

Synopses of Movie Narratives: a Video-Language Dataset for Story Understanding

Yidan Sun¹, Qin Chao¹, Yangfeng Ji² and Boyang Li^{1*}

¹School of Computer Science and Engineering, Nanyang Technological University, Singapore.

²Department of Computer Science, University of Virginia.

*Corresponding author(s). E-mail(s): boyang.li@ntu.edu.sg;

Contributing authors: SUNY0053@e.ntu.edu.sg; CHAO0009@e.ntu.edu.sg;
yangfeng@virginia.edu;

Abstract

Despite recent advances of AI, story understanding remains an open and under-investigated problem. We collect, preprocess, and publicly release a video-language story dataset, Synopses of Movie Narratives (SYM_{ON}), containing 5,193 video summaries of popular movies and TV series. SYMON captures naturalistic storytelling videos for human audience made by human creators. As a prototypical and naturalistic story dataset, SYMON features high coverage of multimodal story events, abundant mental-state descriptions, and large semantic gaps between the visual and the textual modalities. We establish benchmarks on video-text retrieval and zero-shot alignment on movie summary videos, which showcase the importance of in-domain data in story understanding. With SYMON, we hope to lay the groundwork for progress in multimodal story understanding.*

Keywords: video-text retrieval, video-language dataset, story understanding

1 Introduction

Stories are complex artifacts that succinctly encode the human experience. The topic of computational story understanding has attracted substantial research interest over the years (Emelin et al, 2020; Chambers and Jurafsky, 2008; Ferraro and Van Durme, 2016; Wang et al, 2021; Chaturvedi et al, 2016; Kim and Klinger, 2019), but multimodal story understanding (Tapaswi et al, 2016; Agrawal et al, 2016; Kim et al, 2017; Papalampidi et al, 2021) remains an under-investigated problem. A particular challenge is the lack of multimodal datasets that contain naturalistic and complete stories (*vis-à-vis* stories created solely for AI and story segments with incomplete dramatic curves) but are sufficiently short to match the compute available at academic research labs.

In this paper, we present a new video-language story dataset, named Synopses of Movie Narratives (SYM_{ON}). The dataset is sourced from “a movie in 5 minutes” videos on YouTube and contains naturalistic storytelling videos that summarize a movie, a TV episode, or an entire season of a TV series. The videos are created by human creators for a human audience. They are composed of selected clips depicting key story events and a narrator’s retelling of the story, providing condensed yet complete storylines. SYMON includes 5,193 user-generated video summaries for a total length of 869 hours. As an illustration of the data, we show one short snippet in Figure 1.

SYM_{ON} is distinguished from other datasets for its extensive coverage of multimodal story events and mental states of story characters as well as an easy-to-process short format. Its contents are representative of

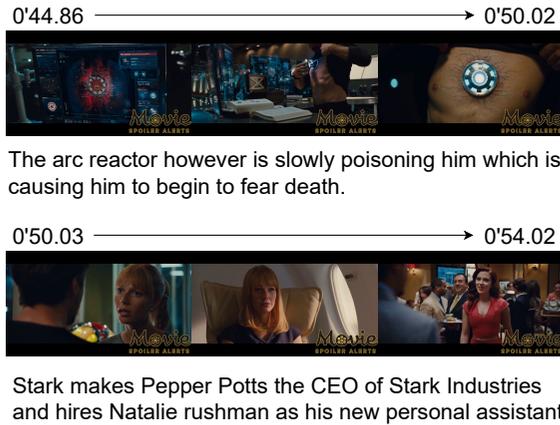


Fig. 1 A short video clip from SYMON, shown as a sequence of frames and the narration text. The timestamps indicate starting and ending times.

storytelling techniques employed in naturalistic multimodal stories. We contend that SYMON is a prototypical story dataset that will facilitate the understanding of crucial story elements.

Although computational story understanding can be traced back to (Schank et al., 1975; Wilensky, 1983), today it is a comparably niche research area. To explain how SYMON facilitates the understanding of crucial story elements, in the following we briefly analyze the story understanding task and the requirements it imposes on datasets.

Converging evidence points to the central role of events and their causal relations in story understanding. First, most definitions of stories in narratology revolve around events and their causal and temporal relations (Tomashevsky, 1925; Rimmon-Kenan, 1983; Prince, 2003; Schmid, 2010; Genette, 2014). A well-known definition due to Forster (1985) is that “the king died and then the queen died” is insufficient but “the king died and then the queen died of grief” qualifies as a story (which Forster called a plot) because the latter introduces a causal relation between the two events. Second, cognitive science research indicates causal relations are key to human story understanding (Graesser et al, 1991; Trabasso and van den Broek, 1985). Third, early works in story generation often assemble events into a causal network (Meehan, 1976; Young et al, 1994; Riedl and Young, 2010).

In Figure 2, we show a causally related network of events from *The Lord of the Rings: The Return of the King*. Following Trabasso et al (1989a), we

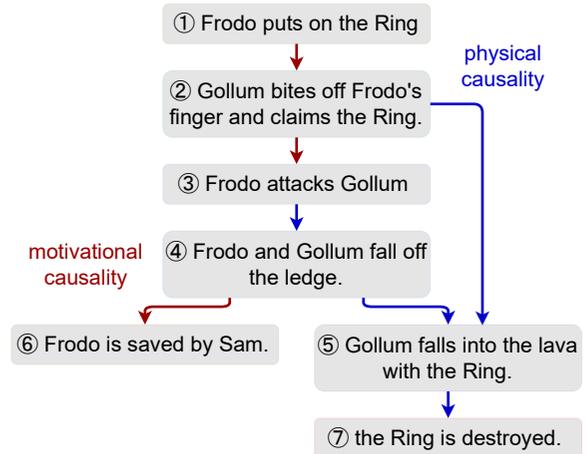


Fig. 2 A event graph of causal connections between 7 events from *The Lord of the Rings*.

differentiate causal relations as either physical or motivational. Physical causality is mediated by physical laws, whereas motivational causality involves the voluntary actions of an individual caused by internal motivation, which may be reaction to external stimuli. For example, Gollum’s desire for the Ring causes him to bite off Frodo’s finger.

An event can be described at several different levels of granularity; choosing the right level is important when building such event networks. Too much detail could hinder the identification of important events that drive the plot. For instance, in the LSMDC dataset, the event “Gollum bites off Frodo’s finger and claims the Ring” is subdivided into Gollum lifting his hand, Gollum biting, Frodo screaming, Gollum celebrating, Gollum holding the Ring, etc (Figure 3), potentially causing a computational system to be overwhelmed and lose the forest for the trees. On the other hand, too little detail is also detrimental. Coincidentally, cognitive research indicates the existence of a privileged level of event granularity that strikes the right balance between efficiency and information loss and is dominant in human cognitive processing (Barker and Wright, 1955; Abbott et al, 1985; Rosch and Lloyd, 1978; Zacks and Tversky, 2001).

We aim to empirically verify if SYMON portrays stories at the privileged level of event granularity. Nevertheless, it is difficult to determine the right granularity of event *a priori*, as event descriptions are also moderated by the story context. For example, if the story is about what happens after the Ring is destroyed, the fight preceding its destruction may be abridged

or omitted. Instead of manually coding event granularity, which would be highly subjective, we adopt human-written story summaries from Wikipedia¹ as gold references, which the cognitive science theory would posit to be at the privileged level. In Section 4, we show that SYMON contains significantly more overlap with the Wikipedia summaries than similar multimodal movie story datasets. This also makes SYMON more likely to contain complete dramatic curves (Li et al, 2018) than other datasets.

Another advantage of SYMON is that it contains abundant descriptions about mental states of story characters. As shown in Figure 2, motivational causality plays an important role in relating events. Character emotions and intentions help in understanding motivation causality (Trabasso et al, 1989b; Riedl and Young, 2010). For example, in Figure 1, Tony Stark’s fear of death motivates him to make arrangements for his company. Experimentally, we show that SYMON describes mental states of story characters, such as their emotions and intentions, much more frequently than comparable datasets (see Section 5).

Yet a third characteristics of SYMON is the use of storytelling techniques in SYMON causes substantial disparity in the semantic content of the visual and textual modalities, which we refer to as the cross-modality semantic gap (see Section 6). These techniques, such as omission of events easily inferred from context and the use of symbolic and indirect visual representation, are widely employed in multimodal storytelling but are rarely covered in traditional video-text datasets. Through a video clip sequencing experiment and a principled probabilistic analysis, we estimate the cross-modality semantic gap of SYMON to be about 31.4%, whereas comparable datasets have gaps that are either too large to bridge or too small to provide a meaningful challenge. Hence, we argue that SYMON is uniquely suited to computational understanding of multimodal storytelling techniques.

We propose cross-modal retrieval and sequential cross-modal alignment as evaluation tasks for story understanding on SYMON. Correctly matching the video clips and text snippets of SYMON requires understanding of the story semantics in order to bridge the aforementioned cross-modality semantic gap. Moreover, retrieval and alignment tasks can be objectively evaluated and annotated, whereas annotating storytelling techniques would require substantial subjective

judgments. Together, the weakly supervised SYMON and the fully annotated YMS dataset (Dogan et al, 2018) form a complete test suite and a new challenge for the multimodal research community.

As a baseline for future research, we present a simple neural network equipped with a history retrieval mechanism for tackling the story understanding tasks. Results indicate that modeling long-range dependencies provides benefits for story understanding.

In summary, we make the following contributions in this work:

- We collect, preprocess, and publish a large-scale movie summary dataset, SYMON. We demonstrate that large-scale and in-domain data boost performance on movie understanding tasks.
- We experimentally characterize the SYMON dataset in three aspects crucial for story understanding: (1) event descriptions at the right level of granularity and story completeness. (2) the amount of mental-state descriptions, and (3) the semantic divergence between text and video. The experiments underscore the differences between SYMON and existing datasets.
- To facilitate future research, we establish baselines for bi-directional text-video retrieval on SYMON and zero-shot video-text alignment on YMS. We demonstrate the effectiveness of modeling long historical context in story understanding.

2 Background and Related Work

2.1 Background on Story and Event Structures

The structuralist school of narratology has offered different variations for the definition of a narrative or a story (Prince, 2003), but most theorists regard events and their relations to be key to story structures. Forster (1985) argued that a minimal story must contain two events that are causally related. Tomashevsky (1925) argued for the necessity of causality whereas Genette (2014) believed one event alone can constitute a story. Schmid (2010) pointed out that it can be difficult to ascertain the existence of causal relations and suggested temporal relations should suffice. It follows that a story understanding dataset should offer opportunities to understand events and their relations.

Given the importance of events in story understanding, we need to ask how detailed should the event descriptions be. One event can be described as a whole or as several parts. For example, the event “going to

¹<https://github.com/markriedl/WikiPlots>

school” has parts like “getting dressed”, “waiting for the school bus”, “getting on the bus” and so on, and each part can be further decomposed. Such part-whole relations form an event partonomy, a cognitive structure that plays important roles in narrative comprehension and event understanding (Zacks and Tversky, 2001; Sridhar et al, 2010).

Cognitive science research shows there is a privileged level in the event partonomy that people naturally adopt when talking and thinking about events (Barker and Wright, 1955; Abbott et al, 1985; Rosch and Lloyd, 1978; Schank and Abelson, 1978; Zacks and Tversky, 2001). The right level in the event partonomy offers enough information at reasonable efficiency. Therefore, we argue that a good story understanding dataset should describe events at this privileged level.

