

Minimal Clips, Maximum Saliency: Long Video Summarization via Key Moment Extraction

Galann Pennec^{∞, ◇, ♥} Zhengyuan Liu^{◇, ♥}

Nicholas Asher^{§, ♥} Philippe Muller^{∞, ♥} Nancy F. Chen^{◇, ♥}

[∞]IRIT, University of Toulouse, France

[◇]Agency for Science, Technology and Research (A*STAR), Singapore

[♥]CNRS@CREATE, Singapore [§]CNRS, IRIT, France

galann.pennec@cnrscnrsatcreate.sg, {liu_zhengyuan, nancy_chen}@a-star.edu.sg

{nicholas.asher, philippe.muller}@irit.fr

Abstract

Vision-Language Models (VLMs) are able to process increasingly longer videos. Yet, important visual information is easily lost throughout the entire context and missed by VLMs. Also, it is important to design tools that enable cost-effective analysis of lengthy video content. In this paper, we propose a clip selection method that targets key video moments to be included in a multimodal summary. We divide the video into short clips and generate compact visual descriptions of each using a lightweight video captioning model. These are then passed to a large language model (LLM), which selects the K clips containing the most relevant visual information for a multimodal summary. We evaluate our approach on reference clips for the task, automatically derived from full human-annotated screenplays and summaries in the MovieSum dataset. We further show that these reference clips (less than 6% of the movie) are sufficient to build a complete multimodal summary of the movies in MovieSum. Using our clip selection method, we achieve a summarization performance close to that of these reference clips while capturing substantially more relevant video information than random clip selection. Importantly, we maintain low computational cost by relying on a lightweight captioning model.

1 Introduction

Vision-Language Models (VLMs) (Bai et al., 2025; Wang et al., 2025; OpenAI, 2024) have demonstrated improved capabilities in processing longer videos, particularly due to efficient pretraining (Li et al., 2024; Weng et al., 2024; Xue et al., 2024; Zhang et al., 2024a; Wei et al., 2025).

However, performing inference on hour-long videos is costly and questions remain about how effectively VLMs handle longer contexts (Fu et al., 2024; Wang et al., 2024a; Zhou et al., 2024; Mangalam et al., 2023). Notably, important visual elements are sometimes lost throughout the video,

often causing VLMs to neglect or completely omit crucial information (Pennec et al., 2025; Zhang et al., 2024b; Nishimura et al., 2024; Shen et al., 2024; Park et al., 2024a).

By observing that not all information in a video is relevant to a task, some strategies maintain a memory over past visual information when processing longer videos (Song et al., 2024; Qian et al., 2024; He et al., 2024; Balazevic et al., 2024; Kahatapitiya et al., 2024). Similarly, in Long Video Understanding, the answer to a question about a video is usually contained within a small subset of key frames retrieved by video content selection methods (Park et al., 2024b; Wang et al., 2024b; Narasimhan et al., 2021).

To the best of our knowledge, most of the above video content selection approaches have been designed for the vision modality alone with a limited focus on multimodal data where different modalities often overlap. Also, video content selection has been widely studied for Long Video Question Answering (LVQA) leaving Multimodal Video Summarization underexplored (Pennec et al., 2025).

In this paper, we make the observation that videos are often highly redundant across modalities, for instance, when what is shown visually is already conveyed through the dialogue or transcripts. We therefore consider the task of visually salient clip selection, meaning that we extract all clips containing relevant visual information that cannot be inferred from the transcripts alone.

We propose a cost-effective clip selection method (as shown in Figure 1) and apply it to multimodal summarization of long videos such as movies from MovieSum (Saxena and Keller, 2024)¹ which offer a reliable testbed due to their rich narratives, diverse multimodal cues, and their need for cross-modal integration.

Unlike LVQA that can assign confidence scores

¹<https://huggingface.co/datasets/rohitsaxena/MovieSum>

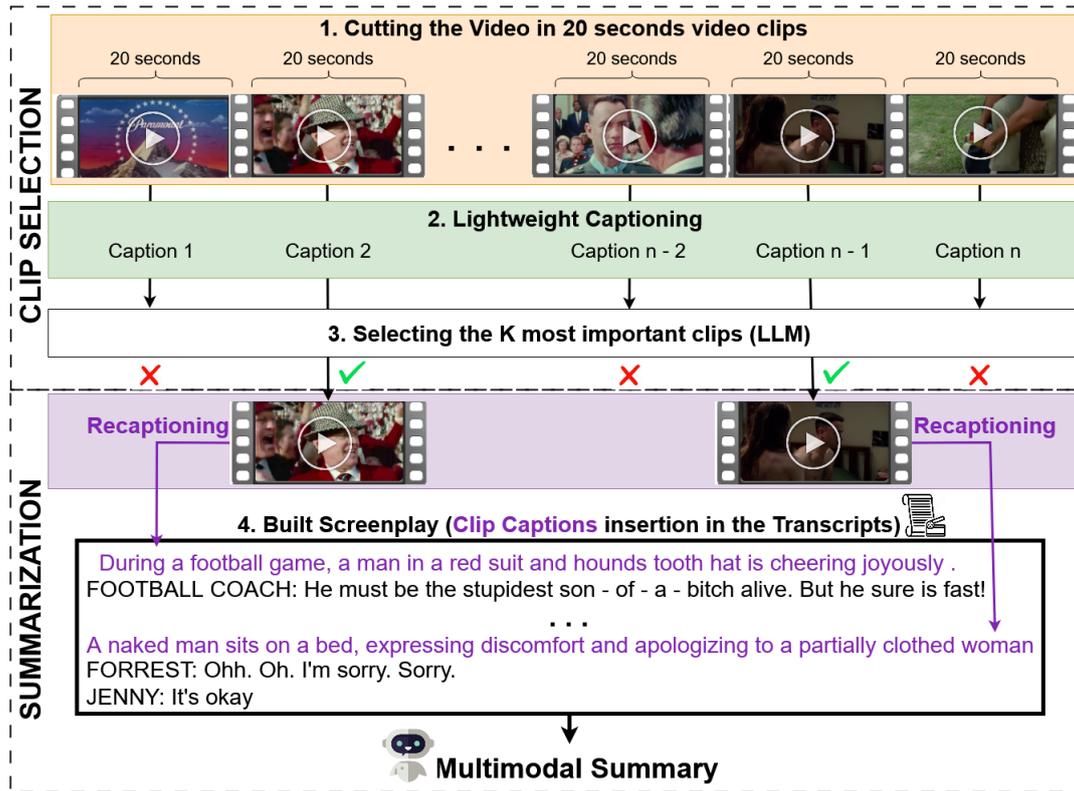


Figure 1: **Our Clip Selection followed by Summarization.** 1) We segment the video into 20-second clips and generate lightweight captions for each. We then feed all the clip captions to an LLM to identify the top K clips that contain visually important information. 2) For summarization, we build a screenplay-like document by inserting the captions of the selected clips into the transcripts at the correct timestamps. We finally summarize these screenplays.

to individual frames or small frame sets, video summarization needs an understanding of the full context to identify key moments. To preserve temporal information, we treat clips, rather than individual frames, as the basic unit (Zhi et al., 2025). Moreover, instead of formulating the task as a binary classification over frames, we define it as selecting the top K most relevant clips from the entire video.

In Figure 1, we divide the video into 20-second clips and generate a caption for each of them using a lightweight captioning model. The resulting captions are then passed to a LLM, which selects the top K most important clips to form the basis of the final multimodal summary.

Our contributions are as follows:

- We propose visually salient clip selection as the task of retrieving all the video clips containing visual information relevant for a multimodal summary that cannot be inferred from the dialogue transcripts alone.
- We introduce a lightweight clip selection strategy (Figure 1) allowing us to retrieve and target key video moments to generate long video

summaries at a lower cost. We evaluate our approach based on reference clips for the task that we infer from MovieSum annotations.

- Using our clip selection strategy, we generate multimodal summaries of entire movies in MovieSum. Our summaries closely match those generated from the reference clips, while retrieving significantly more relevant visual information than random clip selection.

2 Related Work

From Text to Multimodal Summarization

Summarization has a long history in the text modality, evolving from early rule-based approaches to neural sequence-to-sequence models (See et al., 2017; Liu et al., 2021; Chen and Yang, 2020) and large language models (Wang et al., 2023; Tang et al., 2023). These methods have been applied to diverse domains ranging from news and dialogues to scientific articles (Feng et al., 2021; Liu and Chen, 2021; Yasunaga et al., 2019), while addressing challenges such as salient span selection (Kedzie et al., 2018; Liu et al., 2019), factual consistency (Nan et al., 2021; Kryściński et al.,

2020), and long-context understanding (Gehrmann et al., 2018; Kryscinski et al., 2022; Liu and Chen, 2022). With the development of multimodal approaches, the field has expanded beyond text to incorporate other modalities (Li et al., 2017). Recent work explores multimodal summarization, integrating textual, audio, speech, and video data to generate feature-rich summaries (Papalampidi and Lapata, 2023; Pennek et al., 2025).

