

HiCM²: Hierarchical Compact Memory Modeling for Dense Video Captioning

Minkuk Kim¹, Hyeon Bae Kim¹, Jinyoung Moon², Jinwoo Choi^{1*}, Seong Tae Kim^{1*}

¹Kyung Hee University, Republic of Korea

²Electronics and Telecommunications Research Institute (ETRI), Republic of Korea
 {asdjklfgh97,hyeonbae.kim}@khu.ac.kr, jymoon@etri.re.kr, {jinwoochoi,st.kim}@khu.ac.kr

Abstract

With the growing demand for solutions to real-world video challenges, interest in dense video captioning (DVC) has been on the rise. DVC involves the automatic captioning and localization of untrimmed videos. Several studies highlight the challenges of DVC and introduce improved methods utilizing prior knowledge, such as pre-training and external memory. In this research, we propose a model that leverages the prior knowledge of human-oriented hierarchical compact memory inspired by human memory hierarchy and cognition. To mimic human-like memory recall, we construct a hierarchical memory and a hierarchical memory reading module. We build an efficient hierarchical compact memory by employing clustering of memory events and summarization using large language models. Comparative experiments demonstrate that this hierarchical memory recall process improves the performance of DVC by achieving state-of-the-art performance on YouCook2 and ViTT datasets.

1 Introduction

Dense video captioning (DVC) is a task that involves temporally localizing event boundaries in untrimmed videos and describing each event in natural language. Compared to video captioning (Lin et al. 2022; Luo et al. 2020; Seo et al. 2022; Gao et al. 2017; Pei et al. 2019; Qi et al. 2019; Wang et al. 2018; Chen et al. 2017; Pan et al. 2016; Rohrbach et al. 2013; Venugopalan et al. 2014, 2015; Tu et al. 2021, 2022, 2023), which generates a single caption for a video clip, DVC presents considerable challenges. This task requires not only generating captions but also automatically localizing multiple important event segments within untrimmed videos. In other words, DVC should accurately identify event boundaries and describe each event segment in natural language, making it a far more complex process.

Recent advancements in vision and language learning have demonstrated remarkable outcomes in cross-modal correlation tasks (Radford et al. 2021; Saharia et al. 2022; Li et al. 2022). Nevertheless, linking natural language with video remains challenging, largely due to the complexities involved in capturing spatiotemporal information (Hayes

*Corresponding Authors.

Copyright © 2025, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

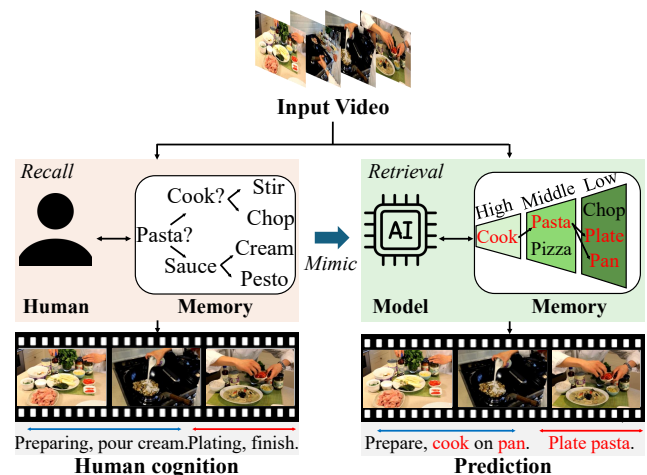


Figure 1: Conceptual figure of the proposed hierarchical compact memory construction. Our method leverages relevant clues from hierarchical compact memory with cross-modal retrieval, effectively mimicking human memory processes for improved event localization and description.

et al. 2022). Video and language learning demands intricate model architectures, tailored training protocols, and significant computational resources (Cheng et al. 2023). DVC involves the additional challenge of associating untrimmed videos with natural language to both localize and describe events (Krishna et al. 2017; Wang et al. 2021; Yang et al. 2023b). To address these challenges, several studies have been explored to leverage prior knowledge through a pre-training and fine-tuning approach (Yang et al. 2023b; Wu et al. 2024). In addition, a study has been proposed (Kim et al. 2024) that enhances DVC by utilizing information stored in memory as an additional source of prior knowledge. This study is inspired by the observation that humans have the ability to identify significant events and describe them by retrieving relevant memories based on cues they have encountered. In cognitive information processing, this phenomenon is called cued recall (Allan and Rugg 1997; Rugg et al. 1998). Notably, when recalling, humans tend to first perceive abstract semantic information, then associate

it with related categories in memory, and ultimately retrieve specific episodes connected to the recognized information, forming a memory hierarchy (Tulving et al. 1972). By recalling memories and drawing on prior knowledge, humans can describe a given situation based on a rich pool of experiences.

Building on this insight, we introduce a novel DVC model, HiCM² (Hierarchical Compact Memory Modeling for DVC). Our model effectively recalls relevant episodes and abstract concepts from a hierarchical and compact external memory, enhancing the quality of captions, as in Figure 1. To mimic the human recalling process, we design an external memory based on knowledge information extracted from the training data. This knowledge information is organized into a hierarchical memory through iterative clustering using a bottom-up approach. During this process, knowledge information at each level is summarized using large language models (LLMs), transforming it into consolidated memory knowledge. This approach combines hierarchical memory and efficient compact memory construction, resembling human memory organization.

