum [6] (ACM MM'25)	✓	✓	X	24.4	N/A	N/A	57.9	0.175	0.236
	CLIP-It [3] (NeurIPS'21)	✓	✓	X	22.8	N/A	N/A	56.3	0.120	0.169
Generic summ.	A2Summ [29] (CVPR'23)	X	✓	✓	21.3	0.145	0.193	58.0	0.169	0.239
	CSTA [60] (CVPR'24)	X	✓	X	23.5	**0.176**	**0.233**	57.7	*0.193*	***0.274***
	PGL-SUM [61] (IEEE ISM'21)	X	✓	X	22.1	0.153	0.203	57.4	0.168	0.241

utilizes the audio transcripts, and two visual-based methods with SotA performance on video summarization benchmarks, namely the CSTA [60] and PGL-SUM [61] methods.

The results of our evaluations are reported in Table 4. The score differences between the proposed SD-MVSum method and the compared ones were tested in terms of statistical significance using a two-tailed paired T-Test with confidence level $\alpha = 0.05$, and were found to be statistically significant, with a few exceptions; for these exceptions, the corresponding compared method's score is shown in italics in Table 4. As a note, this statistical significance testing is not applicable when comparing the proposed SD-MVSum with generic summarization methods on SM-VideoXum, due to the difference in the number of generated video summaries (generic summarization methods produce one tenth of the summaries that SD-MVSum and other script-driven methods produce, for this dataset). Nevertheless, these performance differences (SD-MVSum vs. generic summarization methods) are particularly pronounced, and also all considered generic summarization methods perform worse than SD-VSum, where the F1 difference between SD-MVSum and SD-VSum is statistically significant.

The comparison between script-driven video summarization methods in Table 4 showcases that SD-MVSum outperforms the other competing methods. This is attributed to the use of multiple embeddings for representing the script (as discussed in [6]), and algorithmic improvements over SD-VSum [6], i.e., the dynamic scaling of the attention weights based on the computed cosine similarity matrix of the input embeddings, and the exploitation of the audio transcripts. These improvements led to statistically significant higher performance in both datasets, and according to all measures, compared to [6], [3]. Further comparison with methods for generic summarization in Table 4 indicates the ability of SD-MVSum to produce video summaries that are more tailored to the users' needs. SD-MVSum outperforms all generic summarization methods on both datasets in terms of F-Score, and performs comparably with CSTA on SM-MrHiSum in terms of τ and ρ.

Following, we conducted a series of ablation studies in order to examine: i) the contribution of each of the key concepts of SD-MVSum, namely, the use of audio transcripts as an auxiliary data source and the introduction of weighted cross-modal attention for modeling dependencies among different data modalities; ii) the impact of using fewer or more heads in the weighted cross-modal attention mechanisms; iii) the influence of smaller or larger batches of training data; and iv) the use of other data fusion approaches for combining the output of the integrated weighted cross-modal attention mechanisms. As before, the differences between SD-MVSum and its variants were tested in terms of statistical significance using the same approach, and were found to be statistically significant, with a few exceptions; for these exceptions, the corresponding variant's score is shown in italics. Moreover, we should stress that in Tables 5-9 reporting the results of these ablations, we always compare script-driven summarization approaches, thus for the SM-VideoXum dataset the evaluation protocol of [62] involving the calculation of the Kendall's τ and Spearman's ρ rank correlation coefficients is not applicable, as explained in Section 5.1.

In our first ablation study, that aims to assess the contribution of the key concepts of SD-MVSum, we considered the following variants of SD-MVSum:

- **Variant #1** does not take into account the audio transcripts, and thus performs script-driven video summarization using only the visual content of the video.
- **Variant #2** does not apply the proposed dynamic scaling of attention weights and follows a more straightforward data fusion approach, similarly to SD-VSum.

The outcomes of this study, presented in Table 5, document the positive contribution of both of the aforementioned key concepts. More specifically, the removal of audio tran-

Table 5. Performance comparison with variants of SD-MVSum on SM-VideoXum and SM-MrHiSum, in terms of F-Score (%, denoted "F1") and Kendall's τ and Spearman's ρ rank correlation coefficients. Best scores in bold. Scores whose difference from those of SD-MVSum is not statistically significant are shown in italics.

		Data modalities				SM-VideoXum	SM-MrHiSum		
Task	Model	Script	Visual	Transcript	Scaling	F1	F1	τ	ρ
Script-driven summ.	SD-MVSum (proposed)	✓	✓	✓	✓	**27.3**	**59.3**	**0.204**	**0.273**
	Variant #1	✓	✓	X	✓	26.6	58.3	0.169	0.230
	Variant #2	✓	✓	✓	X	25.1	58.4	*0.195*	*0.265*

Table 6. Performance of SD-MVSum on SM-VideoXum and SM-MrHiSum in terms of F-Score (%) and Kendall's τ and Spearman's ρ rank correlation coefficients, for different numbers of attention heads. Best scores in bold. The differences between the scores attained using the proposed number of attention heads and all other examined options are statistically significant.