It is also worth mentioning the structuralist perspective of the story arc, or the constituents of a complete story. Aristotle (circa. 335 BC) stated that a story contains a beginning, a middle and an ending. Freytag (1863) proposed a five-element structure consisting of the exposition, rising actions, the climax, falling actions, and the denouement. The most important variable in Freytag’s structure is dramatic tension, which rises to its climax through rising actions before falling to the denouement. Several prominent variations of story arcs exist and Li et al (2018) attempted to offer a unifying annotation scheme. Inspired by these theories, we argue that a good story understanding dataset should contain complete stories that cover as much of the story arc as possible.

2.2 Background on Emotion, Motivation, and Intention

The creation of story characters with realistic motivation and behaviors is a hallmark of storytelling excellence (Burgess, 2022; Murakami, 2022). Cognitive studies indicate that emotions play a central role in human motivation and decision making (Naqvi et al, 2006). More specifically, the cognition-appraisal-coping framework (Smith et al, 1990; Gratch and Marsella, 2004) posits that external stimuli are processed by various cognitive processes, whose results are appraised for potential emotional arousal. Some emotions, especially negative ones, require regulation and coping; these emotions motivate either external behaviors that attempt to change the environment, or psychological adaptation of the person’s own stances and relations with the external world.

Due to space limitations, the above discussion of emotion and motivation is vastly simplified. But it is apparent that emotions, motivations, and intentions of story characters are a crucial aspect of stories. Indeed, Trabasso et al (1989b) proposed motivation as one type of causal relations in stories; Riedl and Young (2010) introduced motivations and intentions into plan-based story generation. Several text-based datasets describe this aspect of stories (Rashkin et al, 2018a,b; Sap et al, 2019). Therefore, we argue that a good story understanding dataset should facilitate computational understanding of how emotions of the story characters create intentions and motivate behaviors.

2.3 Datasets for Event and Story Understanding.

Existing datasets annotate the temporal aspects of events, such as temporal precedence and duration (UzZaman et al, 2013; Chambers et al, 2014; Ning et al, 2020; Zhou et al, 2021; Vashishtha et al, 2019, 2020), and causal relations between events (O’Gorman et al, 2016; Roemmele et al, 2011).

Several datasets explore individual components of stories, including sentence ordering (Gangal et al, 2021), social norms and moral consequences (Emelin et al, 2020), plausible antecedent events (Brahman et al, 2021), intentions and effects on mental states (Rashkin et al, 2018b), high-level story structures (Ouyang and McKeown, 2015; Li et al, 2018), and story character descriptions (Brahman et al, 2021). Sap et al (2019) consider relations between events, persona, and mental states. Some datasets aim at summarization of screenplays or conversation transcripts (Gorinski and Lapata, 2015; Papalampidi et al, 2020; Chen et al, 2021).

Some QA datasets are conditioned on comprehension of story texts, such as MCTest (Richardson et al, 2013), NarrativeQA (Kočíský et al, 2018), and FriendsQA (Yang and Choi, 2019). Multimodal counterparts include MovieQA (Tapaswi et al, 2016), TVQA (Lei et al, 2018), and Pororo (Kim et al, 2017). However, not every question in the QA datasets requires in-depth narrative understanding.

2.4 Video-Text Movie Story Datasets.

A number of datasets supply story content extracted from movies. The Large-Scale Movie Description Challenge (LSMDC) (Rohrbach et al, 2017) provides meticulously detailed language descriptions of the video content initially intended for the visual impaired.

Table 1 Comparison of video description datasets with story content.

	Video hours	#Videos (#Clips)	#Sent	Vocab.
CMD (Bain et al, 2020)	1,270	3,605 (33,976)	35,681	15,272
MovieNet ² (Huang et al, 2020)	2,000	1,100		
LSMDC Rohrbach et al (2017)	147	200 (128,085)	128,118	22,500
Pororo (Kim et al, 2017)	20.5	171 (16,066)	43,394	
MovieGraph (Vicol et al, 2018)	94.0	51 (7,637)	20,849	
SYM _{ON} (Ours)	869	5,193	683,611	40,116

Although these descriptions are highly accurate, they are not representative of real-world storytelling.

YouTube Movie Summary (YMS) (Dogan et al, 2018) contains 94 YouTube movie summary videos with human-narrated storylines. The Condensed Movies Dataset (CMD) (Bain et al, 2020) gathers 7 to 11 key clips from each movie with one-sentence descriptions for each clip. Pororo (Kim et al, 2017) captures 20-minute cartoon episodes, in-show conversations, and human-written descriptions. MovieNet (Huang et al, 2020) annotates 2000 hours of movies with extensive annotations and aligned movies scripts. However, due to copyright considerations, at the time of writing, the video release is still pending.

Other types of video annotations have been explored, including semantic roles and event relations (Sadhu et al, 2021), character relationships and types of speech (Wu and Krahenbuhl, 2021), and movie graphs (Vicol et al, 2018).

In this work, we collect a large-scale, readily available, multi-reference dataset of human-curated movie summaries, SYM_{ON}. The dataset can be leveraged for various story understanding and generation tasks such as sequential text localization, story generation from video, and movie summarization. To our knowledge, SYM_{ON} is the largest dataset for short naturalistic storytelling videos.

3 Collection and Preprocessing

In this section, we describe the procedures for collecting and preprocessing of the SYM_{ON} dataset, and also report some data statistics.

3.1 Dataset Collection and Statistics

We apply the following procedure for data collection. First, we manually identify relevant YouTube channels by searching with keywords such as “movie summary”,

“movie recap”, and “movie shortened”. We download all videos from the identified channels and accompanying subtitles, which may be written by humans or automatically generated by YouTube. In subsequent processing, we do not distinguish between manual and automatic transcriptions, but we include a label in the dataset for automatic transcript. Videos without subtitles are excluded. Finally, we perform rule-based extraction of movie names from metadata and subtitles and discard videos that are not movie summaries.

This yields a total of 5,193 videos with an average length of 9.5 minutes and a total length of 869 hours. On average, the narration in one video contains 1,717 words or 131 sentences. The overall vocabulary size is 40,116. SYM_{ON} contains summaries for 2,440 movies and TV series, of which 432 have more than 1 summary. The most popular movie, *The conjuring*, has 12 summaries. On average, one movie or TV series in the 432 has 2.79 summaries. Compared to existing movie story datasets (see Table 1), SYM_{ON} has the largest collection of short storytelling videos as while as the largest vocabulary.

3.2 Preprocessing

Subtitle Masking. Some videos have subtitles embedded in the video. In tasks like text-to-video retrieval, the embedded subtitles may become a shortcut feature, causing neural networks to rely on only optical character recognition.

To eliminate shortcuts, we locate embedded subtitles and mask them out. For efficiency, we randomly sample 100 frames from each video and apply an accurate text detection technique (Baek et al, 2019). Observing that the subtitles are almost always at the same location in a single video, we take the minimum bounding box that can cover all embedded subtitles in all 100 frames as the masked region; we set all pixels in the region to black.

Punctuation Restoration. We acquire subtitle texts from YouTube directly. Sometimes the texts are the result of automatic speech recognition, which cannot recognize punctuation. To fix this, for every unpunctuated narration text, we generate punctuation with the approach proposed by Alam et al (2020).

Scene Segmentation. Later experiments require temporal segmentation of videos based on camera cuts. For this purpose, we run the dataset through the network of (Souček and Lokoč, 2020), which detects hard camera cuts. A scene, defined as a continuous shot between two cuts, lasts 2.2 seconds on average. This is similar to CMD, another movie dataset, whose scenes last 2.4 seconds on average. However, average scenes in ActivityNet (Caba Heilbron et al, 2015) and Kinetics-400 (Kay et al, 2017) last for 11.1 seconds and 30 seconds respectively. This shows camera cuts in movies are much more frequent than the user-generated videos in ActivityNet and Kinetics.

4 Dataset Analysis: Event Granularity and Story Coverage

As discussed in Section 2.1, to assist an AI system in acquiring the ability to understand stories like humans, it is desirable for the dataset to represent events at a level of granularity akin to human cognitive representation. This level of granularity should be neither too concise nor too detailed. From the perspective of event causality, the right level of event granularity should omit details that do not advance the storyline. In addition, it is desirable for the dataset to contain complete stories spanning the entire story arc.

In this section, we argue that SYMON provides event descriptions at the right level of granularity and relatively complete stories. We first present a qualitative comparison between three video-text story datasets. After that, we quantify the similarity between human written story summaries from Wikipedia and textual descriptions of video-text story datasets.

4.1 Qualitative Example

Figure 3 compares SYMON with two other video-text story datasets, CMD and LSMDC, as well as WikiPlots story summaries, using the same scene taken from *The Lord of the Rings: The Return of the King*. We identify 7 major plot events that form a sub-story arc with rising and falling tension:

- Frodo puts on the Ring. (rising action)

- Gollum bites off Frodo’s finger and claims the Ring. (rising action)
- Frodo attacks Gollum. (rising action)
- Frodo and Gollum fall off the ledge. (climax)
- Gollum falls into the lava with the Ring. (climax)
- Frodo is saved by Sam. (falling action)
- The Ring is destroyed. (falling action)

Despite having the shortest video run time (31 seconds), SYMON covers all 7 events. In comparison, the CMD text describes the final outcome, the destruction of the ring, but omits the rest completely. The gap between the brief CMD text and the 40-second video is hence substantial. The third dataset, LSMDC, is designed to help the audience with visual impairment. As a result, its text offers a blow-by-blow recount of the video. Although LSMDC texts match well with the video and may be useful for understanding individual shots, excessive details may complicate story understanding at the event level and obscure the big picture. For example, the description “Gollum suddenly lifts his hands to his face and bites” and subsequent sentences render it difficult to infer that Gollum reclaims the ring this way, whereas in WikiPlots and SYMON this is explicitly explained.

From a cognitive science perspective (see Section 2), LSMDC descriptions sit at a lower level on the event partonomy than SYMON, and CMD descriptions sit at a higher level. SYMON has the best alignment with the human-written summary from WikiPlots, which we argue is at the privileged level of event partonomy that balances information and efficiency. In the next section, we verify this empirically.

4.2 Quantitative Assessment of Story Coverage

We computationally estimate the extent to which stories in CMD, LSMDC, and SYMON cover the sentences in WikiPlots. We use a three-step procedure for computing story coverage. First, we match movie summary in our dataset to their WikiPlots summaries by name³.

Second, we estimate if a sentence from the video narration is equivalent to a sentence in WikiPlots using a distance metric. We experiment with two distance metrics between input sentences a and b . The first metric is the average probability of a entailing b and b entailing a from the natural language inference classifier of Nie et al (2020). Alternatively,

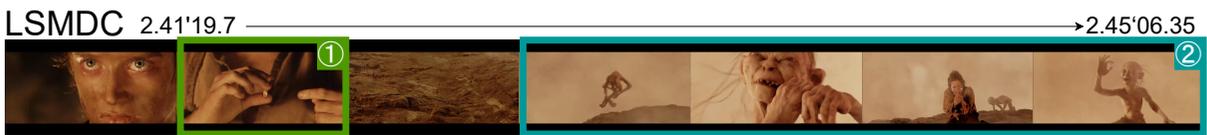
³For story coverage evaluation on SYMON, only movies whose plot summaries can be found in WikiPlots is included.



⑦ The Ring of Power sinks beneath the lava and is finally destroyed.



① he succumbs to the ring and claims it as his own, putting it on his finger. ② gollum finds the invisible frodo and attacks him, biting his finger off to reclaim the ring. ③ frodo attacks gollum in an attempt to reclaim the ring, and in the ensuing struggle, ④ they both fall off the ledge. ⑤ gollum falls into the lava with the ring and dies. ⑥ frodo clings to the side of the ledge and is rescued by sam. ⑦ as the ring disintegrates into the lava.