Video Content Selection Identifying important content from long videos has been addressed mostly in LVQA. Most of the time, the question to answer is used to query and retrieve relevant information (usually frames) from the whole video, whether in a zero-shot setting (Huang et al., 2025; Park et al., 2024b; Wang et al., 2024c), through pre-training (Yu et al., 2023, 2025; Korbar et al., 2024) or via agentic approaches (Wang et al., 2024b; Yang et al., 2024b; Zhi et al., 2025).

Efficient Long Video Summarization Although a summary can take the form of a video, such as a TV show recap or movie trailer (Singh et al., 2024; Papalampidi et al., 2021; Chen et al., 2024), the present work generates long video summaries in text form instead. Existing approaches to the task uniformly sample frames or clips throughout the original video either at a fixed rate (Liu et al., 2025; Atri et al., 2021) or aligned with the scenes or dialogue utterances (Mahon and Lapata, 2024a; Papalampidi and Lapata, 2023). This uniform sampling results in inefficient video context management and VLMs easily missing out on valuable information by treating all video moments as equally important. Noticing the variability in the importance of video moments, some approaches adopt alternative clip selection strategies for the video-to-text summarization task. For instance, Pennek et al. (2025) retrieves all video clips without any dialogue, arguing that they correlate with key visual moments of a movie or TV show.

In this work, we propose a clip selection method that identifies in zero-shot the top K visually salient clips containing important information to include in a multimodal video summary, and study its impact on the end summary.

3 Clip Selection for Multimodal Video Summarization

We approach multimodal summarization in two steps, treating clip selection as an intermediate task

for summary generation. The complete pipeline is presented in Figure 1 and detailed in Section 3.1. In Section 3.2, we further explain how clip selection is evaluated using the gold screenplay and groundtruth summary of a movie.

3.1 Pipeline Overview

As shown in Figure 1, we first segment the video into 20-second clips. Each clip is captioned using a lightweight VLM, and the resulting captions are all passed to an LLM for selection of the K clips containing important visual information. Clip selection is performed in either zero-shot or two-shot settings, depending on the prompts provided in Appendix A.1.

Following Pennek et al. (2025); Mahon and Lapata (2024b), we build a screenplay-like document that efficiently represents the video’s multimodal content, by combining the dialogue transcripts together with visual descriptions, for later summarization. To do so, we recaption the selected clips using a second, more robust VLM and insert them into the transcripts at the proper timestamp. As in (Mahon and Lapata, 2024a), we could infer the timestamp of the transcripts utterances by aligning them with the corresponding audio in the video.

We finally summarize these screenplays using a customized prompt that incites the LLM to focus on multimodal cues from both the video captions and dialogue (see Appendix A.3 for prompt details). We also place the marker ‘Caption:’ at the beginning of every clip caption in the screenplay to further facilitate the identification of important video content by the LLM.

3.2 Clip Selection Reference

Given the human-written screenplay and a reference summary of a movie we can extract all clips containing important visual information for a good multimodal understanding. Those clips serve as a reference for the task of clip selection in our evaluation (section 4.2). We proceed in three steps as follows. The first two steps are performed by an LLM in zero-shot, given the prompts in Appendix A.4.

Step 1: Fact Identification We decompose the groundtruth summary into a list of all its facts, each fact conveying a single piece of information (roughly equivalent to a simple clause).

Step 2: Visual Fact Classification We classify each groundtruth summary fact as Visual (referring to the video) or Textual (referring to the dialogue). For each fact, we ask the LLM to

retrieve the information from within the human-written screenplay by specifically quoting the line. If the information comes from a clip caption in the screenplay, the fact is considered as `Visual`. If it comes from the dialogue between the characters, we instead classify the fact as `Textual`.

Step 3: Reference Clips For every `Visual` fact in **Step 2**, we locate the video segment that visually conveys this information. Using the screenplay timestamps, we define the clip to begin at the utterance immediately preceding the caption containing the `Visual` fact and to end at the utterance immediately following it.

4 Experimental Setup

4.1 Datasets

We conduct our experiments on MovieSum (Saxena and Keller, 2024), a summarization dataset of 2200 movies between 1950 and 2023, with equal splits of 200 movies each for validation and testing. The films span a diverse range of genres (comedy, drama, thriller, ...) and have an average runtime of two hours. It includes detailed summaries (635 words on average) referencing both video and dialogue modalities as well as long human-written screenplays (25K words on average). Structurally, these screenplays are documents that interweave the dialogue transcripts with corresponding visual descriptions. We report all our experiments on the test split for which we purchased the videos.

4.2 Clip Selection Metrics

Similar to previous work (Miech et al., 2019; Krishna et al., 2017; Lei et al., 2020), we evaluate clip selection performance using Recall@K. The Recall@K denotes the ratio of reference clips retrieved by a clip selection method when we fix the number of selected clips to K .

A reference clip r is deemed retrieved if the Intersection-over-Reference (IoR) between r and a predicted clip p is greater than τ , where τ is a fixed threshold. We define the IoR score as follows.

$$\text{IoR}(p, r) = \frac{|p \cap r|}{|r|}$$

where $|p \cap r|$ denotes the temporal intersection length between p and r .

4.3 Summarization Metrics

We report the summarization performance on both traditional and task-specific metrics.

Traditional Metrics We report ROUGE-1 (r_1), ROUGE-2 (r_2), and ROUGE-Lsum (r_{lsum}) using the `python-rouge` package, as well as METEOR scores computed with the `meteor_score` function from `nlk.translate`.

MFACTSUM We evaluate multimodal performance using MFACTSUM metric (Pennec et al., 2025), which measures how effectively a multimodal summary captures the relevant information from both the video and dialogue. The metric computes two components: visual fact recall, assessing visual understanding, and textual fact recall, assessing textual understanding. The final multimodal score, MFACTSUM, is obtained by averaging the two above components. Specifically, visual (resp. textual) fact recall refers here to the proportion of groundtruth summary facts originating from the video (resp. the dialogue) that are supported by the predicted summary.

MFACTSUM metrics rely on the following information: a decomposition of the groundtruth summary into facts, classified as `Visual` or `Textual`, and an assessment of whether these facts are supported by the predicted summary. The decomposition and classification can be done as described in section 3.2 by prompting an LLM in zero-shot. The final step of the evaluation uses the same LLM to judge if the predicted summary supports the `Visual` and `Textual` facts. The prompt is given in Appendix A.5.

4.4 Implementation Details

We generate screenplay summaries for all the 200 movies from MovieSum test split using our pipeline in Figure 1. We use either Qwen2.5-Omni-3B or Qwen2.5-Omni-7B as the lightweight captioning model, Gemini 2.5 Flash-Lite as the recaptioning model and Gemini 2.5 Flash for both the clip selection, summarization as well as for our evaluations with MFACTSUM. While we choose Qwen2.5-Omni (Xu et al., 2025) for its high accuracy at a lower cost, Gemini 2.5 Flash-Lite (Comanici et al., 2025) offers strong multimodal capabilities, making it well-suited for high quality recaptioning. We also discuss results when replacing the summarization LLM by either Gemini 1.5 Flash (Reid et al., 2024) or Qwen2.5-72B-Instruct (Yang et al., 2024a) in Appendix B. We disallow the thinking process and the use of external websites whenever using Gemini’s API² in all

²<https://aistudio.google.com/>

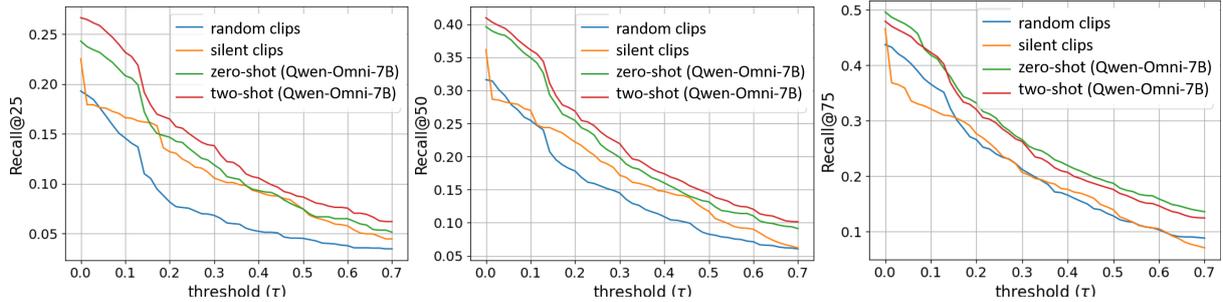


Figure 2: **Recall@K across varying thresholds τ** . Our clip selection outperforms other baselines regardless of the chosen threshold τ used for the IoR matching.

our experiments. Because of the high API costs, results are presented for a single run only.