	SM-VideoXum	SM-MrHiSum		
Att. heads	F1	F1	τ	ρ
4	25.7	58.2	0.181	0.243
8 (prop.)	**27.3**	**59.3**	**0.204**	**0.273**
16	*27.2*	58.4	0.185	0.251

scripts from the pool of input data (Variant #1) leads to a consistent drop in the script-driven video summarization performance across both datasets and according to all measures (being more pronounced on SM-MrHiSum), pointing out the usefulness of audio transcripts when used as an auxiliary source of information. Moreover, scaling the computed attention weights by the utilized cross-modal attention mechanisms using a fixed value - instead of performing dynamic scaling with the help of matrix $\mathbf{S}$ - leads to a similar performance drop in terms of F-Score on SM-MrHiSum, i.e., close to 1%, and an even higher drop on SM-VideoXum, namely $> 2\%$ (the observed small differences in τ, ρ were not found to be statistically significant). Such a finding demonstrates the strong contribution of the proposed weighted cross-modal attention mechanism in finding better dependencies among data from different modalities.

In our second ablation study, we examined the effect of using fewer or more attention heads in the weighted cross-modal attention mechanisms. In particular, starting from the number of heads in SD-VSum (8 in total), we measured the performance of SD-MVSum when using half and double of them. The results, reported in Table 6, indicate that the use of 8-head weighted cross-modal attention mechanisms is the optimal choice for both datasets, and based on both of the employed evaluation protocols.

In our third ablation study, we investigated the effect of using smaller or larger batches of training data. For SM-VideoXum, we started with a batch size equal to the one in [6] (4 videos, which means 40 video-script pairs) and then

Table 7. Performance of SD-MVSum on SM-VideoXum in terms of F-Score (%), for different training data batch sizes. Best scores in bold. The differences between the scores attained using the proposed batch size and all other examined batch sizes are statistically significant.

	SM-VideoXum
Batch size	F1
2	26.7
4 (prop.)	**27.3**
8	26.4

Table 8. Performance of SD-MVSum on SM-MrHiSum in terms of F-Score (%) and Kendall's τ and Spearman's ρ rank correlation coefficients, for different training data batch sizes. Best scores in bold. Scores, whose difference from those attained using the proposed batch size is not statistically significant, are shown in italics.

	SM-MrHiSum		
Batch size	F1	τ	ρ
32	58.4	0.176	0.237
64 (prop.)	**59.3**	**0.204**	**0.273**
128	58.7	*0.195*	*0.263*

we halved and doubled it. For SM-MrHiSum, we started with a similarly large batch size (64 video-script pairs), and we again halved and doubled it. The outcomes of this study, presented in Tables 7 and 8, show that medium-sized batches of training data lead to the best performance on both datasets.

In our fourth ablation study, we tried other approaches for fusing the embeddings in the output of the integrated cross-modal attention mechanisms. As alternatives to the proposed concatenation and the subsequent linear layer for dimensionality reduction (see Fig. 1), we examined the use of mean pooling or max pooling. The findings, shown in Table 9, demonstrate that the employed data fusion approach is clearly more suitable compared to the tested alternatives, on both datasets.

Concerning dataset construction, we conducted an ablation study to investigate the impact of using scripts generated by different SotA Multimodal LLMs. More specifi-

Table 9. Performance of SD-MVSum on SM-VideoXum and SM-MrHiSum in terms of F-Score (%) and Kendall's τ and Spearman's ρ rank correlation coefficients, for different data fusion approaches. Best scores in bold. Scores, whose difference from those attained using the proposed fusion approach is not statistically significant, are shown in italics.

	SM-VideoXum	SM-MrHiSum		
Data fusion	F1	F1	τ	ρ
Mean pooling	26.4	58.5	0.182	0.247
Max pooling	26.2	*59.0*	0.184	0.248
Concat. (prop.)	**27.3**	**59.3**	**0.204**	**0.273**

cally, we trained and evaluated SD-MVSum using scripts obtained by:

- Prompting LLaVA-NeXT-Video-7B to "describe the important scenes in this video" (prompt #1), as in [6].
- Prompting Qwen3-VL-8B-Instruct to "describe the important scenes in this video" (prompt #1).
- Prompting Qwen3-VL-8B-Instruct to "describe the scenery and the main persons and activities shown in the video" (prompt #2).

The findings are presented in Table 10. A comparison between the scripts obtained by LLaVA-Next-Video-7B and Qwen3-VL-8B-Instruct for the same prompt (prompt #1), documents the competency of the latter model to create scripts that are more useful for learning the task of script-driven multimodal video summarization, a finding that is more pronounced on the videos of the SM-VideoXum dataset. In addition, prompting the Qwen3-VL-8B-Instruct model with a more detailed and specific prompt for generating these scripts (prompt #2) leads to further measurable advancements in the summarization performance on both datasets, according to all measures. The observed differences on SM-MrHiSum after modifying the prompt to the Qwen3-VL-8B-Instruct model were not found to be statistically significant; however, the differences when using LLaVA-Next-Video-7B (as in [6]) for generating the scripts were statistically significant, according to the attained rank correlation coefficients. These findings highlight the added value of the released datasets.