① Sam looks at Frodo, the Ring has finally taken him. Sam screams as Frodo puts the Ring on! He vanishes. Footprints moving across the ash covered cavern. Suddenly! Gollum smashes a rock down on Sam's head, knocking him to the ground! Gollum leaps on to the invisible Frodo! Close on: Frodo's footprints. ② Gollum clawing frantically, riding on the back of the invisible Frodo. Gollum suddenly lifts his hands to his face and bites. Frodo materializes as he drops to his knees, clutching his bleeding hand. He screams. The fires below roar in anger, red light blazes, and all the cavern is filled with a great glare and heat. Gollum dancing gleefully. Perilously close to the edge of the chasm. Gollum still triumphant on the edge of the precipice. The fiery lava casting an evil light. Frodo rises slowly, his eyes lock on Gollum. Time seems to stand still as each regards the other. The Ring glinting in Gollum's hand. ③ Suddenly! Frodo lunges at Gollum. Frodo thuds into Gollum, lifting him off his feet. ④ People lock in struggle plummet over the edge of the chasm. Sam watches in horror as people disappear into the chasm. ⑤ Gollum falls into the crack of doom. Instantly engulfed in the churning lava! ⑥ Angle on: Sam leans into the crack of doom. Frodo clinging onto the rock face with his one hand. Sam desperately reaches for him. Frodo is too far down. Frodo struggles to reach Sam. He can't reach with his bleeding hand, and falls back, still clinging valiantly to the rock with his good hand. Sam looks at Frodo imploringly. Frodo reaches up once more and this time Sam grasps his hand tightly. ⑦ The Ring sit on the river of lava for a brief moment then sinks away.

WikiPlots

① There, he succumbs to the Ring's power, putting it on his finger. ② But Gollum manages to bite off his finger and reclaim it. ③ They struggle together ④ and both fall off the ledge. ⑥ Frodo manages to cling on, and is pulled up by Sam, ⑤ but Gollum falls and dies; ⑦ the Ring, which fell with him, disintegrates in the lava.

Fig. 3 Compare textual granularity of CMD, LSMDC SYMoN, and WikiPlots. The frames are respectively sampled from the SYMoN, CMD, and LSMDC video summaries of the movie *The Lord of the Ring: The Return of the King*. Video-text correspondence are manually marked by colors and numbers. We present the WikiPlots summary of the same movie segment as the last section.

Table 2 Estimated alignment percentage with two distance metrics and Dynamic Time Warping.

Distance Metric	CMD	LSMDC	SYMON
Symmetric Entailment	10.6%	18.1%	44.3%
SBERT Cosine	28.7%	38.3%	81.2%

from Sentence-Bert (SBERT) (Reimers and Gurevych, 2019), we compute the cosine distance between the sentence embeddings.

Finally, we find the best correspondence between two texts using Dynamic Time Warping (DTW) (Berndt and Clifford, 1994). Briefly, DTW utilizes dynamic programming to find the minimum-cost correspondence between two sequences, the WikiPlots sentence sequence A , and the narration sentence sequence B . We refer readers to Appendix A for details of the DTW algorithm. DTW has two model parameters, δ_A and δ_B , which denote the costs for skipping a sentence in sequences A and B . We determine them from manual annotations, as described below.

We manually labeled 1,000 sentences from CMD for correspondence to Wikiplots stories, 1,000 sentences from SYMON, and 2,000 sentences from LSMDC, which amount to 83 texts from CMD, 8 texts from SYMON, and 4 texts from LSMDC, due to different text lengths. A second annotator labeled a small portion of data from each dataset to compute inter-rater reliability. Cohen’s Kappa on SYMON, CMD, and LSMDC are 0.86, 0.59, and 0.33 respectively. We believe the low agreement on LSMDC is caused by the much longer lengths of LSMDC texts than Wikiplots stories, which led to difficulties in annotating precise correspondence.

On this annotated dataset, we use grid search to find the optimal δ_A and δ_B that maximize the accuracy of DTW, which is defined as

$$\text{Accuracy} = \frac{1}{2} \left(\frac{M^{\text{wiki}}}{N^{\text{wiki}}} + \frac{M^{\text{text}}}{N^{\text{text}}} \right). \quad (1)$$

Here M^{wiki} and N^{wiki} are the number of correctly matched and the total number of WikiPlots sentences, respectively. M^{text} and N^{text} are the number of correctly matched and the total number of sentences in the video narration. We do not directly optimize alignment percentage because doing so will cause DTW to inflate alignment percentage by matching sentences incorrectly.

With the optimal δ_A and δ_B , we perform DTW again and calculate story coverage as the proportion of WikiPlots sentences matched with narration sentences,

$$\text{Coverage} = \frac{1}{K} \sum_i^K \frac{M_i^{\text{wiki}}}{N_i^{\text{wiki}}}, \quad (2)$$

where the summation is over all WikiPlots summary texts and K is the number of WikiPlots movies appearing in the video dataset. In the i^{th} WikiPlots text, M_i^{wiki} denotes the number of matched sentences and N_i^{wiki} denotes the total number of sentences.

Table 2 shows the story coverage results. Though the two distance metrics result in different coverage numbers, the ranking between methods remains unchanged. Of the three datasets, SYMON provides the highest story coverage. LSMDC takes second place, partially benefiting from the extreme lengths of its texts. The results demonstrate that SYMON stories cover more major plot points than the other two datasets. As most WikiPlots summaries are complete stories, and are probably at the privileged level of event paratomy, the findings suggests that SYMON contains relatively complete story arcs and has suitable granularity for story understanding.

5 Dataset Analysis: Mental-state Descriptions

As explained in Section 2.2, characters’ mental state such as emotions, intentions, and motivations play crucial roles in the causal network of events and in story understanding. In this experiment, we measure the frequency of words related to emotions, intentions, and motivations in the text associated with the videos. For emotional words, we adopt the WordNet-feelings dataset (Siddharthan et al, 2018), which includes 11,387 emotion-related words identified by human experts. For motivation and intention words, we find 200 nearest neighbors of the words “motivation” and “intention” using 300-dimensional fastText embedding⁴ (Bojanowski et al, 2017) trained on Wikipedia and Common Crawl. We stop at 200 words as we find additional neighbors to be mostly irrelevant.

Table 3 reports word frequencies for every thousand words in four video-language datasets. We observe that SYMON employs mental-state words with the highest

⁴Acquired from <https://github.com/facebookresearch/fastText/blob/master/docs/crawl-vectors.md>.

Table 3 Frequency of words related to emotion, motivation, and intention per one thousand words.

	Emotion	Motivation	Intention
ActivityNet Captions	27.5	0.51	2.7
LSMDC	33.5	0.62	2.8
CMD	38.9	1.41	9.4
SYM _{ON} (all text)	57.6	1.58	23.9
SYM _{ON} (excl. commentary)	61	1.78	24

frequency. Interestingly, SYMON contains 2.5 times as many intention-related words as the next dataset, CMD, whose textual descriptions also focus on story content. LSMDC, containing literal descriptions of movie clips, is ranked at the third position. As a sanity check of the metric we adopt, we compute the same statistics for ActivityNet Captions (Krishna et al, 2017), containing matter-of-fact descriptions of actions in generic user-uploaded videos. As we expect, ActivityNet Captions uses the least of mental-state words. Overall, the ranking is consistent with the nature of the datasets, as story text describes mental states more often than literal descriptions of generic videos. SYMON appears to be the most prototypical story dataset.

One possible caveat of the above analysis is that it does not exclude narrator commentary, which may have inflated the word frequencies. We further recruited human annotators from Amazon Mechanical Turk (AMT) to label 36,955 sentences (8.5% of all text from SYMON) as either commentary or non-commentary. The workers annotated 8.9% of sentences as commentary. Over the remaining 91.1% or 33,666 sentences, we recompute the frequency statistics and find marginal changes. We conclude that the high frequencies are indeed due to the story content rather than narrator commentary. Details of the AMT experiment are in Appendix B.

6 Dataset Analysis: Cross-modality Semantic Gap

Multimodal storytelling often utilize the modalities in a complementary manner. For example, Iyyer et al (2017) study comics and find that imagery and text are interwoven in the storytelling and understanding any single modality only is insufficient for story understanding.

Similarly, in SYMON, the story narration and the video often are not simple duplicates, but provide complementary information. As a result, the text in SYMON often does not exactly match the video content, i.e. cross-modality semantic gap. In this section,



Fig. 4 This example shows three frames from *Silence of the Lambs*. The text (kill) describes the effect of the video (shooting).

we qualitatively analyze the cause and quantitatively estimate the semantic gap.

6.1 Categories of the Semantic Gap

Through manual examination of the data, we provide a categorization of the mismatch between the narration text and the video. Though the list of categories is likely incomplete, we believe it can provide insight into the types of multimodal storytelling techniques utilized by videos in SYMON.

Moving along the causal chain. Instead of describing precisely the event shown in the video, the narrator describes its immediate cause or effect. For example, in Figure 4, the video shows Gum getting shot and then lying on the ground coughing up blood. The text describes the event as “Clarisse kills Gum.” The expectation is that the viewer can infer that “being killed” is the direct effect of getting shot and therefore can match the text with the video. By utilizing this strategy, the narrator avoids stating the obvious and provides information that complements the video.

Symbolic representation of events. The storytelling in SYMON sometimes employs a symbolic object to represent the occurrence of an event. For example, in Figure 5, the video shows Dolores Umbridge sitting on the headmaster’s chair while the narrator states “Umbridge becomes the new headmistress”. Here, the chair represents the event of Umbridge becoming headmistress. The audience is expected to have commonsense knowledge and be aware of the association between the high-back armchair and someone’s status as the headmaster or headmistress.

Reference to long-range context. Narrator may make references to previous events when explaining the current event, so that long-range context is needed to match current video segment with text. For example, in Figure 6 the video shows a ship full of wine barrels moving toward a town and the text says the company is smuggled into Asgaroth. To understand the connection

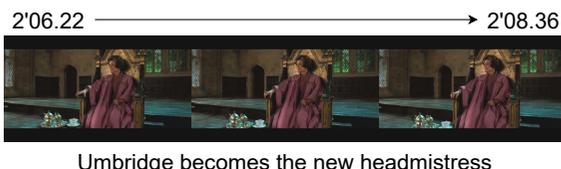


Fig. 5 An example from *Harry Potter and the Order of the Phoenix*. A symbolic object, the chair, is used to represent the event Dolores Umbridge becoming headmistress.

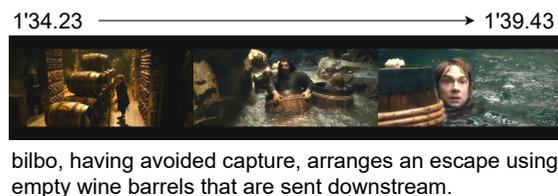


Fig. 6 Example from the film *The Hobbit*, The video segment on the first line gives the context (the dwarfs escaping in barrels), which is required to understand the connection between the video and the text on the second line.

between video and text, the audience needs to recall earlier in the video, the dwarfs hid themselves in the wine barrels.

Explicitly stating the implied. The text explicitly explains events that were only implied in the original movie or TV show. In a full-length movie, some events may be omitted and left for the audience to deduce. However, it is much harder to deduce missing events in a short and fast-tempo summary video. Therefore, these events usually need to be explicitly explained in SYMON. Since the events are not present in the movie, there is no matching imagery to the narration and the semantic gap is inevitable. In the example shown in Figure 7, the video shows Lector watching Chilton disembarking from a plane. In the movie, the audience knows that Lector is a cannibal and therefore can understand that Lector means to eat Chilton. In the recap video, the event “implying Lector is going to eat him” is explicitly stated in the text to help audiences better understand the story.