We always fix the target summary length to 1000 words in the prompt to all our baselines and models (see Appendix A.3). Also, we truncate the output summary to 1000 words for fair comparison between all settings. We do so because we are aware that some summarization metrics including the visual and textual recall as well as MFACTSUM can increase mechanically with the summary length.

5 Clip Selection Experiments

5.1 Baseline Methods

In all our baselines, we fix K as the number of selected clips. We assign different values of K (25, 50, and 75) in practice, as shown in Figure 2.

Random Clips We randomly select K non-overlapping video clips of 20 seconds from the whole video.

Silent Clips All video clips that occur during a pause in the dialogue are considered (Pennec et al., 2025). Such clips are then sorted by decreasing duration and the K first are chosen. This heuristic baseline is motivated by the fact that silent scenes from a movie or TV show often highlight key visual moments and actions impacting the storyline.

Our Clips This corresponds to our clip selection method in either zero-shot or two-shot settings. In the main design (Figure 1), clip selection is performed by the LLM on captions generated by a lightweight captioning model (either Qwen2.5-Omni-3B or Qwen2.5-Omni-7B). For comparison, Table 1 also reports results, instead, on the gold screenplay captions. In the latter case, the LLM is given all the captions present in the gold screenplay and is prompted to select the K most visually relevant ones using the same prompt as in the main setting (Appendix A.1).

	R@25	R@50	R@75
random clips	6.83	14.50	21.22
silent clips	10.57	17.18	20.68
<i>ours zero-shot (Qwen2.5-Omni-3B)</i>	11.14	18.34	25.40
<i>ours two-shot (Qwen2.5-Omni-3B)</i>	10.82	19.85	25.63
<i>ours zero-shot (Qwen2.5-Omni-7B)</i>	11.89	19.89	26.60
<i>ours two-shot (Qwen2.5-Omni-7B)</i>	13.82	21.95	26.25
<hr/>			
<i>ours zero-shot (gold screenplay captions)</i>	39.56	51.22	66.79
<i>ours two-shot (gold screenplay captions)</i>	39.79	55.04	67.33

Table 1: **Evaluation of clip selection methods on the MovieSum test set.** We report the Recall@K (R@K) for all studied clip selection strategies relative to the reference clips from section 3.2. We also provide the scores when using the gold screenplays captions instead of Qwen2.5-Omni captions in our method. We fix the threshold τ to 0.3 in the Recall@K computation. Note that there is on average 354 clips of 20 seconds in a movie from the MovieSum dataset. Therefore, $K = 25$, $K = 50$ and $K = 75$ respectively corresponds, for our method, to retrieving about 7%, 14% and 28% of the total movie duration length.

5.2 Results

We report the Recall@K of various clip selection methods in both Figure 2 and Table 1. In Figure 2, our clip selection method (both zero-shot and two-shot) outperforms all tested baselines regardless of the chosen threshold τ for the IoR matching. Also, using the two-shot examples further improves the performance of our method for lower values of K ($K = 25$ and $K = 50$). For larger K ($K = 75$), the silent clip selection becomes less precise and performs close to the random selection baseline.

We also report the exact scores for different values of K when the threshold τ is fixed to 0.3 in Table 1. The results highlight that the quality of the captions plays an important role for the task. Using the gold screenplay captions instead of Qwen2.5-Omni captions significantly boosts the Recall@K. Similarly, using a larger captioning model like Qwen2.5-Omni-7B instead of Qwen2.5-Omni-3B

	vis-rec	text-rec	MFS	r1	r2	rlsum	METEOR
Transcripts (no clips)	14.42	26.89	20.65	44.66	10.35	42.64	32.12
Filtered Gold Screenplay (avg. 6% clips)	32.84	35.63	34.23	45.73	13.63	43.90	36.24
Gold Screenplay (all clips)	34.47	35.48	34.97	47.43	11.88	45.34	34.06

Table 2: **We only need about 6% of the video information present in the gold screenplay to approximate a multimodal summary of the entire movie.** The filtered gold screenplay is obtained by keeping only the captions for the reference clips, accounting for about 6% of all the captions. Evaluations are conducted on the MovieSum test set. We report the visual recall (vis-rec), textual recall (text-rec) and MFactSum denoted as MFS. We also include ROUGE-1 (r1), ROUGE-2 (r2), ROUGE-Lsum (rlsum) and METEOR. Best results are in **bold**.

consistently improves the performance of our clip selection approach.

5.3 Human Evaluation of the Clip Selection Reference

We validate the clip selection reference described in Section 3.2 through human evaluation conducted by two of the co-authors. The first annotator evaluated four randomly selected movies from the MovieSum test set: The Shining (1980), The Dark Knight (2008), The Imitation Game (2014), and Black Panther (2018). This evaluation covers 108 reference clips in total, providing a statistically meaningful sample for our analysis. We compute the agreement with the second annotator only on the movie The Dark Knight.

The human annotators are asked to watch each movie entirely and manually construct a human clip reference following the same procedure as in Section 3.2. More precisely, given the groundtruth summary facts identified during **Step 1**, the annotators retrieve all the video clips that support those facts (**Steps 2 & 3**). We find this step to have an accuracy of 84.6% between our two annotators.

On the four movies, our clip selection reference achieves an F1 score of 86.5% against the first human reference (see Appendix C).

6 Multimodal Video Summarization Experiments

6.1 Baseline Methods

We study the video summarization task under various settings, with results reported in Table 3. For every setting, we always use the same exact summarization prompt defined in Appendix A.3.

Transcripts (no video) We generate summaries from the transcripts alone.

Built Screenplay (K Clips) Following our pipeline (Figure 1), we build the screenplay from the K selected clips and generate a screenplay sum-

mary. We also replace the clip selection component in our pipeline with alternative clip selection baselines from Section 5.1 such as random clips or silent clips.

Built Screenplay (reference clips) We build the screenplay from the reference clips we identified in Section 3.2. We simply feed the reference clips directly into our summarization pipeline (Figure 1). This setting serves us as an upperbound as we inject the best possible clips into our pipeline.

Gold Screenplay We generate summaries from the screenplay annotations given in MovieSum.

6.2 Results

MovieSum summaries are highly multimodal We discover that a third of the summary content refers to video information. More precisely, we identify 35 Visual facts on average in MovieSum summaries.

6% of the video is enough for a complete movie summary Despite being highly multimodal, MovieSum summaries can be effectively built using a small fraction (6%) of the gold screenplay captions. Specifically, using only the captions from the reference clips (21 clips on average) provides a summary nearly as informative as one built from the full screenplay (Table 2). This finding strongly motivates the use of clip selection in long video summarization.

Our clip selection leads to summaries that better include multimodal information The summarization results in Table 3 show that our clip selection method outperforms the other clip selection baselines especially on the visual recall, textual recall and MFACTSUM metrics. In particular, we are able to retrieve substantially more relevant visual information (visual recall) than the random clip selection baseline. Remarkably, our performance is even close to that of the best possible clips (screenplay of the reference clips). As noted

	vis-rec	text-rec	MFS	r1	r2	rlsum	METEOR
Transcripts (no video)	14.42	26.89	20.65	44.66	10.35	42.64	32.12
Built Screenplay (25 clips)							
random clips	15.44	31.93	23.69	46.11	10.61	44.07	33.27
silent clips	16.33	32.61	24.47	46.24	10.95	44.11	33.24
<i>our clips zero-shot (Qwen2.5-Omni-7B)</i>	20.81	35.64	28.23	46.29	11.20	44.16	33.59
<i>our clips two-shot (Qwen2.5-Omni-7B)</i>	<u>21.05</u>	34.01	27.53	46.90	11.55	44.69	33.93
Built Screenplay (50 clips)							
random clips	15.68	30.91	23.29	46.11	10.51	43.88	33.21
silent clips	17.53	31.82	24.67	46.10	<u>11.32</u>	43.95	33.36
<i>our clips zero-shot (Qwen2.5-Omni-7B)</i>	20.97	33.63	27.30	<u>46.54</u>	11.22	<u>44.49</u>	33.63
<i>our clips two-shot (Qwen2.5-Omni-7B)</i>	20.60	<u>34.22</u>	27.41	46.22	10.97	44.10	33.46
Built Screenplay (75 clips)							
random clips	14.68	30.90	22.79	45.35	10.34	43.15	32.73
silent clips	19.43	32.42	25.92	46.00	10.84	43.83	33.37
<i>our clips zero-shot (Qwen2.5-Omni-7B)</i>	20.87	31.25	26.06	46.04	10.75	43.87	33.13
<i>our clips two-shot (Qwen2.5-Omni-7B)</i>	22.25	33.45	<u>27.85</u>	46.53	10.80	44.32	33.55
Built Screenplay (reference clips)	22.43	35.38	28.90	47.28	11.67	45.21	34.14
Gold Screenplay	34.47	35.48	34.97	47.43	11.88	45.34	34.06

Table 3: **Summarization results on the MovieSum test set.** Except for the gold screenplay, all built screenplays in the Table are produced using Gemini 2.5 Flash-Lite as the recaptioning model. We always use Gemini 2.5 Flash as the summarization model. Column descriptions are the same as in Table 2. Best results are in **bold**.

by (Pennec et al., 2025), we found improvements to be less pronounced on traditional metrics such as ROUGE or METEOR as those metrics are not primarily designed for multimodality. All the above results were found to be similar when using other summarization models instead (see Appendix B).