5.4. Qualitative analysis

To further evaluate the contribution of audio transcripts in the script-driven video summarization outcome, we performed a qualitative analysis that was based on manual observation of the generated summaries by our SD-MVSum method and the SD-VSum method [6] that uses just the visual content of the video, for a set of sampled videos from the SM-VideoXum and SM-MrHiSum datasets. One of the examined samples is presented in Fig. 5. The upper part provides a keyframe-based representation of the content of the full-length video, and the tabular structure beneath shows the utilized input data and the generated video summary by each method. As can be seen, both methods focused on parts of the video presenting cheer-leading routines (either during training or at a competition) and ignored less relevant parts showing e.g., interviews, thus being aligned with the user script. However, SD-VSum puts emphasis on parts of the video showing the group's training in an indoor area (choosing 7 relevant video fragments) and focuses less on parts of the video presenting the team's participation at the competition (selecting 2 relevant video fragments). On the contrary, SD-MVSum produces a more comprehensive summary, showing parts from the training process (including 4 relevant video fragments), capturing the essence of the competition stage (using 2 relevant fragments), and presenting the team's performance during the competition (selecting 3 relevant fragments). This example, demonstrates that the use of audio transcripts allowed SD-MVSum to spot more effectively video parts showing the team's preparation and participation at a competition, and generate a summary that is more aligned with the viewer's needs, as indicated by the significantly higher F-Score.

Another example from our qualitative analysis using a travel VLOG from YouTube (available at: https://www.youtube.com/watch?v=F2RLoK16U4k) is presented in Fig. 6. The upper part provides a keyframe-based representation of the content of the full-length video, and the tabular structure beneath shows the utilized input data and the generated video summary by the SD-MVSum and SD-VSum methods. As can be seen, both methods provide a complete synopsis of the travel and focus on iconic places and monuments of the city. However, SD-MVSum pays more attention to video parts showing local dishes and pastries, producing a video summary that is more tailored to the user script.

6. Conclusions

In this paper, we presented the SD-MVSum method for script-driven multimodal video summarization, which takes into consideration the relevance of the user-provided script with both the visual and the spoken content in the video. This relevance is modeled using a new weighted cross-modal attention mechanism, which exploits the semantic similarity between paired modalities and applies a dynamic scaling to promote the most relevant video parts to the user's script. To assist the training and evaluation of script-driven multimodal video summarization methods, we extended two large-scale datasets for video summarization (S-VideoXum, MrHiSum) to make them suitable for the task. Our quantitative and qualitative evaluations showcased the competitiveness of SD-MVSum against other SotA methods for script-driven and generic video summarization.

Table 10. Performance of SD-MVSum on SM-VideoXum and SM-MrHiSum datasets in terms of F-Score (%) and Kendall's τ and Spearman's ρ rank correlation coefficients, for different script-generation approaches. Best scores in bold. Scores, whose difference from those attained using the proposed script generation model & prompt is not statistically significant, are shown in italics.

	SM-VideoXum	SM-MrHiSum		
Script-generation	F1	F1	τ	ρ
LLaVA-Next-Video-7B & prompt #1	25.2	*58.8*	0.187	0.253
Qwen3-VL-8B-Instruct & prompt #1	27.0	*58.9*	*0.199*	*0.269*
Qwen3-VL-8B-Instruct & prompt #2 (proposed)	**27.3**	**59.3**	**0.204**	**0.273**