Mental state descriptions. The narration explicitly describes the mental state of the characters. In a



The film ends with Lector saying that he is having an old friend over for dinner. and we see that Lector is following Chilton, implying that Lector is going to eat him.

Fig. 7 Example from film *Silence of the Lambs*, the text describes an event implied in the movie (Lector eating Chilton). Because there is no matching video of this event the video shows Chilon disembarking a plane.



the hobbits begin climbing stairs carved in the cliff face that leads to a secret tunnel into mordor. unaware that gollum plans to kill them and take the ring

Fig. 8 An example from the movie *Lord of the Rings: Return of the King*. The video does not show Gollum’s intention but the text explains it to further the story.

short 5-minute summary video, the characters’ emotions, intentions, and motivations need to be directly explained to help audiences understand the story. In the example shown in Figure 8, the video shows the hobbits climbing stairs with Gollum. Gollum’s intention to kill the hobbits and claim the ring is not apparent, but is however crucial for understanding the later events in the story. The narrator makes it explicit.

Summarizing dialogues in the original. The narration may summarize the content of a dialogue in the original movie or TV show, while the video shows two people talking. See an example in Figure 9. In most videos in SYMON, dialogues in the original movie are replaced with a narrated summary. This is because the original dialogue is usually long and contains a lot of information; it is not easy to comprehend the relation between the dialogue and other events in a short amount of time. The summarization omits details irrelevant to the plot and helps the viewer grasp the key content quickly. At the same time, the video often only shows people conversing. The narrator sometimes attempts to show images related to the dialogue content, but for abstract ideas like Voldemort’s intention, finding relevant imagery can be challenging, leaving substantial semantic gaps between the two modalities.



Sirius Black mentions that Voldemort is after an object he did not have during his previous attack

Fig. 9 This example shows two frames from the movie *Harry Potter and the Order of the Phoenix*, the video shows two people having a conversation. The text explains the content of the conversation.

Discussion. The examples and analysis above illustrate how videos in SYMON deliberately utilize the two modalities in a complementary manner to achieve storytelling goals. This is in stark contrast to most existing visual-language datasets, such as Charades (Sigurdsson et al, 2016), ActivityNet Captions (Krishna et al, 2017), and LSMDC (Rohrbach et al, 2017), which are carefully curated or selected so that the two modalities match as closely as possible.

Such semantic gaps are not merely the result of SYMON being short summaries of stories. Rather, most of the causes for the semantic gap we analyzed (omission of events on the causal chain, use of symbolism, reference to past events, etc) are classic film-making techniques ubiquitous to storytelling videos in the wild. For example, in the famous montage of the *Godfather*, the scene cuts between a series of brutal murders to Michael Corleone standing at an altar as the godfather of Connie’s baby. Though never explicitly shown, the audience is expected to infer the cause of the murder is Michael’s order. In addition, Micheal’s status as the baby’s godfather symbolizes the fact that he has become the godfather of a criminal organization. In this sense, SYMON provides a representative benchmark for story video comprehension in the wild. Next, we report experiments designed to estimate the extent of video-text correspondence.

6.2 Quantifying the Semantic Gap

In this section, we first introduce the task of video segment ordering, whose results enable principled statistical analysis of the semantic gap between the two modalities. After that, we analyze the performance of neural networks.

Similar to event / sentence ordering (Liu et al, 2018; Devlin et al, 2019), we predict the correct ordering of two consecutive video segments separated by a hard camera cut. This is a binary classification task with two outcomes: video segment 1 preceding segment 2 or vice versa. To create a balanced label distribution, we

Table 4 Performance of human annotators on the video segment sequencing task.

		Video-text		Total	Estimated Gap
		Correct	Wrong		
SYMÓN	Video-only	Correct 75.6%	Wrong 2.7%	82.8%	31.4%
		Correct 14.5%	Wrong 2.7%	17.2%	
Total		90.1%	9.9%		
CMD	Video-only	Correct 50.5%	Wrong 12.6%	63.1%	69.9%
		Correct 24.0%	Wrong 12.9%	36.9%	
Total		74.5%	25.5%		
LSMDC	Video-only	Correct 90.3%	Wrong 0.8%	93.0%	22.9%
		Correct 6.2%	Wrong 0.8%	7%	
Total		96.5%	3.5%		

Table 5 Cohen’s Kappa for SYMON, CMD, and LSMDC.

	SYMÓN	CMD	LSMDC
Text + Video	87.9%	61.7%	90.5%
Video Only	82.7%	66.3%	86.9%

randomly flip the ordering of the two video segments. We extract the text description that spans the same duration as the two video segments and expand the text to sentence boundaries. From 5,193 (4,733 for training) videos, we generate 136,220 (122,751 for training) data points, each consisting of two video segments and one text snippet.

Human Annotation and Analysis. We use human annotation to estimate the semantic gap on SYMON, CMD, and LSMDC. For each dataset, we recruited 4 human annotators to label 1,000 pairs of video segments both with and without text. We assigned two annotators to the text + video condition and two to the video-only condition. In each condition, both annotators labeled the same 200 data points to compute inter-rater agreement, see Table. 5. Note that the Cohen’s Kappa for CMD is lower than LSMDC and SYMON. This is because human annotators found it difficult to identify the order of a significant number of video pairs in CMD and resorted to random guessing.

We show the results in Table 4. In the video-only condition, the human annotators are correct 82.8% / 63.1% / 93% of the time; when additional text is available, they achieve an accuracy of 90.1% / 74.5% / 96.5%. On 2.7% / 12.9% / 0.8% of the data, annotators in both conditions made the wrong decision.

We conduct a probabilistic analysis to quantify the cross-modality semantic gap. First, we assume that, when humans do not know the answer, they can make random guesses that are 50% correct; knowing the answer, they always answer correctly. Thus, we need to distinguish knowing the correct answer from making the correct decision. We let K_{text} denote the humans in the text-video condition know the right answer and K_{video} denote the humans in the video-only condition know the right answer. Further, we assume that random guesses are independent. We want to determine if a text provides grounding to the corresponding video segments. Note that the effect of grounding is observable only when humans cannot determine the ordering from the video segments alone but can do so with the help of additional text. This is captured by the probability $P(K_{text} | \neg K_{video})$. When the video segments already contain order information, the text may still provide grounding but we cannot discern its effect from the experiment. With the above assumptions, we can show that

$$P_{\text{SYMON}}(K_{text} | \neg K_{video}) = 68.6\% \quad (3)$$

Thus, we estimate 68.6% of the texts provide sufficiently accurate descriptions of both video segments, from which sequencing decisions can be made. Similarly, we calculate the percentage of video grounding in text to be 30.1% and 77.1% for CMD and LSMDC respectively. We refer interested readers to details in Appendix C.

The statistical analysis largely agrees with our observations. CMD provides one-sentence summary for each 1-3 minute video, leaving much video content not described and a large gap (1 - 30.1% = 69.9%) between the two modalities. In comparison, textual descriptions in LSMDC are extremely detailed and closely match the video. The estimated gap of LSMDC (1 - 77.1% = 22.9%) is higher than our expectation, but the estimate may have very high variance and the actual semantic gap is likely much smaller. This is because that the gap is computed from the percentage of video segments wrongly ordered in both experimental conditions, which is a meager 0.8% for LSMDC and may easily be dominated by inevitable measurement noise. Had this value been 0, the estimated gap would be 0.

Finally, the semantic gap in SYMON is neither too large nor too small. We believe the moderate semantic gap reflects the use of storytelling techniques in the data and provides an interesting yet attainable challenge for computational understanding.

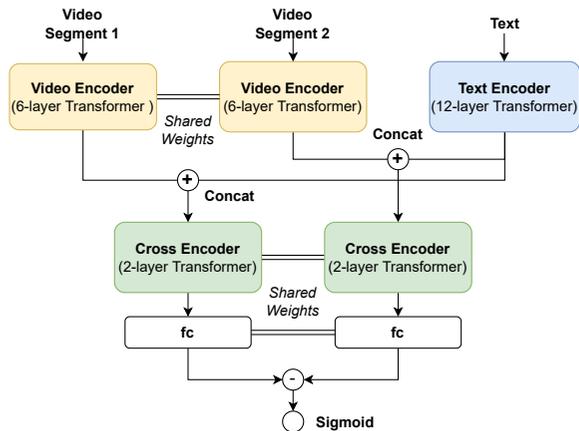


Fig. 10 The network architecture for the temporal ordering task. The double vertical lines indicate weight sharing between modules.

Neural Network Baselines. Similarly to the human annotators, we subject neural networks to two experimental conditions, the video-only and the text + video. We finetune three pretrained modules, the text encoder, the video encoder, and the cross-modality encoder from the UniVL model (Luo et al, 2020). The two video segments are encoded separately and their features are concatenated with the encoded text feature. After that, the two groups of features go through the cross encoder independently, yielding feature vectors f_1 and f_2 . With parameter w , the prediction is

$$P(\hat{y} = 1) = \sigma(w^\top f_1 - w^\top f_2), \quad (4)$$

where $\sigma(\cdot)$ is the sigmoid function and \hat{y} is the predicted class. Figure 10 shows the overall network architecture. For the network in the video-only condition, we replace the textual feature fed into the cross encoder with an all-zero vector and keep the rest of the network unchanged.

To cover as much data as possible, we adopt a special dataset split, containing Set A of 2,444 videos, Set B of 2,289 videos, and a validation set of 500 videos. Each network is trained on Set A and tested on Set B, and then trained on Set B and tested on A. We report the average test accuracy. We tuned hyperparameters extensively on the validation set and select the training epoch with the highest validation accuracy. To avoid test data leak, we put all videos of the same movie or movie franchise into the same set.

We create another baseline from the pretrained Merlot model (Zellers et al, 2021), which has been trained to sequence video frames. We extract the center frame from the two video segments and pass them with the

Table 6 Temporal order prediction accuracy of the text-aware and visual-only models.

	Text + Video	Video Only
Merlot	52.6%	-
UniVL-SYMON	69.1%	63.4%
Human annotator	90.1%	82.8%

corresponding text to the Merlot model, which then predicts which frame comes first. We finetune the Merlot model on SYMON. More settings and details can be found in the Appendix.

Table 6 lists the results. The network we trained (UniVL-SYMON) makes correct sequencing decisions 63.4% of the time from video only and 69.1% after incorporating textual information. Merlot achieves 52.6% with both video and text. Both perform substantially below human accuracy, indicating room for improvement.

Object and Action Analysis. We examine the match between text and video with object detection and action recognition. First, from the UniVL-SYMON sequencing models, we sort texts into helpful and unhelpful based on how much adding text improves the confidence of the correct answer. We take 5% most helpful texts, which improve the confidence the most, and 5% least helpful texts. Meanwhile, we run FasterRCNN (Girshick, 2015) trained on Open Images V4 (Kuznetsova et al, 2020) to detect 600 object classes on video frames, and 3D-ResNet (Hara et al, 2018) trained on Kinetics-700 (Kay et al, 2017) to detect 700 action classes. After that, we match the identified objects and actions to the texts. The Appendix contains more details.

Table 7 shows that the most helpful texts contain relatively 18.8% more recognizable objects and 25.0% more actions than the most unhelpful texts. This suggests that textual references to objects and actions in the video may have contributed to the temporal ordering task. In addition, detected objects and actions appear rather scarcely in the text descriptions, which contain 83 words on average. We attribute this observation to the semantic gaps explain earlier.

We further investigate the phenomenon by inspecting two examples (Figure 11) from the 5% most helpful texts and the 5% least helpful texts, respectively. We observe that the helpful text mentions objects, such as the cauldron and the book, from both video segments, thereby providing ordering information. In comparison, the unhelpful text mentions rare object and action

Table 7 Number of words that match exactly the detected object names or action names per text description.