Subsequently, the proposed model extracts potential event candidates from untrimmed videos and retrieves relevant information from the external memory to provide diverse and semantically rich data. Unlike the bottom-up memory construction, information retrieval follows a top-down approach, starting from high-level abstract information and moving towards lower-level details. We first select relevant information from the higher levels and then iteratively refine the search by focusing on related memories from the lower levels, thus facilitating an efficient hierarchical memory read.

Our main contributions can be summarized as:

- Inspired by the human cognitive process, we introduce a novel DVC method that utilizes cross-modal retrieval from the hierarchical compact memory bank. This approach combines hierarchical organization with compact representations, enabling effective recall of relevant episodes and abstract concepts while minimizing redundancy. To the best of our knowledge, this is the first study to design hierarchical memory structures for DVC.
- To efficiently extract and utilize hierarchical information, we propose a novel hierarchical memory retrieval method. Our approach leverages a top-down retrieval strategy, starting from high-level abstract information and progressively accessing detailed lower-level data. This method ensures that meaningful and relevant information is effectively retrieved from different levels of the hierarchical memory.
- To demonstrate the effectiveness of our hierarchical memory retrieval method, we have conducted comprehensive experiments on YouCook2 and ViTT datasets. Our model achieves state-of-the-art on two DVC datasets. Experimental results show that our method could significantly improve the performance by exploiting relevant information from hierarchical memory.

2 Related Work

2.1 Dense Video Captioning

DVC tasks comprise localizing events within untrimmed videos and generating descriptive captions for each localized event (Krishna et al. 2017). Traditionally, a two-stage approach has been employed, where events in the video are first localized, and then captions are generated for each event (Iashin and Rahtu 2020a,b). This "localize-then-describe" strategy addresses the two sub-tasks independently. However, this two-stage approach has a limitation in that it fails to consider the interaction between the event localization and event captioning tasks, as they are performed independently. To address this issue, recent research has introduced end-to-end approaches that enable joint learning by considering the interaction between the two tasks. PDVC (Wang et al. 2021) defines DVC as a set prediction task that predicts event intervals and captions simultaneously, using the DETR (Carion et al. 2020) architecture to jointly perform event localization and captioning in parallel. Vid2Seq (Yang et al. 2023b), on the other hand, defines DVC as a sequence-to-sequence problem, adopting an end-to-end approach that outputs a single sequence containing both timestamps and captions. This research also incorporates speech-to-text technology, using transcribed audio information as input to leverage multimodal data. (Zhou et al. 2024) proposes a visual memory and streaming decoding approach. DIBS (Wu et al. 2024) introduces a new pretraining framework. CM² (Kim et al. 2024) applies retrieval-augmented generation (RAG), inspired by how humans process visual information, retrieving relevant information from an external memory pre-constructed with visual data to provide additional context to the model. However, when constructing the external memory, all captions from the target benchmark dataset are included, leading to the potential inclusion of redundant and noisy information. To address this, our research aims to construct an external memory that filters out redundant or irrelevant information by utilizing clustering and LLMs.

2.2 Retrieval-Augmented Generation

Recently, RAG has gained significant popularity for its effective application in LLMs, where it references external knowledge to address issues such as hallucinations (Lewis et al. 2020). RAG is particularly effective at reducing the generation of factually incorrect content, making it a valuable approach in various visual captioning tasks (Gao et al. 2023). This approach involves retrieving relevant information from an external memory bank based on visual data and providing it to the model (Ramos et al. 2023; Kim et al. 2024). However, the performance of RAG depends on how constructed the memory and it has not been explored to efficiently construct memory banks for effective DVC. In this study, we designed a system to effectively retrieve semantic information, positively impacting model performance. To achieve this, we employed a clustering and compact representation method to reduce redundant information, in contrast to existing approaches that utilize all caption data in an external memory bank. Additionally, in contrast to tra-

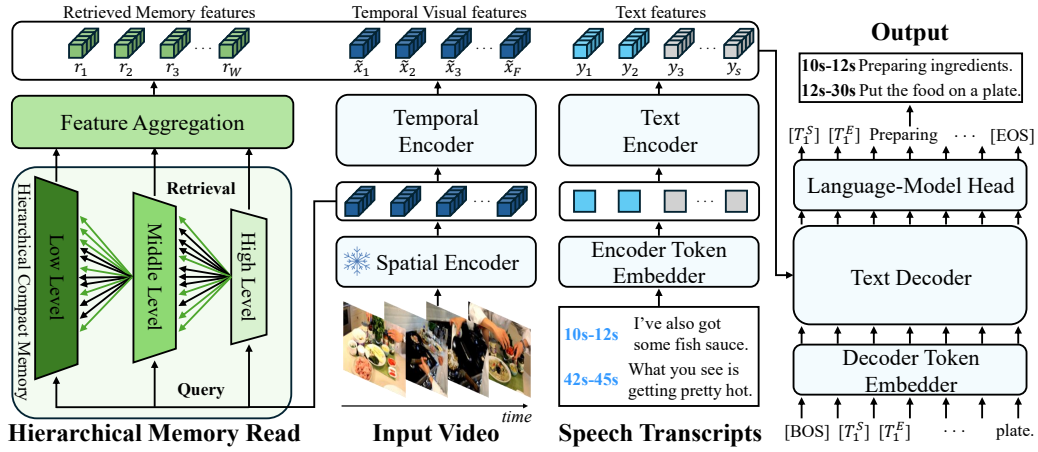


Figure 2: **Overview of HiCM²**. We show the overall architecture. HiCM² approaches the DVC task in a memory retrieval-augmented generation manner using hierarchical compact memory. We conduct video-to-text cross-modal retrieval at each hierarchical level using input video features obtained through a pre-trained spatial encoder. The visual features are temporally encoded through a temporal encoder. The time information of speech is converted into time tokens for time tokenization, and the speech text is tokenized for speech encoding. Each encoded feature vector is concatenated and fed into the cross-attention layer of the text decoder. Finally, we obtain a sequence consisting of the start time, end time, and caption.

ditional single-layer memory structures that read information based on similarity, we utilized a hierarchical clustering method to construct a hierarchical memory structure. This allows for the retrieval of information in a stepwise manner, starting from abstract to more detailed information.