Keyframe-based representation of the video content

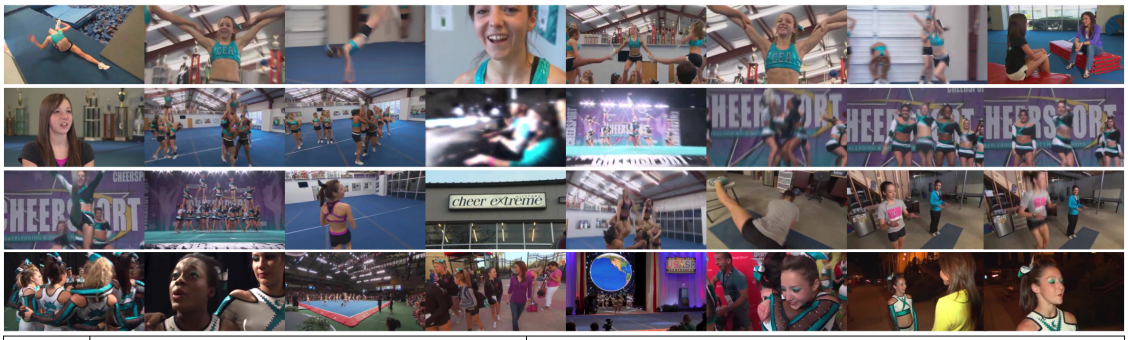

Method	Input	Generated Summary
SD-MVSum	**Video frames** **User script:** *"The video opens with a close-up of a young cheerleader in a blue and white uniform, smiling and waving to the camera. She is standing in front of a backdrop of trophies and awards, indicating a competitive or celebratory setting. The scene then transitions to a wider shot of a cheerleading competition stage. A group of cheerleaders, dressed in matching blue and white uniforms, are performing a routine on a blue mat. The stage is decorated with large banners and a prominent globe display with the year "2011" illuminated on it. The cheerleaders are executing synchronized movements, including jumps, arm gestures, and formations. The audience, visible in the foreground, is watching intently, with some individuals holding cameras to capture the performance. The overall atmosphere is energetic and festive, with bright lighting and a sense of excitement and celebration."* **Audio transcripts**: [...] [54:50] The teams performed two and a half minute routines of high flying stunts, powerful tumbling and impressive teamwork. [65:36] Mikayla trains at Cheer Extreme in Raleigh, North Carolina. [72:04] Aside from countless hours of training year-round with her team, Mikayla works out on her own with gruelling exercises like this. [...]	 **F-Score = 50.7%**
SD-VSum	**Video frames** **User script**: *"The video opens with a close-up of a young cheerleader in a blue and white uniform, smiling and waving to the camera. She is standing in front of a backdrop of trophies and awards, indicating a competitive or celebratory setting. The scene then transitions to a wider shot of a cheerleading competition stage. A group of cheerleaders, dressed in matching blue and white uniforms, are performing a routine on a blue mat. The stage is decorated with large banners and a prominent globe display with the year "2011" illuminated on it. The cheerleaders are executing synchronized movements, including jumps, arm gestures, and formations. The audience, visible in the foreground, is watching intently, with some individuals holding cameras to capture the performance. The overall atmosphere is energetic and festive, with bright lighting and a sense of excitement and celebration."*	 **F-Score = 26.9%**

Figure 5. An indicative sample from our qualitative analysis. The upper part provides a keyframe-based representation of the content of the full-length video, and the tabular structure beneath shows the utilized input data and the generated video summary by each method.

Keyframe-based representation of the video content

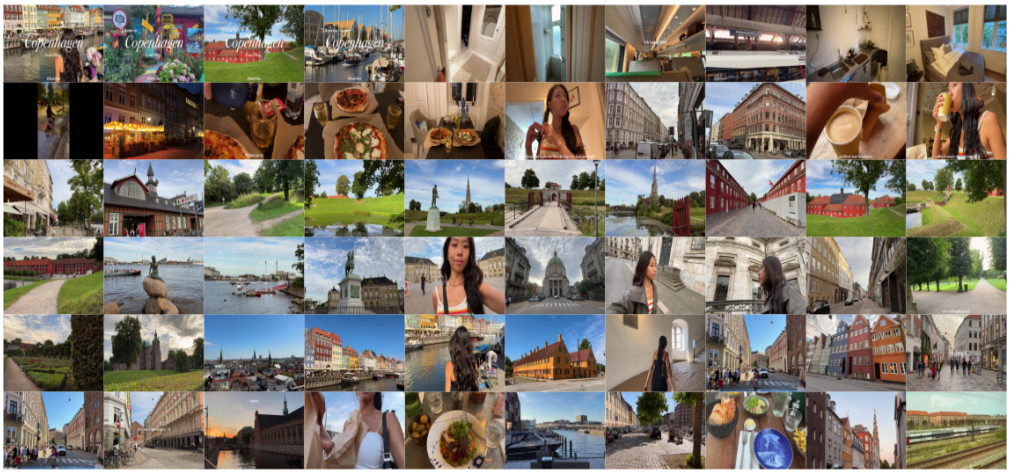

Method	Input	Generated Summary
SD-MVSum	• **Video frames** • **User script**: *"The video showcases the city's cultural heritage through iconic monuments and statues. It also highlights its traditional cuisine with local dishes and pastries."* • **Audio transcripts**: [...] **[01:30]** From there, we headed to one of Copenhagen's most iconic landmarks—The Little Mermaid statue. **[01:50]** Next, I visited Amalienborg, the royal palace, where I admired the equestrian statue of King Frederik V in the center of the square. **[03:10]** For dinner, we went to Restaurant Puk, where I tried some hearty Danish food. [...]	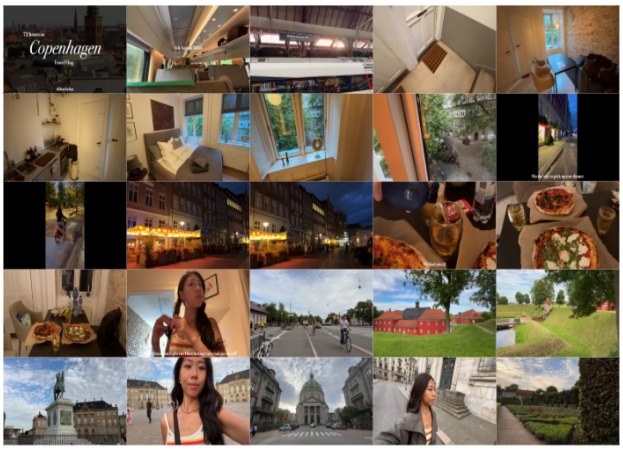
SD-VSum	• **Video frames** • **User script**: *"The video showcases the city's cultural heritage through iconic monuments and statues. It also highlights its traditional cuisine with local dishes and pastries."*	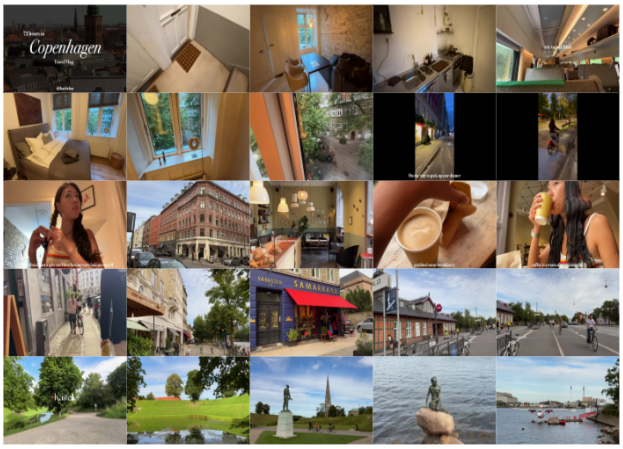