Choice of K on the Summarization Performance Although the clip selection improves with larger values of K (see Recall@ K in Table 1), this observation does not apply to the quality of the end summary. Indeed, the summarization performance reported in Table 3 seems to saturate rather than monotonically increase with K .

Since our summaries are constrained to a fixed target length (1000 words), we believe that growing values of K does not necessarily yield better summaries, as additional clips often exceed what the LLM can effectively leverage given the summary length constraint.

Importance of the captioning quality The quality of the captions used to build the screenplay has a critical role. Indeed, summaries generated using Gemini 2.5 Flash-Lite for recaptioning capture significantly less visual information (visual recall) than those generated from the gold screenplay as input (Table 3).

7 Discussion

Ablation Study: Effect of Recaptioning on Summarization Quality Table 4 examines the impact of the recaptioning step on the overall summarization performance of our pipeline (Figure 1). In particular, we observe a consistent decrease in the visual recall when no recaptioning is being performed, indicating the importance of this step for capturing important visual information.

This finding reveals a clear division of labor between the two captioning stages in our pipeline. While lightweight captions are sufficient for identifying salient clips, they often miss finer visual details that are crucial for building accurate multimodal summaries. This design balances efficiency and accuracy: most of the video is processed cheaply, while the few clips that matter are described in depth to boost summary quality.

Comparing Video Segmentation Approaches for Clip Selection A natural alternative to our fixed 20-second clips in Section 3.1 is scene segmentation, which divides the video into shorter scenes that better align with semantic shifts in the narrative. We infer these scenes in a zero-shot manner using the method from (Mahon and Lapata, 2025).

On our test set, the average scene duration is 73 seconds. In order to match our original setup, we

	vis-rec	text-rec	MFS	r1	r2	rlsum	METEOR
Built Screenplay (25 clips)							
w/o recaptioning	19.31	36.71	28.01	47.07	11.64	44.80	33.79
with recaptioning (ours)	21.05	34.01	27.53	46.90	11.55	44.69	33.93
Built Screenplay (50 clips)							
w/o recaptioning	19.96	34.02	26.99	46.14	11.05	43.99	33.30
with recaptioning (ours)	20.60	34.22	27.41	46.22	10.97	44.10	33.46
Built Screenplay (75 clips)							
w/o recaptioning	21.15	34.73	27.94	46.79	11.58	44.71	33.66
with recaptioning (ours)	22.25	33.45	27.85	46.53	10.80	44.32	33.55

Table 4: **Effect of recaptioning on the summarization pipeline performance.** Recaptioning of visually significant moments with a stronger model (Gemini 2.5 Flash Lite) directly improves how well the generated summary captures important visual information (visual recall). In the above, we always perform clip selection using Qwen2.5-Omni-7B as the lightweight captioning model and summarization using Gemini 2.5 Flash. Column descriptions are the same as in Table 2. Best results are in **bold**.

	R@25	R@50	R@75
our clips zero-shot			
w/o scene segmentation	11.89	19.89	26.60
with scene segmentation	10.94	18.07	23.64
our clips two-shot			
w/o scene segmentation	13.82	21.95	26.25
with scene segmentation	10.36	18.40	24.98

Table 5: **Effect of scene segmentation on our clip selection.** Scene segmentation does not positively impact the performance of our clip selection. In the above, we use Qwen2.5-Omni-7B as the lightweight captioning model. Column descriptions are the same as in Table 1. Best results are in **bold**.

further subdivide each scene into shorter segments by uniformly splitting them so that the average segment is now of 20 seconds. This is to ensure that the two approaches are comparable while we still benefit from the scene boundaries given by the scene segmentation.

Despite being more natural, scene-based segmentation did not outperform our fixed 20-second clips (Table 5). Since clip selection is done at the caption level, we believe that performance depends less on whether segment boundaries match with the scenes or are chosen arbitrarily.

Subjectivity of the Clip Selection and Summarization Tasks To evaluate the subjectivity of the clip selection task across different summary sources, we also collect summaries from The Movie Spoiler³. From the 200 movies in our test set, we successfully retrieve 54 corresponding summaries. Following the same procedure described in A.4, we infer the clips for those summaries. The

³<https://themoviespoiler.com/>

Movie Spoiler summaries are longer and we infer twice as many clips from those on average. We compute the overlap between the clips found in the two summary sources. The Movie Spoiler summaries recover about 48.8% of the reference clips in MovieSum summaries. Also, when evaluated against the MovieSum reference, they achieved a visual recall of 66.2% and textual recall of 61.1%.

8 Conclusion

This paper tackles the dual challenges of long video summarization: the high computational cost and the risk of missing crucial visual information. We propose a cost-effective, clip selection for the task. Our method performs initial captioning of short video segments at a lower cost followed by selection of key visual moments by an LLM for inclusion into the multimodal summary. Our experiments on the MovieSum dataset demonstrated that a small fraction of the movies, about 6% of their content, is sufficient to generate a comprehensive multimodal summary, validating the core principle of our approach. Second, our proposed clip selection method significantly outperforms the tested baselines, capturing substantially more relevant visual information than random clip selection. Crucially, the summaries built from our selected clips achieve a performance close to those generated from a perfect set of reference clips, demonstrating the robustness of our selection strategy.

Future work could extend this methodology to other multimodal generative tasks and domains, and explore different selection criteria. Overall, our findings suggest that focusing on minimal yet highly salient clips offers an efficient paradigm for understanding long-form video content.

Limitations

The performance of clip selection is closely tied to the quality of the lightweight captioning (Section 5.2), suggesting that improvements in smaller VLMs could yield further gains.

Adaptive clip selection strategies that dynamically choose K based on the video duration and density would be useful to explore. In the meantime, our experiments reveal the limited impact of varying K on the end summary and this is mainly due to the fixed length of the generated summary. Such adaptive strategies for varying K could be particularly beneficial in an unconstrained summarization setting, where the summary length is not fixed and this could be investigated in future work.

While our method outperforms the random clip selection baseline, it still incurs a computational cost, both in generating captions and choosing the K best clips. This cost is still lower than processing videos end-to-end using a high-performing VLM such as Gemini.

Acknowledgments

We thank the anonymous reviewers for their feedback. This research is supported by the National Research Foundation, Prime Minister’s Office, Singapore under its Campus for Research Excellence and Technological Enterprise (CREATE) programme. Any opinions, findings and conclusions or recommendations expressed in this material are those of the authors and do not reflect the views of the National Research Foundation, Singapore.