3 Method

The goal of this study is to enhance event localization and caption generation in untrimmed videos by leveraging prior knowledge. We introduce HiCM², a novel framework that utilizes cross-modal retrieval from a hierarchical compact memory. To achieve this, we construct the memory using a hierarchical organization combined with compact representations, enabling the effective recall of relevant episodes and abstract concepts while minimizing redundancy (Section 3.1). Additionally, we propose a method for effective hierarchical memory retrieval. By employing a top-down approach, we begin with high-level abstract information and gradually access finer details from lower levels. This ensures the effective extraction of meaningful information from the hierarchical memory, improving both event localization and caption generation in DVC (Section 3.2).

Figure 2 shows the overall architecture of our method. For the given input video, the model generates segment and caption pairs $\{(T_n^S, T_n^E, S_n)\}_{n=1}^N$ where N denotes the number of events detected by our method and T_n^S and T_n^E denote the start and the end timestamp of n -th event. S_n denotes the generated captions for n -th event segment. Details of dense event prediction will be introduced in Section 3.3.

3.1 Hierarchical Compact Memory Modeling

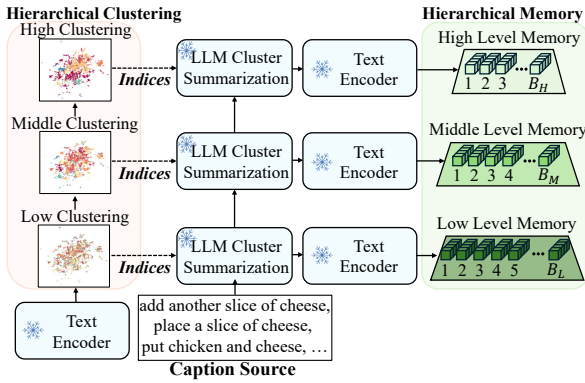
In this section, we describe the method for constructing a hierarchical memory and generating compact representations. To recall relevant episodes and abstract concepts, we

build a hierarchical memory using iterative clustering in a bottom-up approach. To minimize redundancy while creating a memory-efficient compact structure, we generate compact representations through LLM-based summarization.


To ensure the memory bank stores high-quality semantic information, we use captions from the training dataset. For segment-level video-to-text retrieval, we utilize the pre-trained CLIP ViT-L/14 model (Dosovitskiy et al. 2020; Radford et al. 2021), which effectively aligns images and text within a shared embedding space.

Building Hierarchical Memory. To construct a hierarchical memory structure, we leverage the hierarchical clustering method FINCH (Sarfranz, Sharma, and Stiefelwagen 2019), which performs iterative clustering using a bottom-up approach. FINCH forms clusters based on the observation that the first neighbor of each data point is sufficient to discover a connected chain among data points. Furthermore, as a hierarchical clustering method, FINCH generates a meaningful clustering tree that interprets data at different levels of granularity, making it an appropriate approach for constructing hierarchical memory. Note that the number of memory levels is automatically determined based on the FINCH clustering results for each dataset.

As illustrated in Figure 3 (a), we first encode the caption sources using the CLIP text encoder. Then, we perform iterative clustering on the encoded sentences until only abstract clusters that can no longer be grouped remain, resulting in a hierarchical clustering structure. Since each level’s clusters are formed by the connections between clusters from the lower level, we obtain clustering indices for each level. These indices are then used to construct the memory components at each hierarchical level. With this, we construct hierarchical memory that contains high-level abstract information and finer details from lower levels.



(a) Hierarchical Memory Construction



(b) Hierarchical Memory Read

Figure 3: The construction and read process of hierarchical compact memory. As illustrated in (a), hierarchical memory is constructed using iterative clustering in a bottom-up approach to recall relevant episodes and abstract concepts, with LLM-based summarization generating compact, memory-efficient representations. As illustrated in (b), We retrieve a segmented visual cue from each temporal anchor, beginning at the high-level, abstract layer, and then recursively exploring lower layers for additional relevant information. K features are retrieved per level, and this process is repeated for all temporal anchors.

Compact Representation with LLM. To minimize redundancy and effectively recall information from memory, we store compact representations generated by summarization using an LLM. For this purpose, we adopt the pre-trained Llama3 70B model (Dubey et al. 2024). We use the LLM to store representative information derived from clusters. We provide the LLM with instructions for summarizing clustered sentences from each hierarchical level. This process results in the creation of a summarized sentence-level feature for each cluster, effectively removing outliers and capturing core semantic information within the cluster. Additionally, through summarization, we reduce approximately 10K memory components to around 2K, constructing a memory-efficient compact structure.

As in Figure 3 (a), we first summarize the sentences of clusters from the caption source using clustered indices for the lower level. From this, we obtain one summarized sentence per cluster. Consequently, at each memory level, an encoded memory component is created for each cluster, thereby obtaining compact representations. Then, the summarized sentences from this level serve as input for the next level’s summarization, where they are further summarized according to the next level’s cluster indices. We repeat this process across all levels of the memory, constructing a compact memory structure that facilitates efficient recall.