Figure 6. An indicative sample from our qualitative analysis. The upper part provides a keyframe-based representation of the content of the full-length video, and the tabular structure beneath shows the utilized input data and the generated video summary by the SD-MVSum and SD-VSum methods.

References

[1] A. Sharghi, B. Gong, and M. Shah, "Query-focused extractive video summarization," in *Computer Vision – ECCV 2016*, B. Leibe, J. Matas, N. Sebe, and M. Welling, Eds. Cham: Springer International Publishing, 2016, pp. 3–19. 1, 2, 3

[2] A. Sharghi, J. S. Laurel, and B. Gong, "Query-focused video summarization: Dataset, evaluation, and a memory network based approach," in *2017 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2017, pp. 2127–2136. 2, 3

[3] M. Narasimhan, A. Rohrbach, and T. Darrell, "CLIP-It! language-guided video summarization," in *Proc. of the 35th International Conference on Neural Information Processing Systems*, ser. NIPS '21. Red Hook, NY, USA: Curran Associates Inc., 2021. 1, 2, 8, 9

[4] H. Wei, B. Ni, Y. Yan, H. Yu, and X. Yang, "Video summarization via semantic attended networks," in *Proc. of the Thirty-Second AAAI Conference on Artificial Intelligence and Thirtieth Innovative Applications of Artificial Intelligence Conference and Eighth AAAI Symposium on Educational Advances in Artificial Intelligence*, ser. AAAI'18/IAAI'18/EAAI'18. AAAI Press, 2018. 1, 2

[5] J.-H. Huang and M. Worring, "Query-controllable video summarization," in *Proc. of the 2020 International Conference on Multimedia Retrieval*, ser. ICMR '20. New York, NY, USA: Association for Computing Machinery, 2020, p. 242–250. [Online]. Available: https://doi.org/10.1145/3372278.3390695 1

[6] M. Mylonas, E. Apostolidis, and V. Mezaris, "SD-VSum: A method and dataset for script-driven video summarization," in *Proc. of the 33rd ACM International Conference on Multimedia*, ser. MM '25. New York, NY, USA: Association for Computing Machinery, 2025, p. 6596–6604. [Online]. Available: https://doi.org/10.1145/3746027.3755821 1, 2, 3, 4, 5, 8, 9, 10, 11

[7] J. Sul, J. Han, and J. Lee, "Mr. HiSum: a large-scale dataset for video highlight detection and summarization," in *Proc. of the 37th International Conference on Neural Information Processing Systems*, ser. NIPS '23. Red Hook, NY, USA: Curran Associates Inc., 2023. 1, 3, 8

[8] A. B. Vasudevan, M. Gygli, A. Volokitin, and L. Van Gool, "Query-adaptive video summarization via quality-aware relevance estimation," in *Proc. of the 25th ACM International Conference on Multimedia*, ser. MM '17. New York, NY, USA: Association for Computing Machinery, 2017, p. 582–590. [Online]. Available: https://doi.org/10.1145/3123266.3123297 2, 3

[9] Y. Zhang, M. C. Kampffmeyer, X. Liang, M. Tan, and E. P. Xing, "Query-Conditioned Three-Player Adversarial Network for Video Summarization," in *Proc. of the 2018 British Machine Vision Conference (BMVC)*, 2018. 2

[10] P. Jiang and Y. Han, "Hierarchical variational network for user-diversified & query-focused video summarization," in *Proc. of the 2019 on International Conference on Multimedia Retrieval*, ser. ICMR '19. New York, NY, USA: Association for Computing Machinery, 2019, p. 202–206. [Online]. Available: https://doi.org/10.1145/3323873.3325040 2

[11] S. Xiao, Z. Zhao, Z. Zhang, X. Yan, and M. Yang, "Convolutional hierarchical attention network for query-focused video summarization," *Proc. of the AAAI Conference on Artificial Intelligence*, vol. 34, no. 07, pp. 12 426–12 433, Apr. 2020. [Online]. Available: https://ojs.aaai.org/index.php/AAAI/article/view/6929 2

[12] S. Xiao, Z. Zhao, Z. Zhang, Z. Guan, and D. Cai, "Query-biased self-attentive network for query-focused video summarization," *IEEE Transactions on Image Processing*, vol. 29, pp. 5889–5899, 2020. 2

[13] R. Krishna, K. Hata, F. Ren, L. Fei-Fei, and J. C. Niebles, "Dense-captioning events in videos," in *2017 IEEE International Conference on Computer Vision (ICCV)*, 2017, pp. 706–715. 2, 3