References

- Yash Kumar Atri, Shraman Pramanick, Vikram Goyal, and Tanmoy Chakraborty. 2021. [See, hear, read: Leveraging multimodality with guided attention for abstractive text summarization](#). *Knowl. Based Syst.*, 227:107152.
- Shuai Bai, Keqin Chen, Xuejing Liu, Jialin Wang, Wenbin Ge, Sibao Song, Kai Dang, Peng Wang, Shijie Wang, Jun Tang, Humen Zhong, Yuanzhi Zhu, Ming-Hsuan Yang, Zhaohai Li, Jianqiang Wan, Pengfei Wang, Wei Ding, Zheren Fu, Yiheng Xu, Jiabo Ye, Xi Zhang, Tianbao Xie, Zesen Cheng, Hang Zhang, Zhibo Yang, Haiyang Xu, and Junyang Lin. 2025. [Qwen2.5-VL technical report](#). *CoRR*, abs/2502.13923.
- Ivana Balazevic, Yuge Shi, Pinelopi Papalampidi, Rahma Chaabouni, Skanda Koppula, and Olivier J. Hénaff. 2024. [Memory consolidation enables long-context video understanding](#). In *Forty-first International Conference on Machine Learning, ICML 2024, Vienna, Austria, July 21-27, 2024*. OpenReview.net.
- Brian Y. Chen, Xiangyuan Zhao, and Yingnan Zhu. 2024. [Personalized video summarization by multimodal video understanding](#). In *Proceedings of the 33rd ACM International Conference on Information and Knowledge Management, CIKM 2024, Boise, ID, USA, October 21-25, 2024*, pages 4382–4389. ACM.
- Jiaao Chen and Diyi Yang. 2020. [Multi-view sequence-to-sequence models with conversational structure for abstractive dialogue summarization](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 4106–4118, Online. Association for Computational Linguistics.
- Gheorghe Comanici, Eric Bieber, Mike Schaeckermann, Ice Pasupat, Noveen Sachdeva, Inderjit S. Dhillon, Marcel Blistein, Ori Ram, Dan Zhang, Evan Rosen, Luke Marris, Sam Petulla, Colin Gaffney, Asaf Aharoni, Nathan Lintz, Tiago Cardal Pais, Henrik Jacobsson, Idan Szpektor, Nan-Jiang Jiang, Krishna Haridasan, Ahmed Omran, Nikunj Saunshi, Dara Bahri, Gaurav Mishra, Eric Chu, Toby Boyd, Brad Hekman, Aaron Parisi, Chaoyi Zhang, Kornrathop Kawintiranon, Tania Bedrax-Weiss, Oliver Wang, Ya Xu, Ollie Purkiss, Uri Mendlovic, Ilia Deutel, Nam Nguyen, Adam Langley, Flip Korn, Lucia Rossazza, Alexandre Ramé, Sagar Waghmare, Helen Miller, Nathan Byrd, Ashrith Sheshan, Raia Hadsell Sangnie Bhardwaj, Pawel Janus, Tero Rissa, Dan Horgan, Sharon Silver, Ayzaan Wahid, Sergey Brin, Yves Raimond, Klemen Kloboves, Cindy Wang, Nitesh Bharadwaj Gundavarapu, Ilia Shumailov, Bo Wang, Mantas Pajarskas, Joe Heyward, Martin Nikoltchev, Maciej Kula, Hao Zhou, Zachary Garrett, Sushant Kaffle, Sercan Arik, Ankita Goel, Mingyao Yang, Jiho Park, Koji Kojima, Parsa Mahmoudieh, Koray Kavukcuoglu, Grace Chen, Doug Fritz, Anton Bulyenov, Sudeshna Roy, Dimitris Paparas, Hadar Shemtov, Bo-Juen Chen, Robin Strudel, David Reitter, Aurko Roy, Andrey Vlasov, Changwan Ryu, Chas Lechner, Haichuan Yang, Zeldia Mariet, Denis Vnukov, Tim Sohn, Amy Stuart, Wei Liang, Minmin Chen, Praynaa Rawlani, Christy Koh, JD Co-Reyes, Guangda Lai, Praseem Banzal, Dimitrios Vytiniotis, Jieru Mei, and Mu Cai. 2025. [Gemini 2.5: Pushing the frontier with advanced reasoning, multimodality, long context, and next generation agentic capabilities](#). *CoRR*, abs/2507.06261.
- Xiachong Feng, Xiaocheng Feng, and Bing Qin. 2021. [A survey on dialogue summarization: Recent advances and new frontiers](#). *arXiv preprint arXiv:2107.03175*.
- Chaoyou Fu, Yuhan Dai, Yondong Luo, Lei Li, Shuhuai Ren, Renrui Zhang, Zihan Wang, Chenyu Zhou, Yunhang Shen, Mengdan Zhang, Peixian Chen, Yanwei Li, Shaohui Lin, Sirui Zhao, Ke Li, Tong Xu, Xiawu Zheng, Enhong Chen, Rongrong Ji, and Xing Sun.

2024. [Video-MME: The first-ever comprehensive evaluation benchmark of multi-modal LLMs in video analysis](#). *CoRR*, abs/2405.21075.
- Sebastian Gehrmann, Yuntian Deng, and Alexander Rush. 2018. [Bottom-up abstractive summarization](#). In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 4098–4109, Brussels, Belgium. Association for Computational Linguistics.
- Bo He, Hengduo Li, Young Kyun Jang, Menglin Jia, Xuefei Cao, Ashish Shah, Abhinav Shrivastava, and Ser-Nam Lim. 2024. [MA-LMM: memory-augmented large multimodal model for long-term video understanding](#). In *IEEE/CVF Conference on Computer Vision and Pattern Recognition, CVPR 2024, Seattle, WA, USA, June 16-22, 2024*, pages 13504–13514. IEEE.
- De-An Huang, Subhashree Radhakrishnan, Zhiding Yu, and Jan Kautz. 2025. [FRAG: frame selection augmented generation for long video and long document understanding](#). *CoRR*, abs/2504.17447.
- Kumara Kahatapitiya, Kanchana Ranasinghe, Jongwoo Park, and Michael S. Ryoo. 2024. [Language repository for long video understanding](#). *CoRR*, abs/2403.14622.
- Chris Kedzie, Kathleen McKeown, and Hal Daumé III. 2018. [Content selection in deep learning models of summarization](#). In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 1818–1828, Brussels, Belgium. Association for Computational Linguistics.
- Bruno Korbar, Yongqin Xian, Alessio Tonioni, Andrew Zisserman, and Federico Tombari. 2024. [Text-conditioned resampler for long form video understanding](#). In *Computer Vision - ECCV 2024 - 18th European Conference, Milan, Italy, September 29-October 4, 2024, Proceedings, Part LXXXVI*, volume 15144 of *Lecture Notes in Computer Science*, pages 271–288. Springer.
- Ranjay Krishna, Kenji Hata, Frederic Ren, Li Fei-Fei, and Juan Carlos Niebles. 2017. [Dense-captioning events in videos](#). In *IEEE International Conference on Computer Vision, ICCV 2017, Venice, Italy, October 22-29, 2017*, pages 706–715. IEEE Computer Society.
- Wojciech Kryściński, Bryan McCann, Caiming Xiong, and Richard Socher. 2020. [Evaluating the factual consistency of abstractive text summarization](#). In *Proceedings of the 2020 conference on empirical methods in natural language processing (EMNLP)*, pages 9332–9346.
- Wojciech Kryscinski, Nazneen Rajani, Divyansh Agarwal, Caiming Xiong, and Dragomir Radev. 2022. [BOOKSUM: A collection of datasets for long-form narrative summarization](#). In *Findings of the Association for Computational Linguistics: EMNLP 2022*, pages 6536–6558, Abu Dhabi, United Arab Emirates. Association for Computational Linguistics.
- Jie Lei, Licheng Yu, Tamara L. Berg, and Mohit Bansal. 2020. [TVQA+: spatio-temporal grounding for video question answering](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics, ACL 2020, Online, July 5-10, 2020*, pages 8211–8225. Association for Computational Linguistics.
- Haoran Li, Junnan Zhu, Cong Ma, Jiajun Zhang, and Chengqing Zong. 2017. [Multi-modal summarization for asynchronous collection of text, image, audio and video](#). In *Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing*, pages 1092–1102, Copenhagen, Denmark. Association for Computational Linguistics.
- Yanwei Li, Chengyao Wang, and Jiaya Jia. 2024. [LLaMA-VID: An image is worth 2 tokens in large language models](#). In *Computer Vision - ECCV 2024 - 18th European Conference, Milan, Italy, September 29-October 4, 2024, Proceedings, Part XLVI*, volume 15104 of *Lecture Notes in Computer Science*, pages 323–340. Springer.
- Dongqi Liu, Chenxi Whitehouse, Xi Yu, Louis Mahon, Rohit Saxena, Zheng Zhao, Yifu Qiu, Mirella Lapata, and Vera Demberg. 2025. [What is that talk about? A video-to-text summarization dataset for scientific presentations](#). In *Proceedings of the 63rd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers), ACL 2025, Vienna, Austria, July 27 - August 1, 2025*, pages 6187–6210. Association for Computational Linguistics.
- Zhengyuan Liu and Nancy Chen. 2021. [Controllable neural dialogue summarization with personal named entity planning](#). In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*, pages 92–106, Online and Punta Cana, Dominican Republic. Association for Computational Linguistics.
- Zhengyuan Liu and Nancy Chen. 2022. [Dynamic sliding window modeling for abstractive meeting summarization](#). In *Proc. Interspeech 2022*, pages 5150–5154.
- Zhengyuan Liu, Angela Ng, Sheldon Lee, Ai Ti Aw, and Nancy F Chen. 2019. [Topic-aware pointer-generator networks for summarizing spoken conversations](#). In *2019 IEEE Automatic Speech Recognition and Understanding Workshop (ASRU)*, pages 814–821. IEEE.
- Zhengyuan Liu, Ke Shi, and Nancy Chen. 2021. [Coreference-aware dialogue summarization](#). In *Proceedings of the 22nd Annual Meeting of the Special Interest Group on Discourse and Dialogue*, pages 509–519, Singapore and Online. Association for Computational Linguistics.
- Louis Mahon and Mirella Lapata. 2024a. [A modular approach for multimodal summarization of TV shows](#).