3.2 Retrieval from Hierarchical Memory

The hierarchical compact memory is structured with a high level containing abstract information and lower levels that provide more detailed information. Consequently, to effectively utilize this well-composed hierarchical compact memory, it is important to systematically search through the levels, moving from abstract to detailed, to retrieve and utilize information in an appropriate manner. For this, we propose a retrieval method that begins by searching for information related to the visual cue within the most abstract, high-level

layer of the hierarchical memory bank and then recursively explores lower layers for additional relevant information.

As illustrated in Figure 3 (b), we divide the input video into W temporal anchors to extract semantic information related to events, as untrimmed videos may contain multiple events. We use CLIP ViT-L/14 for frame-level visual feature extraction and compress these features by averaging within each anchor, resulting in segment-level visual features. These features then serve as queries to retrieve relevant information from the external memory for each anchor. We then calculate the similarity between the segment-level visual cues and the memory features in the high-level layer, selecting the most similar features. Subsequently, the corresponding clusters containing these selected features are tracked down to the lower layers, where we calculate the similarity between the visual cue and the more detailed clusters that result from the merging process. This recursive process continues across the lower layers to obtain increasingly specific information. For the relevant features obtained, we perform feature aggregation through two rounds of average pooling. First, we conduct average pooling on the most similar features selected at each level to create a single representative feature. Then, we take one feature from each level and perform average pooling again to integrate the information from all levels. This entire process is repeated for each of the W temporal anchors. Finally, we obtain the retrieved text features $\mathbf{r} = \{r_a\}_{a=1}^W$.

3.3 Dense Event Prediction

Architecture. This section describes a framework of structure that integrates visual features with retrieved text features to enhance event localization and captioning.

As illustrated in Figure 2, the visual encoder processes a sequence of F frames. The spatial encoder outputs embeddings for each frame, which are then fed into a temporal encoder composed of transformer blocks to account for tem-

poral interactions between frames. This process yields the final temporal visual features $\tilde{\mathbf{x}} = \{\tilde{x}_i\}_{i=1}^F$ as the output of the visual encoder. We use CLIP ViT-L/14 as the backbone to extract spatial embeddings from the input frames, keeping it in a frozen state for ease of use.

For memory retrieval, the spatial embeddings extracted by the visual encoder are used to generate segment-level temporal anchors. The embeddings for each frame are divided into W temporal anchors, which serve as queries to the hierarchical memory bank. Each anchor retrieves relevant hierarchical information from the memory based on similarity. Then, we get $\mathbf{r} = \{r_a\}_{a=1}^W$. We also use speech information. For the text of speech transcript, the encoder’s token embedder independently embeds each token, producing semantic embeddings. For the given time information of speech, we begin with a text tokenizer that has a vocabulary size of v , to which we add b time tokens, thus expanding the tokenizer to encompass $v + b$ tokens. These time tokens denote relative timestamps within a video, which we achieve by dividing a video into b evenly spaced timestamps. Specifically, we employ the SentencePiece tokenizer (Kudo and Richardson 2018), which features a vocabulary size of $v = 32,128$ and includes $b = 100$ time tokens. Subsequently, the transformer-based text encoder calculates the interactions between the tokens of the speech transcripts, outputting text features $\mathbf{y} = \{y_j\}_{j=1}^s$. The text decoder uses the retrieved memory features, temporal visual features, and text features that compose the encoder embeddings to autoregressively generate a target event sequence consisting of segment and caption pairs.

Training and Inference. During training, we train HiCM² by minimizing the following loss function:

$$\mathcal{L}_\theta(x, y, r, z) = - \frac{1}{\sum_{c=1}^{L-1} w_c} \sum_{c=1}^{L-1} w_c \log p_\theta(z_{c+1} | x, y, r, z_{1:c}), \quad (1)$$

where L represents the length of the target sequence in the decoder, w_c indicates the weight assigned to the c -th token in the sequence, uniformly set to $w_c = 1 \forall c$ in practice, θ symbolizes the trainable parameters within the model, and p_θ is the probability distribution generated by the model over the vocabulary of text and time tokens. Our approach, based on the sequence-to-sequence structure of the Vid2Seq model (Yang et al. 2023b), fine-tunes a model pre-trained on approximately 1.8 million videos (Yang et al. 2023a). During inference, given visual features x , given text features y , and retrieved text input r , our model predicts a target event sequence z . For both training and inference, we conducted retrieval using the same hierarchical compact memory bank.

4 Experiments

To evaluate the effectiveness of our method, we conducted comparative experiments. Section 4.1 describes the experimental setup used in this study. Section 4.2 highlights the role of memory retrieval in DVC. Section 4.3 presents a comparison with state-of-the-art methods. Section 4.4 provides ablation studies to validate the contribution of each component in our model. We also provide qualitative results.

Memory Construction	Hierarachy	CIDEr	METEOR	SODA _c	F1
No Memory	✗	66.29	12.41	9.87	31.08
All Training Captions	✗	67.90	12.49	10.38	32.31
Clustering	✓	67.15	12.97	10.17	32.30
Clustering+LLM(Ours)	✓	71.84	12.80	10.73	32.51

Table 1: Effect of hierarchical compact memory retrieval on YouCook2.