[14] G. Mujtaba, A. Malik, and E.-S. Ryu, "Ltc-sum: Lightweight client-driven personalized video summarization framework using 2d cnn," *IEEE Access*, vol. 10, pp. 103 041–103 055, 2022. 2

[15] J.-H. Huang, L. Murn, M. Mrak, and M. Worring, "Query-based video summarization with pseudo label supervision," in *2023 IEEE International Conference on Image Processing (ICIP)*, 2023, pp. 1430–1434. 2

[16] Y. Guo, J. Xing, X. Hou, S. Xin, J. Jiang, D. Terzopoulos, C. Jiang, and Y. Liu, "Cfsum: A transformer-based multimodal video summarization framework with coarse-fine fusion," in *ICASSP 2025 - 2025 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 2025, pp. 1–5. 2

[17] M. Rodrigo, C. Cuevas, and N. García, "Text-guided sports highlights: A clip-based framework for automatic video summarization," *IEEE Access*, vol. 13, pp. 210 996–211 009, 2025. 2

[18] J. Lin, H. Hua, M. Chen, Y. Li, J. Hsiao, C. Ho, and J. Luo, "Videoxum: Cross-modal visual and textural summarization of videos," *IEEE Transactions on Multimedia*, vol. 26, pp. 5548–5560, 2024. 3, 8

[19] J. Qiu, J. Zhu, W. Han, A. Kumar, K. Mittal, C. Jin, Z. Yang, L. Li, J. Wang, D. Zhao, B. Li, and L. Wang, " MMSum: A Dataset for Multimodal Summarization and Thumbnail Generation of Videos ," in *2024 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*. Los Alamitos, CA, USA: IEEE Computer Society, Jun. 2024, pp. 21 909–21 921. [Online]. Available: https://doi.ieeecomputersociety.org/10.1109/CVPR52733.2024.02069 3

[20] M. Narasimhan, A. Nagrani, C. Sun, M. Rubinstein, T. Darrell, A. Rohrbach, and C. Schmid, "TL;DW? Summarizing instructional videos with task relevance and cross-modal saliency," in *Computer Vision – ECCV 2022*, S. Avidan, G. Brostow, M. Cissé, G. M. Farinella, and T. Hassner, Eds. Cham: Springer Nature Switzerland, 2022, pp. 540–557. 2

[21] S. Palaskar, J. Libovický, S. Gella, and F. Metze, "Multimodal abstractive summarization for how2 videos," in *Proc. of the 57th Annual Meeting of the Association for Computational Linguistics*, A. Korhonen, D. Traum, and L. Màrquez,

Eds. Florence, Italy: Association for Computational Linguistics, Jul. 2019, pp. 6587–6596. [Online]. Available: https://aclanthology.org/P19-1659/ 2

[22] M. Sanabria, F. Precioso, and T. Menguy, "Hierarchical multimodal attention for deep video summarization," in *2020 25th International Conference on Pattern Recognition (ICPR)*, 2021, pp. 7977–7984. 2

[23] S.-H. Zhong, J. Lin, J. Lu, A. Fares, and T. Ren, "Deep semantic and attentive network for unsupervised video summarization," *ACM Trans. Multimedia Comput. Commun. Appl.*, vol. 18, no. 2, Feb. 2022. [Online]. Available: https://doi.org/10.1145/3477538 2

[24] D. M. Argaw, S. Yoon, F. C. Heilbron, H. Deilamsalehy, T. Bui, Z. Wang, F. Dernoncourt, and J. S. Chung, "Scaling Up Video Summarization Pretraining with Large Language Models," in *2024 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*. Los Alamitos, CA, USA: IEEE Computer Society, Jun. 2024, pp. 8332–8341. [Online]. Available: https://doi.ieeecomputersociety.org/10.1109/CVPR52733.2024.00796 2

[25] M. J. Lee, D. Gong, and M. Cho, "Video summarization with large language models," in *2025 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2025, pp. 18 981–18 991. 3

[26] H. Liu, C. Li, Q. Wu, and Y. J. Lee, "Visual instruction tuning," in *Advances in Neural Information Processing Systems*, A. Oh, T. Naumann, A. Globerson, K. Saenko, M. Hardt, and S. Levine, Eds., vol. 36. Curran Associates, Inc., 2023, pp. 34 892–34 916. 3

[27] H. Touvron, L. Martin, K. Stone, P. Albert, A. Almahairi, and et al., "Llama 2: Open foundation and fine-tuned chat models," 2023. [Online]. Available: https://arxiv.org/abs/2307.09288 3

[28] X. Fu, J. Wang, and Z. Yang, "MM-AVS: A full-scale dataset for multi-modal summarization," in *Proc. of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, K. Toutanova, A. Rumshisky, L. Zettlemoyer, D. Hakkani-Tur, I. Beltagy, S. Bethard, R. Cotterell, T. Chakraborty, and Y. Zhou, Eds. Online: Association for Computational Linguistics, Jun. 2021, pp. 5922–5926. [Online]. Available: https://aclanthology.org/2021.naacl-main.473/ 3