- In *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, ACL 2024, Bangkok, Thailand, August 11-16, 2024, pages 8272–8291. Association for Computational Linguistics.
- Louis Mahon and Mirella Lapata. 2024b. [Screenwriter: Automatic screenplay generation and movie summarisation](#). *CoRR*, abs/2410.19809.
- Louis Mahon and Mirella Lapata. 2025. [Parameter-free video segmentation for vision and language understanding](#). *CoRR*, abs/2503.01201.
- Karttikeya Mangalam, Raiymbek Akshulakov, and Jitendra Malik. 2023. [Egoschema: A diagnostic benchmark for very long-form video language understanding](#). In *Advances in Neural Information Processing Systems 36: Annual Conference on Neural Information Processing Systems 2023, NeurIPS 2023, New Orleans, LA, USA, December 10 - 16, 2023*.
- Antoine Miech, Dimitri Zhukov, Jean-Baptiste Alayrac, Makarand Tapaswi, Ivan Laptev, and Josef Sivic. 2019. [HowTo100M: Learning a text-video embedding by watching hundred million narrated video clips](#). In *2019 IEEE/CVF International Conference on Computer Vision, ICCV 2019, Seoul, Korea (South), October 27 - November 2, 2019*, pages 2630–2640. IEEE.
- Feng Nan, Ramesh Nallapati, Zhiguo Wang, Cicero Nogueira dos Santos, Henghui Zhu, Dejiao Zhang, Kathleen McKeown, and Bing Xiang. 2021. [Entity-level factual consistency of abstractive text summarization](#). In *Proceedings of the 16th Conference of the European Chapter of the Association for Computational Linguistics: Main Volume*, pages 2727–2733, Online. Association for Computational Linguistics.
- Medhini Narasimhan, Anna Rohrbach, and Trevor Darrell. 2021. [CLIP-It! language-guided video summarization](#). In *Advances in Neural Information Processing Systems 34: Annual Conference on Neural Information Processing Systems 2021, NeurIPS 2021, December 6-14, 2021, virtual*, pages 13988–14000.
- Taichi Nishimura, Shota Nakada, and Masayoshi Kondo. 2024. [On the audio hallucinations in large audio-video language models](#). *CoRR*, abs/2401.09774.
- OpenAI. 2024. [Hello GPT-4o](#). Accessed: 2024-11-6.
- Pinelopi Papalampidi, Frank Keller, and Mirella Lapata. 2021. [Movie summarization via sparse graph construction](#). In *Thirty-Fifth AAAI Conference on Artificial Intelligence, AAAI 2021, Thirty-Third Conference on Innovative Applications of Artificial Intelligence, IAAI 2021, The Eleventh Symposium on Educational Advances in Artificial Intelligence, EAAI 2021, Virtual Event, February 2-9, 2021*, pages 13631–13639. AAAI Press.
- Pinelopi Papalampidi and Mirella Lapata. 2023. [Hierarchical3D adapters for long video-to-text summarization](#). In *Findings of the Association for Computational Linguistics: EACL 2023, Dubrovnik, Croatia, May 2-6, 2023*, pages 1267–1290. Association for Computational Linguistics.
- Jean Park, Kuk Jin Jang, Basam Alasaly, Sriharsha Mopidevi, Andrew Zolensky, Eric Eaton, Insup Lee, and Kevin Johnson. 2024a. [Assessing modality bias in video question answering benchmarks with multimodal large language models](#). *CoRR*, abs/2408.12763.
- Jongwoo Park, Kanchana Ranasinghe, Kumara Kahatapitiya, Wonjeong Ryoo, Donghyun Kim, and Michael S. Ryoo. 2024b. [Too many frames, not all useful: Efficient strategies for long-form video QA](#). *CoRR*, abs/2406.09396.
- Galann Pennec, Zhengyuan Liu, Nicholas Asher, Philippe Muller, and Nancy F. Chen. 2025. [Integrating video and text: A balanced approach to multimodal summary generation and evaluation](#). *CoRR*, abs/2505.06594.
- Rui Qian, Xiaoyi Dong, Pan Zhang, Yuhang Zang, Shuangrui Ding, Dahua Lin, and Jiaqi Wang. 2024. [Streaming long video understanding with large language models](#). In *Advances in Neural Information Processing Systems 38: Annual Conference on Neural Information Processing Systems 2024, NeurIPS 2024, Vancouver, BC, Canada, December 10 - 15, 2024*.
- Machel Reid, Nikolay Savinov, Denis Teplyashin, Dmitry Lepikhin, Timothy P. Lillicrap, Jean-Baptiste Alayrac, Radu Soricut, Angeliki Lazaridou, Orhan Firat, Julian Schrittwieser, Ioannis Antonoglou, Rohan Anil, Sebastian Borgeaud, Andrew M. Dai, Katie Millican, Ethan Dyer, Mia Glaese, Thibault Sottiaux, Benjamin Lee, Fabio Viola, Malcolm Reynolds, Yuanzhong Xu, James Molloy, Jilin Chen, Michael Isard, Paul Barham, Tom Hennigan, Ross McIlroy, Melvin Johnson, Johan Schalkwyk, Eli Collins, Eliza Rutherford, Erica Moreira, Kareem Ayoub, Megha Goel, Clemens Meyer, Gregory Thornton, Zhen Yang, Henryk Michalewski, Zaheer Abbas, Nathan Schucher, Ankesh Anand, Richard Ives, James Keeling, Karel Lenc, Salem Haykal, Siamak Shakeri, Pranav Shyam, Aakanksha Chowdhery, Roman Ring, Stephen Spencer, Eren Sezener, and et al. 2024. [Gemini 1.5: Unlocking multimodal understanding across millions of tokens of context](#). *CoRR*, abs/2403.05530.
- Rohit Saxena and Frank Keller. 2024. [MovieSum: An abstractive summarization dataset for movie screenplays](#). In *Findings of the Association for Computational Linguistics, ACL 2024, Bangkok, Thailand and virtual meeting, August 11-16, 2024*, pages 4043–4050. Association for Computational Linguistics.
- Abigail See, Peter J Liu, and Christopher D Manning. 2017. [Get to the point: Summarization with pointer-generator networks](#). In *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 1073–1083.