4.1 Experimental Settings

Dataset. We utilized two DVC benchmark datasets, YouCook2 (Zhou, Xu, and Corso 2018) and ViTT (Huang et al. 2020), for both training and evaluation. The YouCook2 dataset includes 2,000 untrimmed videos depicting cooking procedures, with an average duration of 320 seconds per video and 7.7 temporally localized sentences per annotation. We employed the standard dataset split for training, validation, and testing purposes. The ViTT dataset includes 8,000 untrimmed instructional videos, each with an average length of 250 seconds and annotated with 7.1 temporally localized short tags. It is worth noting that we utilized approximately 10%~20% fewer videos than the original dataset, as we only included those accessible on YouTube (Yang et al. 2023b).

Evaluation Metrics. We evaluated our method on two sub-tasks in DVC. By using the official evaluation tool (Wang, Zheng, and Yu 2020), we measured the quality of the generated captions with the CIDEr (Vedantam, Lawrence Zitnick, and Parikh 2015), BLEU4 (Papineni et al. 2002), and METEOR (Banerjee and Lavie 2005) metrics, which compare the generated captions to the ground truth across IoU thresholds of 0.3, 0.5, 0.7, 0.9. Additionally, to assess storytelling ability, we used the SODA_c metric (Fujita et al. 2020). For event localization, we calculated the average precision, average recall, and F1 score (the harmonic mean of precision and recall), averaging these metrics over IoU thresholds of 0.3, 0.5, 0.7, 0.9.

Implementation Details. For both datasets, video frames were extracted at a rate of 1 frame per second. The sequences were then either subsampled or padded to ensure a total of F frames, where F was set to 100. Both the text encoder and decoder are initialized using a pretrained T5-Base model (Raffel et al. 2020).

The hierarchical memory is organized into 4 levels for YouCook2 and 5 layers for ViTT. For YouCook2, we use 10,337 sentences to construct compact memory across levels, with {1758, 313, 68, 8} memory units allocated per level. For ViTT, we utilize 34,599 sentences, distributing compact memory across levels as follows: {4487, 741, 114, 12, 3} memory units per level. For YouCook2, the number of anchors is set to 10 and the number of retrieved features is set to 10 for each anchor level. For ViTT, the number of anchors is set to 30 and the number of retrieved features is set to 30 for each anchor level.

4.2 Effect of Hierarchical Memory Retrieval

To evaluate the effectiveness of memory retrieval with hierarchical compact memory, we conducted comparative exper-

Method	PT	YouCook2(val)				ViTT(test)			
		CIDEr	METEOR	SODA _c	BLEU4	CIDEr	METEOR	SODA _c	BLEU4
PDVC	✗	29.69	5.56	4.92	1.40	-	-	-	-
CM ²	✗	31.66	6.08	5.34	1.63	-	-	-	-
Streaming V2S	✓	32.90	7.10	6.00	-	25.2	5.80	10.00	-
DIBS	✓	44.44	7.51	6.39	-	-	-	-	-
Vid2Seq [†]	✓	<u>66.29</u>	<u>12.41</u>	<u>9.87</u>	<u>5.64</u>	<u>48.84</u>	<u>9.51</u>	<u>14.99</u>	<u>0.71</u>
HiCM ² (Ours)	✓	71.84	12.80	10.73	6.11	51.29	9.66	15.07	0.86

Table 2: **Performance of event captioning.** Bold means the highest score. PT denotes pretraining from the additional video datasets. [†] denotes results reproduced from official implementation in our environment. All methods use CLIP backbone.

iments on YouCook2. As in Table 1, we compare four different memory retrieval methods. “No Memory” refers to the test results of an encoder-decoder structure without memory retrieval. “All Training Captions” involves storing all memory captions as features in a flat memory structure without any hierarchical organization. “Clustering” utilizes the FINCH algorithm (Sarfranz, Sharma, and Stiefelhagen 2019) to create a hierarchical structure, using averaged features of clustered sentences as memory elements. In this method, a bottom-up approach is applied, where layers above the lowest level are formed by averaging features from the previous level, based on clustering result indices. “Clustering + LLM” involves generating representative captions for each cluster through summarization with LLM, which are then used as memory elements. This approach, implemented in a bottom-up fashion, represents the hierarchical compact memory.

As shown in Table 1, the results demonstrate that storing and retrieving all training captions in memory leads to a performance improvement compared to using no memory at all. Further enhancements are observed when employing hierarchical clustering to create a structured memory, with additional performance gains achieved by generating prototypes through LLM-based summarization. This suggests that while average pooling after hierarchical clustering provides some benefit, the use of LLMs for summarization is more effective in creating semantic prototypes, resulting in more substantial performance improvements.

4.3 Comparison with State-of-the-art-Methods

In Table 2, we compare our method with state-of-the-art approaches (Wu et al. 2024; Zhou et al. 2024; Yang et al. 2023b) on both the YouCook2 and ViTT datasets. As shown in Table 2, the methods are categorized based on their use of pretraining, with performance measured across several key metrics, including CIDEr, METEOR, SODA_c, and BLEU4. The results indicate that methods employing pretraining generally outperform those that do not, highlighting the importance of pretraining in improving caption quality. Our method achieves the best scores across the four metrics among the models that use pre-trained knowledge. By using prior knowledge from external hierarchical memory, our method improves the quality of caption generation across all metrics on both datasets. This performance advantage underscores the effectiveness of our approach in generating accurate and contextually appropriate captions.