[29] B. He, J. Wang, J. Qiu, T. Bui, A. Shrivastava, and Z. Wang, " Align and Attend: Multimodal Summarization with Dual Contrastive Losses ," in *2023 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*. Los Alamitos, CA, USA: IEEE Computer Society, Jun. 2023, pp. 14 867–14 878. [Online]. Available: https://doi.ieeecomputersociety.org/10.1109/CVPR52729.2023.01428 3, 8, 9

[30] S. Yeung, A. Fathi, and L. Fei-Fei, "Videoset: Video summary evaluation through text," *ArXiv*, vol. abs/1406.5824, 2014. [Online]. Available: https://api.semanticscholar.org/CorpusID:9338736 3

[31] M. Gygli, H. Grabner, H. Riemenschneider, and L. Van Gool, "Creating summaries from user videos," in *Computer Vision – ECCV 2014*, D. Fleet, T. Pajdla, B. Schiele, and T. Tuytelaars, Eds. Cham: Springer International Publishing, 2014, pp. 505–520. 3

[32] Y. Song, J. Vallmitjana, A. Stent, and A. Jaimes, "Tvsum: Summarizing web videos using titles," in *2015 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2015, pp. 5179–5187. 3

[33] K. Cizmeciler, E. Erdem, and A. Erdem, "Leveraging semantic saliency maps for query-specific video summarization," *Multimedia Tools Appl.*, vol. 81, no. 12, p. 17457–17482, May 2022. [Online]. Available: https://doi.org/10.1007/s11042-022-12442-w 3

[34] S. Abu-El-Haija, N. Kothari, J. Lee, P. Natsev, G. Toderici, B. Varadarajan, and S. Vijayanarasimhan, "YouTube-8M: A large-scale video classification benchmark," *CoRR*, vol. abs/1609.08675, 2016. [Online]. Available: http://arxiv.org/abs/1609.08675 3

[35] H. Li, J. Zhu, C. Ma, J. Zhang, and C. Zong, "Multi-modal summarization for asynchronous collection of text, image, audio and video," in *Proc. of the 2017 Conference on Empirical Methods in Natural Language Processing*, M. Palmer, R. Hwa, and S. Riedel, Eds. Copenhagen, Denmark: Association for Computational Linguistics, Sep. 2017, pp. 1092–1102. [Online]. Available: https://aclanthology.org/D17-1114/ 3

[36] J. Zhu, H. Li, T. Liu, Y. Zhou, J. Zhang, and C. Zong, "MSMO: Multimodal summarization with multimodal output," in *Proc. of the 2018 Conference on Empirical Methods in Natural Language Processing*, E. Riloff, D. Chiang, J. Hockenmaier, and J. Tsujii, Eds. Brussels, Belgium: Association for Computational Linguistics, Oct.-Nov. 2018, pp. 4154–4164. [Online]. Available: https://aclanthology.org/D18-1448/ 3

[37] R. Sanabria, O. Caglayan, S. Palaskar, D. Elliott, L. Barrault, L. Specia, and F. Metze, "How2: A Large-scale Dataset for Multimodal Language Understanding," in *Proc. of Visually Grounded Interaction and Language (ViGIL), NeurIPS*, Montréal, Canada, 2018. [Online]. Available: https://hal.science/hal-02431947 3

[38] M. Li, X. Chen, S. Gao, Z. Chan, D. Zhao, and R. Yan, "VMSMO: Learning to generate multimodal summary for video-based news articles," in *Proc. of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, B. Webber, T. Cohn, Y. He, and Y. Liu, Eds. Online: Association for Computational Linguistics, Nov. 2020, pp. 9360–9369. [Online]. Available: https://aclanthology.org/2020.emnlp-main.752/ 3

[39] E. Apostolidis, E. Adamantidou, A. I. Metsai, V. Mezaris, and I. Patras, "Video summarization using deep neural networks: A survey," *Proceedings of the IEEE*, vol. 109, no. 11, pp. 1838–1863, 2021. 4

[40] A. Vaswani, N. Shazeer, N. Parmar, J. Uszkoreit, L. Jones, A. N. Gomez, Ł. Kaiser, and I. Polosukhin, "Attention is all you need," *Advances in neural information processing systems*, vol. 30, 2017. 5

[41] B. Li, K. Zhang, H. Zhang, D. Guo, R. Zhang, F. Li, Y. Zhang, Z. Liu, and C. Li, "Llava-next: Stronger llms supercharge multimodal capabilities in the wild," May 2024. [Online]. Available: https://llava-vl.github.io/blog/2024-05-10-llava-next-stronger-llms/ 5

[42] Qwen3-VL-Team, "Qwen3-VL Technical Report," *arXiv preprint arXiv:2511.21631*, 2025. 5, 6

[43] Silero-Team, "Silero VAD: pre-trained enterprise-grade voice activity detector (vad), number detector and language classifier," https://github.com/snakers4/silero-vad, 2024. 6

[44] A. Radford, J. W. Kim, T. Xu, G. Brockman, C. McLeavey, and I. Sutskever, "Robust speech recognition via large-scale weak supervision," in *Proc. of the 40th International Conference on Machine Learning*, ser. ICML'23. JMLR.org, 2023. 6