- Yuhan Shen, Linjie Yang, Longyin Wen, Haichao Yu, Ehsan Elhamifar, and Heng Wang. 2024. [Exploring the role of audio in video captioning](#). In *IEEE/CVF Conference on Computer Vision and Pattern Recognition, CVPR 2024 - Workshops, Seattle, WA, USA, June 17-18, 2024*, pages 2090–2100. IEEE.
- Aditya Kumar Singh, Dhruv Srivastava, and Makarand Tapaswi. 2024. ["previously on..." from recaps to story summarization](#). In *IEEE/CVF Conference on Computer Vision and Pattern Recognition, CVPR 2024, Seattle, WA, USA, June 16-22, 2024*, pages 13635–13646. IEEE.
- Enxin Song, Wenhao Chai, Guan hong Wang, Yucheng Zhang, Haoyang Zhou, Feiyang Wu, Haozhe Chi, Xun Guo, Tian Ye, Yanting Zhang, Yan Lu, Jenq-Neng Hwang, and Gaoang Wang. 2024. [MovieChat: From dense token to sparse memory for long video understanding](#). In *IEEE/CVF Conference on Computer Vision and Pattern Recognition, CVPR 2024, Seattle, WA, USA, June 16-22, 2024*, pages 18221–18232. IEEE.
- Yuting Tang, Ratish Puduppully, Zhengyuan Liu, and Nancy Chen. 2023. [In-context learning of large language models for controlled dialogue summarization: A holistic benchmark and empirical analysis](#). In *Proceedings of the 4th New Frontiers in Summarization Workshop*, pages 56–67, Singapore. Association for Computational Linguistics.
- Bin Wang, Zhengyuan Liu, and Nancy Chen. 2023. [Instructive dialogue summarization with query aggregations](#). In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*, pages 7630–7653, Singapore. Association for Computational Linguistics.
- Weihan Wang, Zehai He, Wenyi Hong, Yean Cheng, Xiaohan Zhang, Ji Qi, Shiyu Huang, Bin Xu, Yuxiao Dong, Ming Ding, and Jie Tang. 2024a. [LVBench: An extreme long video understanding benchmark](#). *CoRR*, abs/2406.08035.
- Xiao Wang, Qingyi Si, Jianlong Wu, Shiyu Zhu, Li Cao, and Liqiang Nie. 2025. [AdaReTaKe: Adaptive redundancy reduction to perceive longer for video-language understanding](#). *CoRR*, abs/2503.12559.
- Xiaohan Wang, Yuhui Zhang, Orr Zohar, and Serena Yeung-Levy. 2024b. [VideoAgent: Long-form video understanding with large language model as agent](#). In *Computer Vision - ECCV 2024 - 18th European Conference, Milan, Italy, September 29-October 4, 2024, Proceedings, Part LXXX*, volume 15138 of *Lecture Notes in Computer Science*, pages 58–76. Springer.
- Ziyang Wang, Shoubin Yu, Elias Stengel-Eskin, Jaehong Yoon, Feng Cheng, Gedas Bertasius, and Mohit Bansal. 2024c. [VideoTree: Adaptive tree-based video representation for LLM reasoning on long videos](#). *CoRR*, abs/2405.19209.
- Xilin Wei, Xiaoran Liu, Yuhang Zang, Xiaoyi Dong, Pan Zhang, Yuhang Cao, Jian Tong, Haodong Duan, Qipeng Guo, Jiaqi Wang, Xipeng Qiu, and Dahua Lin. 2025. [VideoRoPE: What makes for good video rotary position embedding?](#) *CoRR*, abs/2502.05173.
- Yuetian Weng, Mingfei Han, Haoyu He, Xiaojun Chang, and Bohan Zhuang. 2024. [LongVLM: Efficient long video understanding via large language models](#). In *Computer Vision - ECCV 2024 - 18th European Conference, Milan, Italy, September 29-October 4, 2024, Proceedings, Part XXXIII*, volume 15091 of *Lecture Notes in Computer Science*, pages 453–470. Springer.
- Jin Xu, Zhifang Guo, Jinzheng He, Hangrui Hu, Ting He, Shuai Bai, Keqin Chen, Jialin Wang, Yang Fan, Kai Dang, Bin Zhang, Xiong Wang, Yunfei Chu, and Junyang Lin. 2025. [Qwen2.5-omni technical report](#). *CoRR*, abs/2503.20215.
- Fuzhao Xue, Yukang Chen, Dacheng Li, Qinghao Hu, Ligeng Zhu, Xiuyu Li, Yunhao Fang, Haotian Tang, Shang Yang, Zhijian Liu, Ethan He, Hongxu Yin, Pavlo Molchanov, Jan Kautz, Linxi Fan, Yuke Zhu, Yao Lu, and Song Han. 2024. [LongVILA: Scaling long-context visual language models for long videos](#). *CoRR*, abs/2408.10188.
- An Yang, Baosong Yang, Beichen Zhang, Binyuan Hui, Bo Zheng, Bowen Yu, Chengyuan Li, Dayiheng Liu, Fei Huang, Haoran Wei, Huan Lin, Jian Yang, Jianhong Tu, Jianwei Zhang, Jianxin Yang, Jiayi Yang, Jingren Zhou, Junyang Lin, Kai Dang, Keming Lu, Keqin Bao, Kexin Yang, Le Yu, Mei Li, Mingfeng Xue, Pei Zhang, Qin Zhu, Rui Men, Runji Lin, Tianhao Li, Tingyu Xia, Xingzhang Ren, Xuancheng Ren, Yang Fan, Yang Su, Yichang Zhang, Yu Wan, Yuqiong Liu, Zeyu Cui, Zhenru Zhang, and Zihan Qiu. 2024a. [Qwen2.5 technical report](#). *CoRR*, abs/2412.15115.
- Zeyuan Yang, Delin Chen, Xueyang Yu, Maohao Shen, and Chuang Gan. 2024b. [VCA: video curious agent for long video understanding](#). *CoRR*, abs/2412.10471.
- Michihiro Yasunaga, Jungo Kasai, Rui Zhang, Alexander R Fabbri, Irene Li, Dan Friedman, and Dragomir R Radev. 2019. [Scisummnet: A large annotated corpus and content-impact models for scientific paper summarization with citation networks](#). In *Proceedings of the AAAI conference on artificial intelligence*, volume 33, pages 7386–7393.
- Shoubin Yu, Jaemin Cho, Prateek Yadav, and Mohit Bansal. 2023. [Self-chained image-language model for video localization and question answering](#). In *Advances in Neural Information Processing Systems 36: Annual Conference on Neural Information Processing Systems 2023, NeurIPS 2023, New Orleans, LA, USA, December 10 - 16, 2023*.
- Sicheng Yu, Chengkai Jin, Huanyu Wang, Zhenghao Chen, Sheng Jin, Zhongrong Zuo, Xiaolei Xu, Zhenbang Sun, Bingni Zhang, Jiawei Wu, Hao Zhang,

and Qianru Sun. 2025. [Frame-Voyager: Learning to query frames for video large language models](#). In *The Thirteenth International Conference on Learning Representations, ICLR 2025, Singapore, April 24-28, 2025*. OpenReview.net.

Peiyuan Zhang, Kaichen Zhang, Bo Li, Guangtao Zeng, Jingkang Yang, Yuanhan Zhang, Ziyue Wang, Hao-ran Tan, Chunyuan Li, and Ziwei Liu. 2024a. [Long context transfer from language to vision](#). *CoRR*, abs/2406.16852.

Yifan Zhang, Weichen Yu, Qingsong Wen, Xue Wang, Zhang Zhang, Liang Wang, Rong Jin, and Tieniu Tan. 2024b. [Debiasing multimodal large language models](#). *CoRR*, abs/2403.05262.

Zhuo Zhi, Qiangqiang Wu, Minghe shen, Wenbo Li, Yinchuan Li, Kun Shao, and Kaiwen Zhou. 2025. [VideoAgent2: Enhancing the LLM-based agent system for long-form video understanding by uncertainty-aware CoT](#). *CoRR*, abs/2504.04471.

Junjie Zhou, Yan Shu, Bo Zhao, Boya Wu, Shitao Xiao, Xi Yang, Yongping Xiong, Bo Zhang, Tiejun Huang, and Zheng Liu. 2024. [MLVU: A comprehensive benchmark for multi-task long video understanding](#). *CoRR*, abs/2406.04264.

A Prompts

A.1 Clip Selection

A.1.1 Clip Selection Prompt

We provide below the prompt for the clip selection with an LLM.

- [MOVIE_NAME](#) is the movie title.
- [<CAPTIONS>](#) refers to all the captions generated for the 20-second video clips using the lightweight captioning model (Qwen2.5-Omni).
- [NB_CAPTIONS](#) is the number of selected clips (same as K).

Here are captions from the movie [MOVIE_NAME](#):

[<CAPTIONS>](#)

What are the [NB_CAPTIONS](#) most important Captions that describe important action or visual event you would include in the existing Summary of the movie [MOVIE_NAME](#)?

Provide your answer in the following way:

1. Caption [caption_number](#): Justification why the Caption describes crucial action for the

summary

2. Caption [caption_number](#): Justification why the Caption describes crucial action for the summary

...

[NB_CAPTIONS](#). Caption [caption_number](#): Justification why the Caption describes crucial action for the summary

Answer:

A.1.2 Two-shot Clip Selection Examples

We annotate and use the following few-shot examples for the clip selection task. Those examples are derived from the movies *Forrest Gump* (1994) and *Wonder Woman* (2017).

Here are captions from the movie Forrest Gump:

Caption 1110000: In the video, a man and woman sit on a bench in a park. The man is wearing a suit and tie while the woman wears casual clothes. They appear to be reading books together as they sit side by side. The man then turns his attention towards the woman and starts talking about something. He mentions that life is like a box of chocolates and you never know what you're going to get. He also comments on how comfortable her shoes must be and suggests she could walk all day in them.

Caption 1130000: Forrest is sitting on a bench outside. He then sits inside a doctor's office with his legs up on the table. The doctor removes Forrest's leg braces and asks him to stand up. Forrest stands up and walks around the room.

Caption 1150000: The dialogue reveals that the woman is explaining the origin of the character's name "Forrest Gump." She mentions that the "Forrest" part of the name comes from an incident where they were related to someone who started a club called the Ku Klux Klan. The woman explains that the "Gump" part of the name was given because sometimes people do things that don't make sense.

Caption 1170000: The video shows a group of boys chasing Forrest Gump as he runs down a dirt road. The boys are shouting at him to run faster, while Forrest continues to run without looking back. One of the boys falls over, but gets up quickly and continues chasing Forrest. The other boys also catch up with Forrest and start to chase him more aggressively. As they get closer, one of the boys throws a rock at Forrest, who ducks to avoid it. Another boy tries to kick him, but misses. The boys continue to chase Forrest until he reaches his home, where his mother is waiting for him. She tells him that miracles happen every day, and that some people may not believe them, but they still exist.

Caption 1190000: The man is running on the field, and he jumps over the fence. He runs to the football field and throws the ball. The coaches are watching him.

Caption 1210000: The video shows a scene where a woman holding a baby sits on a bench next to another woman who is reading a book. A man in a suit is sitting on the other side of the bench with his suitcase beside him. The woman with the baby stands up and walks away from the bench while talking to the man. She then sits back down on the bench and continues talking to him. In the background, there is a bus passing by. The dialogue includes the woman asking if the bus is the number nine, but the man corrects her and says it's the number four. They also have a conversation about someone named Wallace getting shot while they were in college.

Caption 1230000: The video shows a woman reading a book to her son on their bed. The boy asks his mother about vacation, and she explains that it is when someone goes somewhere and never comes back.

What are the 3 most important Captions that describe important action or visual event you would include in a Summary of the movie Forrest Gump?