Method	PT	YouCook2(val)			ViTT(test)		
		F1	Recall	Precision	F1	Recall	Precision
PDVC	✗	26.81	22.89	32.37	-	-	-
CM ²	✗	28.43	24.76	<u>33.38</u>	-	-	-
Streaming V2S	✓	24.10	-	-	35.40	-	-
DIBS	✓	<u>31.43</u>	26.24	39.81	-	-	-
Vid2Seq [†]	✓	31.08	<u>30.38</u>	31.81	46.21	45.89	<u>46.53</u>
HiCM ² (Ours)	✓	32.51	32.51	32.51	<u>45.98</u>	<u>45.00</u>	47.00

Table 3: **Performance of event localization.** Bold means the highest score. [†] denotes results reproduced from official implementation in our environment.

High	Middle	Low	CIDEr	METEOR	SODA _c	F1
		✓	66.29	12.41	9.87	31.08
		✓	67.75	12.33	10.35	32.27
✓	✓		67.59	12.45	10.37	32.23
		✓	67.21	12.21	10.28	31.98
	✓	✓	68.61	12.35	10.51	32.07
✓		✓	69.81	12.63	10.54	33.00
✓	✓		68.05	12.33	10.41	32.87
✓	✓	✓	71.84	12.80	10.73	32.51

Table 4: **The ablation study for the use of hierarchical levels memory.** Among the four hierarchical levels generated from the YouCook2 training data, the high level corresponds to level 4, the middle level includes levels 2 and 3, and the low level corresponds to level 1.

In Table 3, we compare the performance of various methods on the event localization task across the YouCook2 and ViTT datasets. The methods are evaluated based on whether they utilize pretraining, with key metrics including F1, Recall, and Precision. Our method, HiCM², demonstrates comparable performance across both datasets. HiCM² effectively leverages hierarchical compact memory, resulting in more accurate event localization. The results highlight the advantages of our approach in capturing complex temporal patterns within untrimmed videos. These findings underscore the robustness and effectiveness of HiCM² in enhancing event localization tasks. The results from both captioning and localization performance indicate that pre-trained prior knowledge and retrieval-augmented prior knowledge complement each other, creating a synergistic effect.

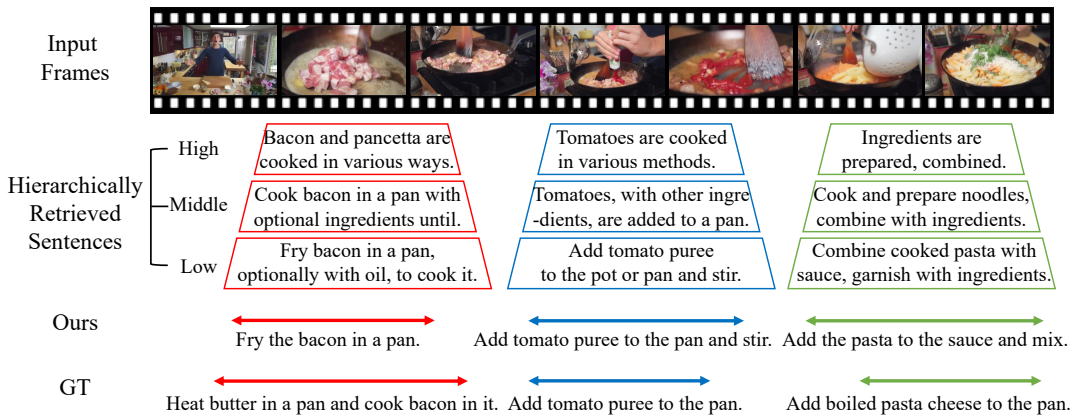


Figure 4: **Example of predictions on the YouCook2 Validation set using our approach.** The hierarchical retrieved sentences shown are examples of retrieval results with the highest semantic similarity at each hierarchical level, corresponding to specific segments of the input frames. Each retrieved sentence is converted into features and utilized in our model’s predictions for the segments. Matching colors indicate the association between retrieved knowledge and prediction.

4.4 Discussion

Effect of Each Level in Hierarchy. Table 4 presents an ablation study evaluating the impact of utilizing hierarchical memory at different levels on performance. The high level is the level that contains the most abstract concepts. The middle level contains more detailed actions and situations, providing a balance between abstract concepts and specific details. The low level focuses on concrete details and events accurately describing specific actions, objects, and events within the scenes. The results clearly demonstrate the importance of incorporating memory across all hierarchical levels, as the best performance is achieved when all three levels (high, middle, and low) are utilized together.

We show the effectiveness of a comprehensive hierarchical memory approach, where each level contributes uniquely to improving the overall performance.

Qualitative Result. Figure 4 shows an example of predictions by our approach, demonstrating effective use of memory retrieval that references meaningful and helpful sentences from memory across different levels, achieved through segment-level video-text retrieval for the given video. As shown in the figure, sentences formed at a high level typically include words that encompass a variety of methods or ingredients. At the middle level, while still abstract about various ingredients, there is a more detailed description of actions and contexts. At the low level, sentences precisely describe the specific subjects and actions appearing in the scene. As a result, our method generates relatively accurate event boundaries and captions with abstract concepts and detailed descriptions.

Effect of Number of Retrieved Features for Each Level.

We also explored the effect of the number of retrieved features per anchor as shown in Table 5. Setting this number to 1 means that only the most similar text from the memory, relative to the visual feature of the temporal anchor, is considered. Increasing this number generally provides the model with more stable and robust semantic information.