	Objects Detected	Actions Recognized
Unhelpful Text	0.16	0.16
Helpful Text	0.19 (+18.8%)	0.20 (+25.0%)

such as cat costume and jewelry robbery, which are difficult for the network to learn. Similarly, matching the text “the mother refuses her son” and the discussion shown in video is not straightforward and would require identity tracking and event understanding.

This case study highlights challenges in understanding realistic videos. It also points to the integration of multiple AI capabilities, including visual recognition on the long tail, long-term identity tracking, and event understanding, as potential future directions.

7 Evaluating Story Understanding: Retrieval and Alignment

In order to evaluate how much computational models understand the multimodal stories of SYMON, we need a task that satisfies the following criteria. First, it should be strongly correlated with the level of story understanding that the models can perform. Second, it should be easy to automatically evaluate. Third, it should be relatively easy for humans to annotate, or is compatible with unsupervised learning, so that we can scale the training data.

We propose that one suitable task is bi-directional text-video retrieval. The task requires retrieving the narration text corresponding to a given video segment and the other way around. Retrieval may seem straightforward, but significant story understanding is required to bridge semantic gaps between the two modalities (as discussed in Section 6.1). In addition, we surmise that annotating text-video correspondence is easier than motivational states or causal relations.

In this section, we establish baselines on the task of video-text retrieval on SYMON and video-text alignment on YMS (Dogan et al, 2018), which serve as references for future research.

7.1 Approach

We employ pretrained text and video encoders from UniVL (Luo et al, 2020). Due to hardware constraints, we choose not to use the cross-encoder to accelerate training and inference. Empirically, the cross-encoder increases the memory footprint and forces us to

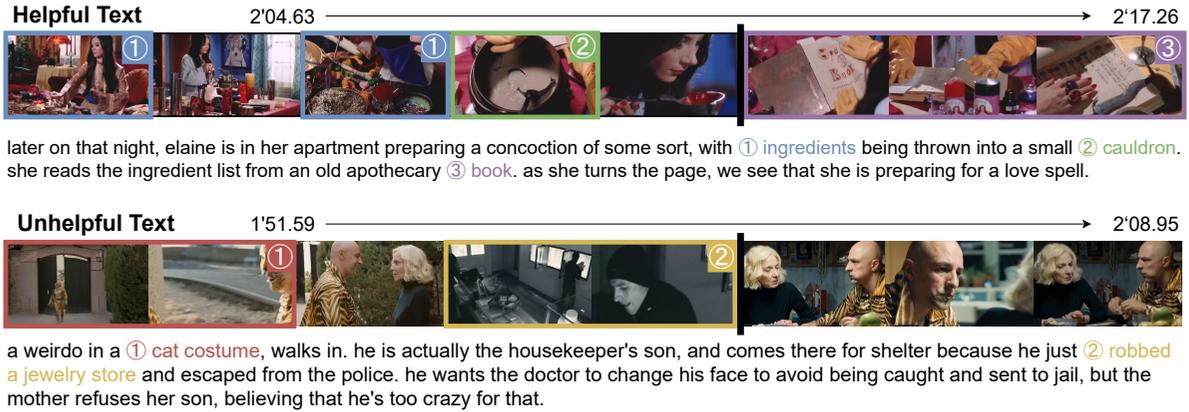


Fig. 11 Examples from the most and least helpful text clusters. Bound boxes of the same color in text and video frame denote video-text correspondence. The black line denotes the boundary between the two video segments to be ordered.

decrease the batch size, leading to degraded generalization for retrieval. In addition, the $O(n^2)$ complexity of exact inference, where n is the number of samples in each modality, leads to high computational cost, and the dot product incurs a much lower constant in the $O(n^2)$ complexity than the cross-encoder.

Following UniVL, we use S3D to extract 1 video feature per second. The video segments are trimmed or appended to 4 seconds. A video segment A is represented as 4 video feature vectors $\{v_1^A, \dots, v_4^A\}$. The texts are trimmed or appended to 64 tokens. The video features and text tokens are then passed through the video / text encoders. The encoded features for the i^{th} is denoted as t_i . The encoded features of the i^{th} video segment is v_i . The similarity between the text and the video is simply the dot product of the two vectors. With randomly sampled negative text features $t_k, k \neq i$ and video features $v_k, k \neq i$, we use the NCELoss (Gutmann and Hyvärinen, 2010):

$$L_{\text{NCE}} = \frac{1}{N} \sum_{i=1}^N -v_i^\top t_i + \log \left(\exp v_i^\top t_i + \sum_{k \neq i}^K \exp v_i^\top t_k + \sum_{k \neq i}^K \exp v_k^\top t_i \right), \quad (5)$$

where N is the total number of training samples and K the number of negative samples. Each video segment is matched to the text description spanning its duration. The text is extended on both ends to the nearest sentence boundaries.

One cause for the semantic gaps is long-range references to previous events. To tackle this issue, we implement a long-term video memory mechanism.

Briefly, we retrieve the most similar video segment from the past 10 video clips (about 44 seconds). Then, we append the video features and text tokens from the memory video clip to the current video features and text tokens respectively before the video / text encoders.

The similarity function used in memory retrieval attempts to match video segments partially and weigh the similarity with the number of people they share. For each text paragraph associated with the video segments, we perform named entity recognition and extract person names. The similarity between segments A and B is

$$\text{sim}(A, B) = (1 + E) \frac{1}{16} \sum_{i=1}^4 \sum_{j=1}^4 \cos(v_i^A, v_j^B), \quad (6)$$

where E is the number of shared participants in the two texts and $\cos(\cdot)$ is the cosine similarity. Intuitively, $\frac{1}{16} \sum_{i=1}^4 \sum_{j=1}^4 \cos(v_i^A, v_j^B)$ measures visual similarity and considers the contribution of each video pair equally. The $(1 + E)$ tracks participants in the video segments, two video segments are more like to be related if the participants overlap.

7.2 Retrieval on SYMON

The training, validation, and test sets contain 4,191 videos, 500 videos, and 502 videos, respectively. No movies or movie franchises appear in two sets simultaneously. The videos are divided into non-overlapping clips, each consisting of two scenes and having a mean duration of 4.4 seconds. YouTube videos often contain introduction and channel information at the beginning and the end, so we exclude 5% at each end of the videos.

For evaluation, we report recall at 1, 5, and 10 items (R@1, R@5, and R@10), and Median Rank (MR). As the video and the text are not exactly matched by time, given a video clip, we consider the three closest sentences as correct answers and vice versa. During evaluation, videos of the same movie are removed to avoid noise.

We compare the original UniVL model, the VideoCLIP (Xu et al, 2021) model without finetuning, UniVL finetuned on SYMON data (UniVL-SYMON), as while as UniVL modal finetuned on SYMON with long-term memory (UniVL-SYMON-memory).

7.3 Transfer to YMS

Without in-domain finetuning, we directly test the model trained on SYMON on the YMS dataset, which contains 94 YouTube movie summary videos with manual annotation of fine-grained video-text alignment. To prevent test data leak, we remove any summary videos for the 94 YMS movies from the training set used in this experiment.

Evaluation. In YMS, a text segment may correspond to multiple video clips, whereas a video clip may correspond to one or zero text segments. During inference, we align every video clip to the text segment with the highest similarity, as computed by the neural network. This creates the desired many-to-one alignment. If the highest similarity falls below a threshold, tuned on the validation set, the video clip is considered as not matching anything.

Following Dogan et al (2018), we use two evaluation metrics. Clip accuracy is defined as the temporal proportion of correctly aligned video segments. Sentence IoU is defined as the intersection-over-union between the aligned video durations and the ground-truth durations.

Baselines. Using the network described in 7.1, we compare the original UniVL model, the VideoCLIP (Xu et al, 2021) model without finetuning, UniVL finetuned on SYMON data with and without long-term memory (UniVL-SYMON/ UniVL-SYMON-memory), as well as the supervised NeuMATCH network without the sequential context (*i.e.*, the minimum distance (MD) baseline from Dogan et al (2018)). Details on baseline implementation can be found in the Appendix. Note that UniVL-SYMON and UniVL-SYMON-memory are trained with two video scenes as the basic unit for retrieval and NeuMATCH-MD uses more finely

Table 8 Retrieval performance on SYMON measured using recall@1, 5, and 10 and median rank.

Model	R@1 (↑)	R@5 (↑)	R@10 (↑)	MR (↓)
<i>Text-to-video Retrieval</i>				
UniVL	0.11	0.39	0.63	4818
VideoCLIP	0.02	0.12	0.29	3740.5
UniVL-SYMON	0.76	2.30	3.58	1268
UniVL-SYMON-memory	0.96	3.27	5.41	793
<i>Video-to-text Retrieval</i>				
UniVL	0.03	0.11	0.19	5687
VideoCLIP	0.03	0.16	0.35	4271.5
UniVL-SYMON	0.68	2.36	3.83	1304
UniVL-SYMON-memory	1.21	3.92	6.12	755

segmented units. As YMS contains fine-grained annotations, NeuMATCH-MD is trained on in-domain data, and UniVL-SYMON and UniVL-SYMON-memory are under the transfer learning regime. Thus, this comparison likely puts our network at a disadvantage.

Test Data. For fair comparison with NeuMATCH-MD, we use the original test set of 15 videos and the original video segmentation by Dogan et al (2018), which was performed by a heuristic. In addition, we also create a new split using 70% of the entire YMS as the test set and 30% as the validation set. In this new setting, we also use the new video segmentation strategy in our preprocessing (Section 3.2). The YMS dataset contains sentence-level annotation and sub-sentence annotations. We report performance on both.

7.4 Results

Retrieval on SYMON The bi-directional retrieval results are shown in Table 8. As we expect, the UniVL network finetuned on SYMON (UniVL-SYMON) outperforms both the original UniVL and the zero-shot VideoCLIP baseline. Furthermore, the additional of long-term context yielded significant performance gain, with average R@10 improving by 55%.

Transfer to YMS Table 9 shows the results of video-text alignment on YMS. Despite the difference in segmentation and the weak supervision from SYMON, UniVL-SYMON outperforms the supervised NeuMATCH-MD baseline. This shows that UniVL-SYMON learns a superior cross-modality distance metric, demonstrating the utility of the large-scale SYMON dataset. With sub-sentence level annotations, UniVL-SYMON outperforms the original UniVL by 2.6% / 1.7% in the original setting and 2.7% / 0.9% in the new setting. With sentence level annotations, UniVL-SYMON outperforms the original UniVL by

2.8% / 2.6% in the original setting and 2.0% / 2.0% in the new setting. Furthermore, UniVL-SYMON outperforms zeroshot VideoCLIP and UniVL across all settings. Considering UniVL and VideoCLIP were trained on the gigantic HowTo100M dataset, we attribute the improvement to the similarity between SYMON and YMS, which highlights the effectiveness of SYMON in the domain of story video understanding.

Furthermore, long-term context improved average clip accuracy by 18%. We concede that adding memory causes a performance drop in the original split with sentence-level annotation. However, in the original YMS split, only 15 videos are in the test set. Therefore, we believe the results on our new split to be more trustworthy.

Discussion Overall, utilizing long-term memory yielded improvements on both tasks suggesting that the ability to model long-range dependencies is beneficial for understanding videos in SYMON. As shown in Section. 6.1, videos in SYMON frequently reference past events at long range. Intuitively, modeling long-range dependencies enable the model to better understand story events. Furthermore, this experiment demonstrates that better story understanding leads to better performance on the retrieval task.

8 Conclusion

In this work, we present a multimodal story dataset, SYMON. We quantitatively characterize the dataset and establish the status of SYMON as the largest and most prototypical multimodal story dataset. Furthermore, we establish multimodal retrieval baselines for SYMON and a zero-shot alignment baseline on YMS to demonstrate the effectiveness of the in-domain SYMON in story understanding. We believe that SYMON will serve as a new challenge for the research community and inspire new advances in multimodal learning. The dataset will be released upon acceptance.