Provide your answer in the following way:

1. Caption caption_number: Justification why the Caption describes crucial action for the summary
2. Caption caption_number: Justification why

the Caption describes crucial action for the summary

3. Caption caption_number: Justification why the Caption describes crucial action for the summary

Answer:

Caption 1130000: Justification: This caption depicts the removal of Forrest's leg braces, a pivotal moment signifying his physical transformation and newfound freedom.

Caption 1170000: Justification: This caption illustrates the bullying Forrest faces and his eventual discovery of his running ability, a recurring motif in the film.

Caption 1190000: Justification: This caption depicts Forrest's accidental entry into the world of football, showcasing his unexpected athletic talent.

Here are captions from the movie Wonder Woman:

Caption 4210000: The scene opens with a man sitting at his desk, looking at his watch. He then turns to face another man standing before him. The man in uniform speaks to the other man, telling him that he will do nothing. The man in uniform then walks away as the other man looks on. The scene ends with the man in uniform walking out of the room.

Caption 4230000: Diana and Steve are walking down the stairs. Steve is talking to Diana. Steve is angry at Diana for not fighting back against Ares. He tells her that she didn't stand her ground because there was no chance of changing Ares' mind. He also tells her that millions of people will die if they don't fight back. He tells her that his people are next. Summary: Steve is angry at Diana for not fighting back against Ares. He tells her that she didn't stand her ground because there was no chance of changing Ares' mind. He also tells her that millions of people will die if they don't fight back. He tells her that his people are next.

Caption 4250000: The video shows a man sitting on a chair in a room. A bomb is thrown

into the room and explodes. The man gets up and runs out of the door. He then talks to another man who is standing outside the door. The man inside the room is coughing and choking on smoke.

What are the 1 most important Captions that describe important action or visual event you would include in a Summary of the movie Wonder Woman?

Provide your answer in the following way:

1. Caption caption_number: Justification why the Caption describes crucial action for the summary

Answer:

Caption 4250000: Justification: This caption depicts a sudden and violent attack, showcasing the dangers faced by the characters and the chaos of the war. It emphasizes the element of surprise and the characters' ability to react quickly to threats. Therefore the Caption depicts important visual action of event.

A.2 Clip captioning Prompts

Below are the prompt templates used for the lightweight captioning with Qwen2.5-Omni and the recaptioning with Gemini 2.5 Flash-Lite. The video clips are processed by both VLMs at one frame per second (1 fps) and including the audio.

A.2.1 Lightweight Captioning with Qwen2.5-Omni

<VIDEO CLIP (1 fps)+ AUDIO>

Describe both the action and Summarize the corresponding dialogue.

A.2.2 Recaptioning with Gemini 2.5 Flash-Lite

<VIDEO CLIP (1 fps)+ AUDIO>

Describe both the video, action and dialogue in one paragraph

A.3 Summarization Prompt

We provide here the prompt we used for generating multimodal summaries in all our experiments.

We explicitly state in our prompt that the produced summary has to be multimodal by including both relevant visual and textual elements from either the transcript lines and the video captions.

We fix the generated summary length to 1000 words in the prompt and truncate the output beyond that limit. Note that the average summary length of the groundtruth summaries in the whole MovieSum dataset (train and test sets) is 635 words.

<TRANSCRIPT> or <SCREENPLAY>

Generate a comprehensive multimodal summary of exactly 1000 words of the movie based on the provided dialogue and the most important visual elements.

Your summary should:

Synthesize information from both the dialogue (transcript) and the important visual events (visual analysis).

Your overall summary should contain exactly 1000 words. Do not refer to external websites, movie databases or plot summaries.

A.4 Clip Selection Reference

A.4.1 Fact Identification

We provide below the prompt for extracting all the facts from the groundtruth summary by first splitting the summary into sentences and then each sentence into facts.

Summary:

<SUMMARY>

For every sentence from the Summary, decompose the sentence in a list of facts (at least 1). Each fact can be only part of a sentence and should convey a single piece of information about the story.

Example:

Sentence 1:

*

Sentence 2:

*

...

Sentence N:

*

A.4.2 Visual Fact Classification

Given the gold screenplay of a movie, we are able to infer which groundtruth summary fact is Visual or Textual.

We prompt an LLM in zero-shot to quote the line from the screenplay that supports a given groundtruth summary fact. If the quoted line belongs to the dialogue, then the fact is classified as Textual. Otherwise, if it corresponds to a clip caption, then the fact is classified as Visual. We provide below the prompt being used for the task of visual fact classification.

Fact Evaluation against Transcripts

Screenplay:

<SCREENPLAY>

For every fact below:

<ALL FACTS>

-> Find the information in the above Screenplay. Quote a line from the Screenplay.

Example:

Fact 1: Recopy the Fact

-> Quoted line from Screenplay

Fact 2: Recopy the Fact

-> Quoted line from Screenplay

...

Fact N: Recopy the Fact

-> Quoted line from Screenplay

Summary:

<SUMMARY>

Task:

For each fact listed below, determine whether the exact meaning of the fact is explicitly present in the summary above.

Instructions:

You must justify your answer by quoting or paraphrasing the relevant part of the summary. If the fact is not explicitly present, even if it seems implied or suggested, you must answer No.

Do not accept facts just because they are likely, inferable, or assumed from context. However, do allow for reasonable paraphrasing or rewording. If the summary conveys the same meaning as the fact using different but equivalent words, answer Yes.

Format:

Fact 1: [Recopy the Fact]

1. Justification (quote or paraphrase from the summary, and explain how it matches the fact)

2. Yes

Fact 2: [Recopy the Fact]

1. Justification

2. No

...

Fact N: [Recopy the Fact]

1. Justification

2. Yes

List of all Facts:

<ALL FACTS>

A.5 MFactSum evaluation

We present below the prompt used to evaluate the visual or textual recall of groundtruth summary facts. Specifically, this prompt tests whether each groundtruth fact is supported by the predicted summary.

B Additional Experiments

In Tables 6 and 7, we report the results using respectively Gemini 1.5 Flash and Qwen2.5-72B-Instruct in place of Gemini 2.5 Flash as the summarization model in our pipeline (Figure 1).

	vis-rec	text-rec	MFS	r1	r2	rlsum	METEOR
Transcripts (no video)	13.17	18.41	15.79	34.19	7.10	32.64	26.52
Built Screenplay (50 clips)							
random clips	14.20	18.68	16.44	33.79	7.12	32.13	26.78
silent clips	14.11	19.54	16.83	34.80	7.41	33.15	27.16
<i>our clips zero-shot (Qwen2.5-Omni-7B)</i>	14.88	19.72	17.30	33.82	7.14	32.17	27.15
<i>our clips two-shot (Qwen2.5-Omni-7B)</i>	16.88	20.00	18.44	34.25	7.40	32.57	27.45
Built Screenplay (reference clips)	16.45	19.04	17.75	34.86	7.37	33.14	27.21
Gold Screenplay	22.78	23.07	22.92	34.87	7.80	33.03	28.41

Table 6: **Evaluation results using Gemini 1.5 Flash for summarization.** Evaluations are made on the MovieSum test set. Column descriptions are the same as in Table 2. Best results are in **bold**.

	vis-rec	text-rec	MFS	r1	r2	rlsum	METEOR
Transcripts (no video)	17.27	23.92	20.59	41.88	10.41	40.08	29.88
Built Screenplay (50 clips)							
random clips	17.69	24.04	20.86	41.80	10.45	39.88	29.80
silent clips	18.56	24.28	21.42	42.20	10.66	40.10	30.16
<i>our clips zero-shot (Qwen2.5-Omni-7B)</i>	19.25	24.32	21.79	41.79	10.58	39.72	29.77
<i>our clips two-shot (Qwen2.5-Omni-7B)</i>	18.71	24.08	21.39	41.66	10.48	39.81	29.56
Built Screenplay (reference clips)	19.44	23.62	21.53	42.15	10.70	40.06	30.01
Gold Screenplay	28.77	27.81	28.29	43.55	11.32	41.43	31.47

Table 7: **Evaluation results using Qwen2.5-72B-Instruct for summarization.** Evaluations are made on the MovieSum test set. Column descriptions are the same as in Table 2. Best results are in **bold**.

C Human Evaluation of the Clip Selection Reference

We report the results of our human evaluation against the first annotator in Table 8.

	Shining (1980)	Dark Knight (2008)	Imitation Game (2014)	Black Panther (2018)	Average/Total
Precision	80.6	72.2	81.8	100.0	83.65
Recall	89.5	90.0	91.7	90.0	90.3
F1 Score	84.8	80.1	86.5	94.7	86.5
Nb reference clips	31	54	11	12	108

Table 8: **Human evaluation of the clip selection reference by the first annotator.** We report the Precision, Recall and F1 scores between the clip selection reference (see Section 3.2) and the human reference on all 4 movies. We also report the number of clips in the clip selection reference for each movie.