#K	CIDEr	METEOR	SODA _c	F1
1	67.47	12.39	10.53	32.27
5	68.15	12.58	10.33	32.51
10	71.84	12.80	10.73	32.51
20	67.97	12.22	10.40	31.56

Table 5: **Effect of the number of selected features in YouCook2.** #K denotes the number of retrieved text features per anchor from the memory bank at each level.

When the number of retrieved features per anchor reaches 10, our approach consistently delivers strong performance across both sub-tasks in the YouCook2 dataset. Note that the number of retrieved features has been uniformly applied across all levels. However, retrieving too many features can introduce noise, as they are selected based on descending similarity, potentially lowering the overall performance.

5 Conclusion

In this study, we introduce a novel DVC model that leverages hierarchical compact memory with cross-modal retrieval, effectively mimicking human memory processes for improved event localization and description. By organizing memory from abstract to detailed information and employing a top-down retrieval process, the model ensures precise and coherent information retrieval. Through comprehensive experiments on the YouCook2 and ViTT datasets, we demonstrate the effectiveness of our hierarchical memory retrieval approach. Notably, HiCM² achieved state-of-the-art performance on both datasets. Our work highlights the potential of memory-augmented models to enhance DVC. Additionally, our findings suggest that the synergy between pre-trained prior knowledge and retrieval-augmented knowledge in our approach could complement existing pretraining efforts in vision-and-language tasks, potentially offering a promising avenue for further improvement in the field.

Acknowledgments

This work was supported in part by the Institute of Information and Communications Technology Planning and Evaluation (IITP) Grant funded by the Korea Government (MSIT) under Grant 2020-0-00004 (Development of Provisional Intelligence Based on Long-term Visual Memory Network), Grant 2022-0-00078, Grant IITP-2024-RS-2023-00258649, Grant RS-2022-00155911, Grant 2021-0-02068, by the National Research Foundation of Korea (NRF) Grant funded by the Korea Government (MSIT) under Grant RS-2024-00334321, and by Center for Applied Research in Artificial Intelligence (CARAI) grant funded by DAPA and ADD (UD230017TD).

References

- Allan, K.; and Rugg, M. 1997. An event-related potential study of explicit memory on tests of cued recall and recognition. *Neuropsychologia*, 35(4): 387–397.
- Banerjee, S.; and Lavie, A. 2005. METEOR: An automatic metric for MT evaluation with improved correlation with human judgments. In *Proceedings of the acl workshop on intrinsic and extrinsic evaluation measures for machine translation and/or summarization*, 65–72.
- Carion, N.; Massa, F.; Synnaeve, G.; Usunier, N.; Kirillov, A.; and Zagoruyko, S. 2020. End-to-end object detection with transformers. In *European conference on computer vision*, 213–229. Springer.
- Chen, S.; Chen, J.; Jin, Q.; and Hauptmann, A. 2017. Video captioning with guidance of multimodal latent topics. In *Proceedings of the 25th ACM international conference on Multimedia*, 1838–1846.
- Cheng, F.; Wang, X.; Lei, J.; Crandall, D.; Bansal, M.; and Bertasius, G. 2023. Vindlu: A recipe for effective video-and-language pretraining. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 10739–10750.
- Dosovitskiy, A.; Beyer, L.; Kolesnikov, A.; Weissenborn, D.; Zhai, X.; Unterthiner, T.; Dehghani, M.; Minderer, M.; Heigold, G.; Gelly, S.; et al. 2020. An image is worth 16x16 words: Transformers for image recognition at scale. *arXiv preprint arXiv:2010.11929*.
- Dubey, A.; Jauhri, A.; Pandey, A.; Kadian, A.; Al-Dahle, A.; Letman, A.; Mathur, A.; Schelten, A.; Yang, A.; Fan, A.; et al. 2024. The llama 3 herd of models. *arXiv preprint arXiv:2407.21783*.
- Fujita, S.; Hirao, T.; Kamigaito, H.; Okumura, M.; and Nagata, M. 2020. SODA: Story oriented dense video captioning evaluation framework. In *Computer Vision—ECCV 2020: 16th European Conference, Glasgow, UK, August 23–28, 2020, Proceedings, Part VI 16*, 517–531. Springer.
- Gao, L.; Guo, Z.; Zhang, H.; Xu, X.; and Shen, H. T. 2017. Video Captioning With Attention-Based LSTM and Semantic Consistency. *IEEE Transactions on Multimedia*, 19(9): 2045–2055.
- Gao, Y.; Xiong, Y.; Gao, X.; Jia, K.; Pan, J.; Bi, Y.; Dai, Y.; Sun, J.; and Wang, H. 2023. Retrieval-augmented generation for large language models: A survey. *arXiv preprint arXiv:2312.10997*.
- Hayes, T.; Zhang, S.; Yin, X.; Pang, G.; Sheng, S.; Yang, H.; Ge, S.; Hu, Q.; and Parikh, D. 2022. Mugen: A playground for video-audio-text multimodal understanding and generation. In *European Conference on Computer Vision*, 431–449. Springer.
- Huang, G.; Pang, B.; Zhu, Z.; Rivera, C.; and Soricut, R. 2020. Multimodal pretraining for dense video captioning. *arXiv preprint arXiv:2011.11760*.
- Iashin, V.; and Rahtu, E. 2020a. A better use of audio-visual cues: Dense video captioning with bi-modal transformer. *arXiv preprint arXiv:2005.08271*.
- Iashin, V.; and Rahtu, E. 2020b. Multi-modal dense video captioning. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition workshops*, 958–959.
- Kim, M.; Kim, H. B.; Moon, J.; Choi, J.; and Kim, S. T. 2024. Do You Remember? Dense Video Captioning with Cross-Modal Memory Retrieval. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 13894–13904.
- Krishna, R.; Hata, K.; Ren, F.; Fei-Fei, L.; and Carlos Niebles, J. 2017. Dense-captioning events in videos. In *Proceedings of the IEEE international conference on computer vision*, 706–715.
- Kudo, T.; and Richardson, J. 2018. Sentencepiece: A simple and language independent subword tokenizer and detokenizer for neural text processing. *arXiv preprint arXiv:1808.06226*.
- Lewis, P.; Perez, E.; Piktus, A.; Petroni, F.; Karpukhin, V.; Goyal, N.; Küttler, H.; Lewis, M.; Yih, W.-t.; Rocktäschel, T.; et al. 2020. Retrieval-augmented generation for knowledge-intensive nlp tasks. *Advances in Neural Information Processing Systems*, 33: 9459–9474.
- Li, J.; Li, D.; Xiong, C.; and Hoi, S. 2022. Blip: Bootstrapping language-image pre-training for unified vision-language understanding and generation. In *International Conference on Machine Learning*, 12888–12900. PMLR.
- Lin, K.; Li, L.; Lin, C.-C.; Ahmed, F.; Gan, Z.; Liu, Z.; Lu, Y.; and Wang, L. 2022. Swinbert: End-to-end transformers with sparse attention for video captioning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 17949–17958.
- Luo, H.; Ji, L.; Shi, B.; Huang, H.; Duan, N.; Li, T.; Li, J.; Bharti, T.; and Zhou, M. 2020. UniVL: A Unified Video and Language Pre-Training Model for Multimodal Understanding and Generation. *arXiv:2002.06353*.
- Pan, Y.; Mei, T.; Yao, T.; Li, H.; and Rui, Y. 2016. Jointly modeling embedding and translation to bridge video and language. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, 4594–4602.
- Papineni, K.; Roukos, S.; Ward, T.; and Zhu, W.-J. 2002. Bleu: a method for automatic evaluation of machine translation. In *Proceedings of the 40th annual meeting of the Association for Computational Linguistics*, 311–318.