9 Limitations

We identify several limitations. First, the preprocessing of the SYMON dataset relies on several automatic techniques. For example, the text description may be written by the original authors but may also come from Google automatic speech recognition. The punctuation restoration model relies on the text. This pipeline architecture may cause errors to propagate. However, given the amount of data, their complexity, and available

Table 9 Alignment performance on YMS. For UniVL-SYMON, we report the average and standard deviation from 5 different random seeds.

	Clip Acc.	Sent. IoU
<i>Original Split (sub-sentence level)</i>		
UniVL	3.3	1.0
VideoCLIP	4.8	0.6
NeuMATCH-MD (Supervised)	4.0	2.4
UniVL-SYMON	5.9 ± 0.3	2.7 ± 0.2
UniVL-SYMON-memory	6.5 ± 0.3	2.6 ± 0.2
<i>New Split (sub-sentence level)</i>		
UniVL	7.4	1.0
VideoCLIP	7.6	0.7
UniVL-SYMON	10.1 ± 0.4	1.9 ± 0.1
UniVL-SYMON-memory	13.5 ± 0.3	2.6 ± 0.1
<i>Original Split (sentence level)</i>		
UniVL	4.6	0.8
VideoCLIP	4.0	1.1
UniVL-SYMON	7.4 ± 0.1	3.4 ± 0.2
UniVL-SYMON-memory	7.5 ± 0.4	2.1 ± 0.2
<i>New Split (sentence level)</i>		
UniVL	5.7	1.3
VideoCLIP	4.9	1.0
UniVL-SYMON	7.7 ± 0.2	3.3 ± 0.2
UniVL-SYMON-memory	8.7 ± 0.3	3.2 ± 0.2

research budget, we find it unrealistic to manually annotate everything. Fortunately, automatic tools improve over time and the preprocessing can be performed again as better tools become available. We opt to make the dataset available now rather than waiting until perfection.

Second, the human annotators in the video sequencing task do not see the entire story video when making decisions. We choose this setting because the story videos are quite long and omitting them simplifies the annotation task without sacrificing predictive accuracy too much. However, this may have lowered our estimates of human performance.

Third, the English-only composition is a limitation of SYMON. In the next iteration, we plan to expand the dataset to include other languages.

Fourth, the movies and TV-shows featured in SYMON comes predominately from western producers. Therefore the depiction of cultural and social conventions in SYMON may be biased towards a western perspective.

References

- Abbott V, Black JB, Smith EE (1985) The representation of scripts in memory. *Journal of memory and language* 24(2):179–199
- Agrawal H, Chandrasekaran A, Batra D, et al (2016) Sort story: Sorting jumbled images and captions into stories. In: EMNLP, pp 925–931
- Alam T, Khan A, Alam F (2020) Punctuation restoration using transformer models for resource-rich and-poor languages. In: W-NUT 2020, pp 132–142
- Aristotle (circa. 335 BC) *Poetics*
- Baek Y, Lee B, Han D, et al (2019) Character region awareness for text detection. In: CVPR, pp 9365–9374
- Bain M, Nagrani A, Brown A, et al (2020) Condensed movies: Story based retrieval with contextual embeddings. In: ACCV
- Barker RG, Wright HF (1955) *Midwest and its children: The psychological ecology of an American town*. Row, Peterson
- Berndt DJ, Clifford J (1994) Using dynamic time warping to find patterns in time series. In: KDD workshop
- Bird S, Klein E, Loper E (2009) Natural language processing with Python: analyzing text with the natural language toolkit
- Bojanowski P, Grave E, Joulin A, et al (2017) Enriching word vectors with subword information. *ACL Transactions* 5:135–146
- Brahman F, Huang M, Tafjord O, et al (2021) "let your characters tell their story": A dataset for character-centric narrative understanding. In: EMNLP Findings
- Burgess A (2022) *Novel*. In: *Encyclopedia Britannica*
- Caba Heilbron F, Escorcia V, Ghanem B, et al (2015) Activitynet: A large-scale video benchmark for human activity understanding. In: CVPR, pp 961–970
- Chambers N, Jurafsky D (2008) Unsupervised learning of narrative event chains. In: *ACL*
- Chambers N, Cassidy T, McDowell B, et al (2014) Dense event ordering with a multi-pass architecture. *ACL Transactions* 2:273–284
- Chaturvedi S, Srivastava S, au2 HDI, et al (2016) Modeling dynamic relationships between characters in literary novels. In: *AAAI*
- Chen M, Chu Z, Wiseman S, et al (2021) Summscreen: A dataset for abstractive screenplay summarization
- Devlin J, Chang MW, Lee K, et al (2019) BERT: Pre-training of deep bidirectional transformers for language understanding. In: *NAACL*, pp 4171–4186
- Dogan P, Li B, Sigal L, et al (2018) LSTM stack-based neural multi-sequence alignment technique (neumatch). [arXiv:1803.00057](https://arxiv.org/abs/1803.00057)
- Emelin D, Bras RL, Hwang JD, et al (2020) Moral stories: Situated reasoning about norms, intents, actions, and their consequences. In: *EMNLP*
- Ferraro F, Van Durme B (2016) A unified bayesian model of scripts, frames and language. In: *AAAI*
- Forster EM (1985) *Aspects of the Novel*, vol 19
- Freytag G (1863) *Die Technik des Dramas*
- Gangal V, Feng SY, Alikhani M, et al (2021) Nareor: The narrative reordering problem. In: *AAAI*
- Genette G (2014) *Nouveau discours du récit*
- Girshick R (2015) Fast r-cnn. In: *ICCV*, pp 1440–1448
- Gorinski PJ, Lapata M (2015) Movie script summarization as graph-based scene extraction. In: *NAACL*, pp 1066–1076
- Graesser AC, Lang KL, Roberts RM (1991) Question answering in the context of stories. *Journal of Experimental Psychology: General* 120:254–277
- Gratch J, Marsella S (2004) A domain-independent framework for modeling emotion. *Cognitive Systems Research* 5(4):269–306

- Gutmann M, Hyvärinen A (2010) Noise-contrastive estimation: A new estimation principle for unnormalized statistical models. In: AISTATS, pp 297–304
- Hara K, Kataoka H, Satoh Y (2018) Can spatiotemporal 3d cnns retrace the history of 2d cnns and imagenet? In: CVPR, pp 6546–6555
- Huang Q, Xiong Y, Rao A, et al (2020) Movienet: A holistic dataset for movie understanding. In: ECCV, pp 709–727
- Iyyer M, Manjunatha V, Guha A, et al (2017) The amazing mysteries of the gutter: Drawing inferences between panels in comic book narratives. In: IEEE Conference on Computer Vision and Pattern Recognition
- Kay W, Carreira J, Simonyan K, et al (2017) The kinetics human action video dataset. [arXiv:1705.06950](https://arxiv.org/abs/1705.06950)
- Kim E, Klinger R (2019) Frowning frodo, wincing leia, and a seriously great friendship: Learning to classify emotional relationships of fictional characters. In: NAACL
- Kim KM, Heo MO, Choi SH, et al (2017) Deepstory: Video story qa by deep embedded memory networks. In: IJCAI, pp 2016–2022
- Kočiský T, Schwarz J, Blunsom P, et al (2018) The NarrativeQA Reading Comprehension Challenge. *ACL Transactions* 6:317–328
- Krishna R, Hata K, Ren F, et al (2017) Dense-captioning events in videos. In: ICCV
- Kuznetsova A, Rom H, Alldrin N, et al (2020) The open images dataset v4. *IJCV* 128(7):1956–1981
- Lei J, Yu L, Bansal M, et al (2018) Tvqa: Localized, compositional video question answering. In: EMNLP
- Li B, Cardier B, Wang T, et al (2018) Annotating high-level structures of short stories and personal anecdotes. In: LREC
- Liu Z, Mitamura T, Hovy E (2018) Graph-based decoding for event sequencing and coreference resolution. In: COLING
- Luo H, Ji L, Shi B, et al (2020) Univl: A unified video and language pre-training model for multimodal understanding and generation. [arXiv:2002.06353](https://arxiv.org/abs/2002.06353)
- Meehan J (1976) *The metanovel: Writing stories by computers*. PhD thesis, Yale University
- Miech A, Zhukov D, Alayrac JB, et al (2019) Howto100m: Learning a text-video embedding by watching hundred million narrated video clips. In: ICCV, pp 2630–2640
- Murakami H (2022) *Where My Characters Come From*
- Naqvi N, Shiv B, Bechara A (2006) The role of emotion in decision making: A cognitive neuroscience perspective. *Current directions in psychological science* 15(5):260–264
- Nie Y, Williams A, Dinan E, et al (2020) Adversarial NLI: A new benchmark for natural language understanding. In: ACL, pp 4885–4901
- Ning Q, Wu H, Han R, et al (2020) Torque: A reading comprehension dataset of temporal ordering questions. In: EMNLP, pp 1158–1172
- O’Gorman T, Wright-Bettner K, Palmer M (2016) Richer event description: Integrating event coreference with temporal, causal and bridging annotation. In: CNS, pp 47–56
- Ouyang J, McKeown K (2015) Modeling reportable events as turning points in narrative. In: EMNLP, pp 2149–2158
- Papalampidi P, Keller F, Frermann L, et al (2020) Screenplay summarization using latent narrative structure. In: ACL, pp 1920–1933
- Papalampidi P, Keller F, Lapata M (2021) Movie summarization via sparse graph construction. In: AAAI, pp 13,631–13,639
- Prince G (2003) *A dictionary of narratology*. University of Nebraska Press
- Rashkin H, Bosselut A, Sap M, et al (2018a) Modeling naive psychology of characters in simple commonsense stories. In: ACL, pp 2289–2299
- Rashkin H, Sap M, Allaway E, et al (2018b) Event2Mind: Commonsense inference on events,