[45] NLLB-Team, "No language left behind: Scaling human-centered machine translation," 2022. [Online]. Available: https://arxiv.org/abs/2207.04672 6

[46] A. Singh, A. Fry, A. Perelman, A. Tart, A. Ganesh, and et al., "OpenAI GPT-5 System Card," 2025. [Online]. Available: https://arxiv.org/abs/2601.03267 6

[47] xAI, "RealWorldQA: a benchmark for real-world spatial understanding," https://huggingface.co/datasets/xai-org/RealworldQA, 2024. 6

[48] L. Chen, J. Li, X. Dong, P. Zhang, Y. Zang, Z. Chen, H. Duan, J. Wang, Y. Qiao, D. Lin, and F. Zhao, "Are we on the right way for evaluating large vision-language models?" in *Proc. of the 38th International Conference on Neural Information Processing Systems*, ser. NIPS '24. Red Hook, NY, USA: Curran Associates Inc., 2024. 6

[49] X. Cheng, W. Zhang, S. Zhang, J. Yang *et al.*, "SimpleVQA: multimodal factuality evaluation for multimodal large language models," in *2025 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2025, pp. 4637–4646. 6

[50] C. Fu, Y. Dai, Y. Luo, L. Li, S. Ren, R. Zhang, Z. Wang, C. Zhou, Y. Shen, M. Zhang, P. Chen, Y. Li, S. Lin, S. Zhao, K. Li, T. Xu, X. Zheng, E. Chen, C. Shan, R. He, and X. Sun, "Video-MME: The first-ever comprehensive evaluation benchmark of multi-modal llms in video analysis," in *2025 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2025, pp. 24 108–24 118. 6

[51] J. Zhou, Y. Shu, B. Zhao, B. Wu, Z. Liang, S. Xiao, M. Qin, X. Yang, Y. Xiong, B. Zhang, T. Huang, and Z. Liu, "MLVU: Benchmarking multi-task long video understanding," in *2025 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2025, pp. 13 691–13 701. 6

[52] K. Hu, P. Wu, F. Pu, W. Xiao, Y. Zhang, X. Yue, B. Li, and Z. Liu, "Video-MMMU: Evaluating knowledge acquisition from multi-discipline professional videos," 2025. [Online]. Available: https://arxiv.org/abs/2501.13826 6

[53] Y. Zhao, H. Zhang, L. Xie, T. Hu, G. Gan, Y. Long, Z. Hu, W. Chen, C. Li, Z. Xu, C. Wang, Z. Shangguan, Z. Liang, Y. Liu, C. Zhao, and A. Cohan, "MMVU: Measuring expert-level multi-discipline video understanding," in *2025 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2025, pp. 8475–8489. 6

[54] M. McKinnon, S. Khaki, C. K. Reddy, and W. Huang, "Window size versus accuracy experiments in voice activity detectors," 2026. [Online]. Available: https://arxiv.org/abs/2601.17270 6

[55] M. Bain, J. Huh, T. Han, and A. Zisserman, "Whisperx: Time-accurate speech transcription of long-form audio," *INTERSPEECH 2023*, 2023. 6

[56] OpenAI, "GPT-4 technical report," 2024. [Online]. Available: https://arxiv.org/abs/2303.08774 6

[57] AI@Meta, "Llama 3 model card," 2024. [Online]. Available: https://github.com/meta-llama/llama3/blob/main/MODEL_CARD.md 6

[58] Qwen, "Qwen2.5 technical report," 2025. [Online]. Available: https://arxiv.org/abs/2412.15115 6

[59] S. Arif, A. J. Khan, M. Abbas, A. A. Raza, and A. Athar, "WER we stand: Benchmarking Urdu ASR models," in *Proc. of the 31st International Conference on Computational Linguistics*, O. Rambow, L. Wanner, M. Apidianaki, H. Al-Khalifa, B. D. Eugenio, and S. Schockaert, Eds. Abu Dhabi, UAE: Association for Computational Linguistics, Jan. 2025, pp. 5952–5961. [Online]. Available: https://aclanthology.org/2025.coling-main.397/ 7

[60] J. Son, J. Park, and K. Kim, " CSTA: CNN-based Spatiotemporal Attention for Video Summarization ," in *2024 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*. Los Alamitos, CA, USA: IEEE Computer Society, Jun. 2024, pp. 18 847–18 856. [Online]. Available: https://doi.ieeecomputersociety.org/10.1109/CVPR52733.2024.01783 9

[61] E. Apostolidis, G. Balaouras, V. Mezaris, and I. Patras, "Combining global and local attention with positional encoding for video summarization," in *2021 IEEE International Symposium on Multimedia (ISM)*, December 2021, pp. 226–234. 9

[62] M. Otani, Y. Nakashima, E. Rahtu, and J. Heikkilä, "Rethinking the evaluation of video summaries," in *2019 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2019, pp. 7588–7596. 8, 9

[63] M. G. Kendall, "The treatment of ties in ranking problems," *Biometrika*, vol. 33, no. 3, pp. 239–251, 1945. 8

[64] S. Kokoska and D. Zwillinger, *CRC standard probability and statistics tables and formulae*. Crc Press, 2000. 8