- intents, and reactions. In: ACL
- Reimers N, Gurevych I (2019) Sentence-bert: Sentence embeddings using siamese bert-networks. In: EMNLP, pp 3982–3992
- Richardson M, Burges CJ, Renshaw E (2013) MCTest: A challenge dataset for the open-domain machine comprehension of text. In: EMNLP, pp 193–203
- Riedl MO, Young RM (2010) Narrative planning: Balancing plot and character. *Journal of Artificial Intelligence Research* 39(1):217–268
- Rimmon-Kenan S (1983) *Narrative Fiction: Contemporary Poetics*. Routledge
- Roemmele M, Bejan CA, Gordon AS (2011) Choice of plausible alternatives: An evaluation of commonsense causal reasoning. In: *AAAI Spring Symposium on Logical Formalizations of Commonsense Reasoning*
- Rohrbach A, Torabi A, Rohrbach M, et al (2017) Movie description. *IJCV* 123(1):94–120
- Rosch E, Lloyd BB (1978) *Principles of categorization*
- Sadhu A, Gupta T, Yatskar M, et al (2021) Visual semantic role labeling for video understanding. [arXiv:2104.00990](https://arxiv.org/abs/2104.00990)
- Sap M, LeBras R, Allaway E, et al (2019) Atomic: An atlas of machine commonsense for if-then reasoning. In: *AAAI*
- Schank RC, Abelson RP (1978) *Scripts, plans, goals, and understanding: An inquiry into human knowledge structures*. *Language* 54(3)
- Schank et al. RC (1975) *Sam – a story understander*. Research Report 43, Department of Computer Science, Yale University
- Schmid W (2010) *Narratology: an introduction*
- Siddharthan A, Cherbuin N, Eslinger PJ, et al (2018) Wordnet-feelings: a linguistic categorisation of human feelings. [arXiv:1811.02435](https://arxiv.org/abs/1811.02435)
- Sigurdsson GA, Varol G, Wang X, et al (2016) Hollywood in homes: Crowdsourcing data collection for activity understanding. In: *ECCV*, pp 510–526
- Smith CA, Lazarus RS, et al (1990) *Emotion and adaptation*
- Souček T, Lokoč J (2020) Transnet v2: An effective deep network architecture for fast shot transition detection. [arXiv:2008.04838](https://arxiv.org/abs/2008.04838)
- Sridhar M, Cohn AG, Hogg DC (2010) Relational graph mining for learning events from video. In: *STAIRS*, pp 315–327
- Tapaswi M, Zhu Y, Stiefelwagen R, et al (2016) MovieQA: Understanding Stories in Movies through Question-Answering. In: *CVPR*
- Tomashevsky BV (1925) *Teoriya literatury*. Krug.[Reprint: Letchworth, 1971], *Poetika*, Moscow and Leningrad
- Trabasso T, van den Broek P (1985) Causal thinking and the representation of narrative events. *Journal of Memory and Language* 24(5):612–630
- Trabasso T, Van den Broek P, Suh SY (1989a) Logical necessity and transitivity of causal relations in stories. *Discourse processes* 12(1):1–25
- Trabasso T, van den Broek P, Suh SY (1989b) Logical necessity and transitivity of causal relations in stories. *Discourse Processes* 12(1):1–25
- UzZaman N, Llorens H, Derczynski L, et al (2013) SemEval-2013 task 1: TempEval-3: Evaluating time expressions, events, and temporal relations. In: *SemEval*, pp 1–9
- Vashishtha S, Van Durme B, White AS (2019) Fine-grained temporal relation extraction. In: *ACL*, pp 2906–2919
- Vashishtha S, Poliak A, Lal YK, et al (2020) Temporal reasoning in natural language inference. In: *EMNLP Findings*
- Vicol P, Tapaswi M, Castrejon L, et al (2018) MovieGraphs: Towards understanding human-centric situations from videos. In: *CVPR*
- Wang H, Chen M, Zhang H, et al (2021) Joint constrained learning for event-event relation extraction. In: *EMNLP*

Wilensky R (1983) Planning and Understanding: A Computational Approach to Human Reasoning

Wu CY, Krahenbuhl P (2021) Towards long-form video understanding. In: CVPR, pp 1884–1894

Xie S, Sun C, Huang J, et al (2018) Rethinking spatiotemporal feature learning: Speed-accuracy trade-offs in video classification. In: ECCV, pp 305–321

Xu H, Ghosh G, Huang P, et al (2021) Videoclip: Contrastive pre-training for zero-shot video-text understanding. [arxiv:2109.14084](https://arxiv.org/abs/2109.14084)

Yang Z, Choi JD (2019) FriendsQA: Open-domain question answering on TV show transcripts. In: SIGDIAL, pp 188–197

Young RM, Pollack ME, Moore JD (1994) Decomposition and causality in partial-order planning. In: AIPS, pp 188–194

Zacks JM, Tversky B (2001) Event structure in perception and conception. Psychological bulletin 127(1):3

Zellers R, Lu X, Hessel J, et al (2021) Merlot: Multimodal neural script knowledge models. Advances in Neural Information Processing Systems 34

Zhou B, Richardson K, Ning Q, et al (2021) Temporal reasoning on implicit events from distant supervision. In: NAACL

A Story Coverage

Dynamic Time Warping. We present the formal definition of the DTW problem: given the WikiPlots sequence of sentences $A = (a_1, \dots, a_N)$ and the video narration sentences $B = (b_1, \dots, b_M)$, we seek the best set of correspondence $\{(a_i, b_{g(i)})\}_{i=1}^N$, where the function $g(i) \in \{\epsilon, 1, \dots, M\}$ returns the index in sequence B that matches sentence a_i in A . Further, we introduce a special symbol ϵ . Setting $g(i) = \epsilon$ signifies that a_i is not matched with any sentence in B .

The DTW algorithm can be understood as finding the shortest path in a graph, where each node (i, j) in the graph represents matching sentence a_i and sentence b_j . The graph contains dummy nodes $(0, 0)$ and $(N + 1, M + 1)$. From node (i, j) , we can transit to node $(i + 1, j + 1)$, which would match a_{i+1} with b_{j+1} and

Instructions

Columns **A** and **B** are different narratives of the same story. For each sentence in column **A** try to find an equivalent sentence in column **B** and put its index in the brackets. If there's not equivalent sentence to be found, leave the bracket empty.

A	B
[0]Audrey Burke (Halle Berry) and her warm and loving husband Brian (David Duchovny) have been happily married 11 years.()	[0]Audrey (Halle Berry) and Jerry (Benicio Del Toro) reunite at her husband's funeral.
[1]they have a 10-year-old daughter named Harper (Alexis Llewellyn) and a 6-year-old son named Dory (Micah Berry).()	[1]Audrey (Halle Berry) and Brian (David Duchovny) share their final moment together just hours before Brian is murdered.
.....

Fig. 12 Annotator Instructions for Story Coverage

incur cost $c(i + 1, j + 1)$.

$$c(i + 1, j + 1) = 1 - P(a_{i+1} \Leftrightarrow b_{j+1}). \quad (7)$$

Here $P(a_{i+1} \Leftrightarrow b_{j+1})$ denotes the probability that sentences a_{i+1} and b_{j+1} are equivalent, as determined by the Natural Language Inference classifier.

Similarly, we can transit from (i, j) to $(i + 1, j)$, which would match a_{i+1} with b_j and incur cost $c(i + 1, j)$. The transition from (i, j) to $(i, j + 1)$ is symmetric. Additionally, we can transit from (i, j) to $(i, j + 1, \epsilon)$, which prevents b_{j+1} from matching anything. From $(i, j + 1, \epsilon)$, we may transit to $(i, j + 2, \epsilon)$, $(i, j + 2)$, or $(i + 1, j + 2)$. The costs of ignoring a sentence in A and B are δ_A and δ_B respectively. With this setup, the best correspondence can be found as the path from $(0, 0)$ to $(N + 1, M + 1)$ with minimum cost. We use grid search to find optimal δ_A and δ_B between 0.1 and 1.2 using manually labeled sentence correspondence.

Annotation instructions. Figure 12 shows the instructions we give to our annotator. Here column **A** is the WikiPlot summary and column **B** is the summary from SYMON or CMD or LSMDC.

Justification for the size of labeled validation set. To find the optimal hyperparameter for DTW, we manually label 1000 sentences from CMD and SYMON, and 2000 sentences from LSMDC.

Although the labeled set is small compared to the size of the dataset, we believe it is sufficient for the purpose of the hyperparameter search, as adding more annotated examples does not change the results significantly. We conduct another experiment with 500

Table 10 Estimated story coverage with different validation set sizes. The display is in the form of CMD validation size / LSMDC validation size / SYMON validation size

Validation Set Size	CMD	LSMDC	SYMON
1000 / 2000 / 1000	10.6%	18.1%	44.3%
500 / 1300 / 500	10.8%	18.1%	37.9%

labeled sentences from CMD and SYMON, and 1300 sentences from LSMDC using exactly the same experimental protocol and symmetric entailment for sentence similarity.

Table. 10 shows the results. Increasing the validation set size from 500/1300/500 to 1000/2000/1000 causes little change in the calculated story coverage. Thus, we conclude that the gain from increasing the labeled validation set size has saturated.

B Mental-state Description

Commentary Labeling Protocol. To estimate the amount of commentary text in SYMON and determine if the high frequency of mental-state description is due to narrator commentary, we recruited workers from AMT to label the sentences as either commentary or non-commentary. We organize the sentences into sets of five. The workers were instructed to read the five sentences and identify sentences that they consider commentary. To ensure quality, we only recruited AMT workers with a “Mechanical Turk Master” qualification and assign at least 3 workers to each set of sentences. We consider a sentence as commentary if 2 or more workers agree that it is commentary. As a verification mechanism, we manually labeled a small test set of 100 sentences. The above AMT labeling protocol achieved 100% accuracy on this test set.

Figure 13 shows the annotation instruction given to the AMT workers along with a set of five sentences.

C Estimating the Semantic Gap

First, we introduce the following notations. C_{text} denotes that the human under the video + text condition provided the right answer, and C_{video} denotes that the human with video only provided the right answer. Alternatively, $\neg C_{text}$ and $\neg C_{video}$ denote the situations where they did not get the right answer respectively.

We reckon that the human labels may be the results of random guessing and would like to differentiate random guesses from actually knowing the right answer. For this reason, we introduce notations K_{text} and

K_{video} , which denote the humans with and without the texts know the right answer respectively. $\neg K_{text}$ and $\neg K_{video}$ denote that the human in those conditions do not know the right answer, respectively.

We further assume that, even without knowing the answer, the human can guess 50% correctly. That is,

$$\begin{aligned} P(C_{text}|\neg K_{text}) &= 0.5, \\ P(\neg C_{text}|\neg K_{text}) &= 0.5, \\ P(C_{video}|\neg K_{video}) &= 0.5, \\ \text{and } P(\neg C_{video}|\neg K_{video}) &= 0.5. \end{aligned} \quad (8)$$

But knowing the answer, the human always makes the correct choice,

$$\begin{aligned} P(C_{text}|K_{text}) &= 1, \\ P(\neg C_{text}|K_{text}) &= 0, \\ P(C_{video}|K_{video}) &= 1, \\ \text{and } P(\neg C_{video}|K_{video}) &= 0. \end{aligned} \quad (9)$$

Finally, we assume the independence relations embedded in the probabilistic graphic model in Figure 14. In particular, given K_{text} and K_{video} , C_{text} and C_{video} are independent. That is, random guesses are independent, but whether humans know the answers in the two conditions is not independent. Given K_{text} , K_{video} and C_{text} are independent, and vice versa.

As an example, we show the calculation of the semantic gap on SYMON, the other two datasets can be calculated similarly. With the above independence relations, we have

$$\begin{aligned} &P(\neg C_{text} \wedge \neg C_{video} \wedge \neg K_{video} \wedge \neg K_{video}) \\ &= P(\neg C_{text}|\neg K_{text})P(\neg C_{video}|\neg K_{video}) \\ &P(\neg K_{text} \wedge \neg K_{video}). \end{aligned} \quad (10)$$

We know that

$$\begin{aligned} &P(\neg C_{text} \wedge \neg C_{video}) \\ &= P(\neg C_{text} \wedge \neg C_{video} \wedge K_{text} \wedge K_{video}) \\ &\quad + P(\neg C_{text} \wedge \neg C_{video} \wedge \neg K_{text} \wedge K_{video}) \\ &\quad + P(\neg C_{text} \wedge \neg C_{video} \wedge K_{text} \wedge \neg K_{video}) \\ &\quad + P(\neg C_{text} \wedge \neg C_{video} \wedge \neg K_{text} \wedge \neg K_{video}). \end{aligned} \quad (11)$$

Instruction

This is a narration of a movie summary video. Part of the text contains commentary that don't describe the movie content. **These commentaries contain the narrator's own feeling and opinions.** Here is an example of a commentary and non-commentary sentences:

This sentence is a commentary:

This is one of the most intricate space stations i've seen in a film in a while.

This sentence is not a commentary:

Logan is confronted by Gabriella's killer, transition cyborg chief of security, Donald Pierce, who is looking for Laura.

Select the sentence you think is a commentary. **There can be 0 or 1 or multiple commentaries in the text. If there are no commentaries, please select 'None of the above'.** Please note that the commentaries are rare. We estimate they make up 5% of the text.

Fig. 13 Annotator Instructions for Commentary Labeling

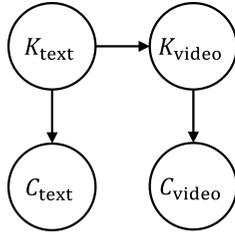


Fig. 14 The directed probabilistic graphical model for analyzing user responses.

We rewrite Eq. 11 with factorizations similar to Eq. 10 and eliminate the zero terms due to Eq. 9,

$$\begin{aligned} & P(\neg C_{text} \wedge \neg C_{video}) \\ &= P(\neg C_{text} | \neg K_{text}) P(\neg C_{video} | \neg K_{video}) \quad (12) \\ & P(\neg K_{text} \wedge \neg K_{video}). \end{aligned}$$

Experimental results show

$$P(\neg C_{text} \wedge \neg C_{video}) = 2.7\%. \quad (13)$$

Therefore,

$$\begin{aligned} & \frac{1}{4} P(\neg K_{text} \wedge \neg K_{video}) = 2.7\%, \quad (14) \\ & P(\neg K_{text} \wedge \neg K_{video}) = 2.7\% \times 4 = 10.8\%. \end{aligned}$$

We know from data that

$$P(\neg C_{video}) = 17.2\%. \quad (15)$$

From our assumptions,

$$\begin{aligned} P(\neg C_{video}) &= P(\neg C_{video} | K_{video}) P(K_{video}) + \\ & P(\neg C_{video} | \neg K_{video}) P(\neg K_{video}) \\ &= 0 + 0.5 P(\neg K_{video}). \end{aligned} \quad (16)$$

Therefore,

$$P(\neg K_{video}) = 2 \times 17.2\% = 34.4\%. \quad (17)$$

Putting things together,

$$\begin{aligned} P(\neg K_{text} | \neg K_{video}) &= \frac{P(\neg K_{text} \wedge \neg K_{video})}{P(\neg K_{video})} \\ &= \frac{10.8\%}{34.4\%} = 31.4\% \end{aligned} \quad (18)$$

$$\begin{aligned} P(K_{text} | \neg K_{video}) &= 1 - P(\neg K_{text} | \neg K_{video}) \\ &= 68.6\% \end{aligned} \quad (19)$$

D Computational Video Temporal Ordering

Preprocessing We generate video segments of exactly 8 seconds separated by a hard camera cut at the midpoint. This is done by selecting a camera cut and cutting the video 4 seconds before and after the camera cut. Each video segment is sampled at 16 frames per second and features are extracted with S3D (Xie et al,

2018) pretrained on HowTo100M. We extract S3D features every second (i.e. from 16 frames), yielding 8 1024-dimensional video features for each video segment. For video features extraction we use a frame size of 112×112 .

We extract the text between the start of the first video segment and the end of the second video segment. To ensure completeness, the text is extended to the nearest sentence boundaries. The maximum number of text tokens is 128. For longer texts, we remove extra tokens from the start and end of the text. For shorter texts, we add zero padding to the end.

Training The models are trained for 30 epochs with learning rate warm-up in the first 6 epochs. Hyperparameters are tuned manually to achieve the highest accuracy on the validation set. The text-aware model is trained with a batch size of 128 and learn rate of 5×10^{-6} and the visual-only model is trained with a batch size of 256 and initial learning rate of 10^{-5} . We apply cosine learning rate decay and the Adam optimizer to all models.

Merlot baseline. For the Merlot baseline, we initialize with weights trained on the YT-Temporal-180M (Zellers et al, 2021) dataset. We finetune the Merlot model on SYMON for 20 epochs with an initial learning rate of 10^{-4} and linear learning rate decay. The finetuning batch size is set to 1024. We use the AdamW optimizer with $\beta_2 = 0.98$. We apply learning rate warm-up for the first 10% of training steps.

Calculating overlap between text description and object / action classes. We first tokenize the text description and use part-of-speech tagging to identify nouns and verbs in the text description (Bird et al, 2009). For matching with object and action detection, we use only the nouns and verbs from the text descriptions and lemmatize them. We also remove common nouns and verbs (“men”, “women”, “person”, and “clothing” for nouns and “is”, “go”, “to”, “get”, “have”, “look”, “walk”, “play”, and “take” for verbs) because these are too frequent and provide little information. For object detection, we adopt the top 10 class predictions for each clip. For action detection, we divide the segment into scenes and take the top 3 class prediction for each scene. Finally, we calculate the number of times the nouns or verbs appear in the detected object or action class names.

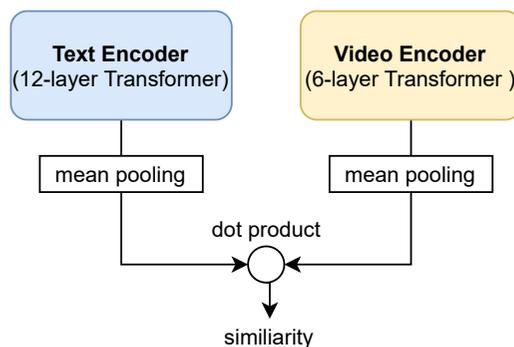


Fig. 15 Retrieval Model

E Retrieval and Alignment

Preprocessing For retrieval, a video segment is extracted every two scenes. For video feature extraction we sample the video at 16 frames per second and use the S3D network (Xie et al, 2018) pretrained on HowTo100M to extract one 1024-dimensional feature every 16 frames. The frame size is 112×112 . We extract 4 features from each clip. If the clip is shorter than 4 seconds, zero paddings are added. If the clip is longer than 4 seconds, we only use the first 4 seconds. We take 64 text tokens for each clip with similar zero padding and truncation. Text is extracted from between the start and the end of the video clip and extended to the nearest sentence boundaries.

Network architecture The video and text encoders consist of 12 and 6 Transformer layers respectively, and are initialized from weights pretrained on HowTo100M. The outputs are then averaged into two 768-dimensional embeddings for video and text. The similarity between a video, text pair is calculated as the dot product of the video and text embeddings.

Training The model is finetuned on SYMON with an initial learning rate of 5×10^{-5} and cosine learning rate decay. We use a batch size of 1024 and train for 20 epochs. The first epoch applies linear warm-up of learning rates. SGD with momentum of 0.9 is used for optimization and a weight decay term of 0.5 is added for regularization. The learning rate, learning rate scheduler, weight decay term and batch size are tuned manually to achieve the maximal R@5 on the validation set. Figure 15 shows the model architecture.

Baselines For the retrieval experiment, we report results from a single training run finetuned on SYMON. For the zero-shot video-text alignment experiment, the

UniVL-SyMoN and UniVL-SyMoN-memory results reported in the main paper are averaged over 5 training runs. The UniVL and VideoCLIP results are from the UniVL / VideoCLIP models without finetuning. The result of the NeuMatch baseline is directly taken from Dogan et al (2018).

For the VideoCLIP baseline, we initialize the VideoClip (Xu et al, 2021) model with weights pre-trained on HowTo100M (Miech et al, 2019). VideoCLIP uses S3D network pretrained on 30fps video data as a feature extractor. Therefore during evaluation, the video is sampled at 30 fps. Similar to UniVL, VideoCLIP is used to calculate the similarity between video and text. For retrieval, the similarity is used to retrieve video from text and vice versa. For alignment, a similarity value is computed for each video-text pair in the full-length video, and each video clip is aligned to the text with the highest similarity.

Long term memory For any given video segment, we find the most similar video segment from 10 preceding video segments using Eq. 6 and append its video features and text tokens to the current video features and text tokens respectively. As a result, the input to the video encoder is 8×1024 vector. We limit the input to the text encoder to 128 tokens.

The model training is the same as training without memory, except that 60% of the video and text tokens are masked during training as a form of data augmentation.

F Example Results and Videos

In Figure 16 we show some examples from SYMoN. Figure 17 shows some example retrieval results from the UniVL and the UniVL+SyMoN models. The first three lines showcases where training on SYMoN improves retrieval performance. The last line show a case where training on SYMoN do not improve performance, possibly due to text that are not visually grounded, such as 'her son says, don't go'.

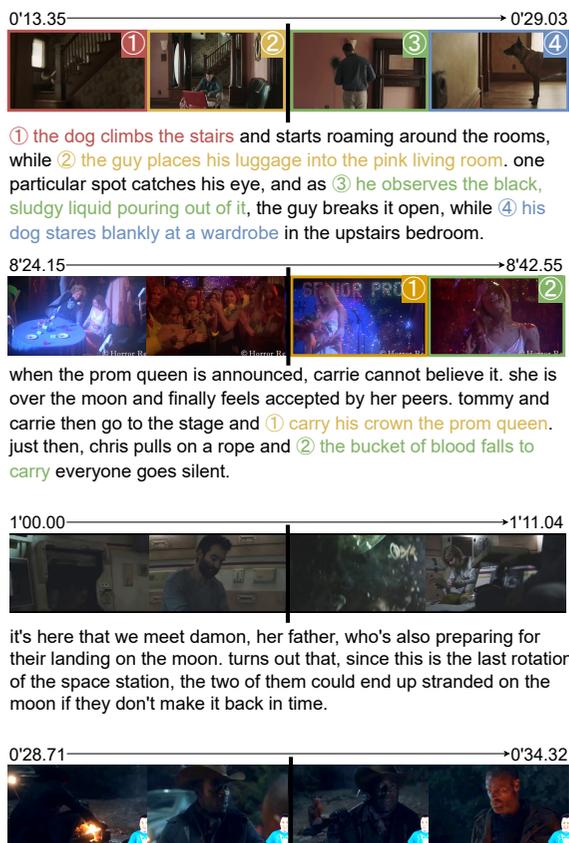


Fig. 16 Examples of video segments and text descriptions from SYMoN. From the top to the bottom, the videos are summaries for the movies *Girl on the Third Floor*, *Carrie*, *Prospect*, and the TV show *Fear the Walking Dead*. The frames are cropped to the scale of 9:16 for uniform display. We increase brightness for line 4 for clear display. Bound boxes of the same color in text and video frame denote video-text correspondence. The black line denotes the boundary between the two video segments to be ordered.

Text	Groud Truth	Top 5 Retrieved videos
<p>learning of the recent increasing malfunction rate of robots, the supervisor raises his concerns, believes that the park should take actions to lower the malfunction rate before it's losing control. after urgent maintenance of the damaged robots, the three worlds are recovered to business the next day.</p>		<p data-bbox="605 364 962 388">Retrieval results from UniVL-SyMoN</p> 
<p>letty then leaves with a motorcycle, and dom reviews the video of mr nobody, spotting a familiar cross symbol that he recognizes from a postcard he received of mexico from a while back. in the morning, when the team is about to board the plane, dom rushes to join his teammates.</p>		<p data-bbox="605 668 962 691">Retrieval results from UniVL-SyMoN</p>  <p data-bbox="605 791 879 815">Retrieval results from UniVL</p> 
<p>things come to a head when Brian and Dom seem to be outnumbered on a bridge. Dom turns around to face them so Brian can escape and be a father. but Brian isn't going to abandon his family like that</p>		<p data-bbox="605 978 962 1001">Retrieval results from UniVL-SyMoN</p>  <p data-bbox="605 1101 879 1125">Retrieval results from UniVL</p> 
<p>anna again, climbing up with her husband, but this time, her son says, sadly, mother, don't go. and here she lost her patience.</p>		<p data-bbox="605 1285 962 1309">Retrieval results from UniVL-SyMoN</p>  <p data-bbox="605 1409 879 1432">Retrieval results from UniVL</p> 

Fig. 17 Example of text to video retrieval results with UniVL model and UniVL finetuned on SYMON. The frames are cropped to a constant scale of 9:16, and the brightness of the frames are adjusted for clear display.