- Pei, W.; Zhang, J.; Wang, X.; Ke, L.; Shen, X.; and Tai, Y.-W. 2019. Memory-attended recurrent network for video captioning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 8347–8356.
- Qi, M.; Wang, Y.; Li, A.; and Luo, J. 2019. Sports video captioning via attentive motion representation and group relationship modeling. *IEEE Transactions on Circuits and Systems for Video Technology*, 30(8): 2617–2633.
- Radford, A.; Kim, J. W.; Hallacy, C.; Ramesh, A.; Goh, G.; Agarwal, S.; Sastry, G.; Askell, A.; Mishkin, P.; Clark, J.; et al. 2021. Learning transferable visual models from natural language supervision. In *International conference on machine learning*, 8748–8763. PMLR.
- Raffel, C.; Shazeer, N.; Roberts, A.; Lee, K.; Narang, S.; Matena, M.; Zhou, Y.; Li, W.; and Liu, P. J. 2020. Exploring the limits of transfer learning with a unified text-to-text transformer. *Journal of machine learning research*, 21(140): 1–67.
- Ramos, R.; Martins, B.; Elliott, D.; and Kementchedjhiya, Y. 2023. Smallcap: lightweight image captioning prompted with retrieval augmentation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2840–2849.
- Rohrbach, M.; Qiu, W.; Titov, I.; Thater, S.; Pinkal, M.; and Schiele, B. 2013. Translating video content to natural language descriptions. In *Proceedings of the IEEE international conference on computer vision*, 433–440.
- Rugg, M. D.; Fletcher, P. C.; Allan, K.; Frith, C. D.; Frackowiak, R.; and Dolan, R. J. 1998. Neural correlates of memory retrieval during recognition memory and cued recall. *Neuroimage*, 8(3): 262–273.
- Saharia, C.; Chan, W.; Saxena, S.; Li, L.; Whang, J.; Denton, E. L.; Ghasemipour, K.; Gontijo Lopes, R.; Karagol Ayan, B.; Salimans, T.; et al. 2022. Photorealistic text-to-image diffusion models with deep language understanding. *Advances in Neural Information Processing Systems*, 35: 36479–36494.
- Sarfraz, S.; Sharma, V.; and Stiefelhausen, R. 2019. Efficient parameter-free clustering using first neighbor relations. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 8934–8943.
- Seo, P. H.; Nagrani, A.; Arnab, A.; and Schmid, C. 2022. End-to-end generative pretraining for multimodal video captioning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 17959–17968.
- Tu, Y.; Li, L.; Su, L.; Gao, S.; Yan, C.; Zha, Z.-J.; Yu, Z.; and Huang, Q. 2022. I 2 Transformer: Intra-and inter-relation embedding transformer for TV show captioning. *IEEE Transactions on Image Processing*, 31: 3565–3577.
- Tu, Y.; Zhou, C.; Guo, J.; Gao, S.; and Yu, Z. 2021. Enhancing the alignment between target words and corresponding frames for video captioning. *Pattern Recognition*, 111: 107702.
- Tu, Y.; Zhou, C.; Guo, J.; Li, H.; Gao, S.; and Yu, Z. 2023. Relation-aware attention for video captioning via graph learning. *Pattern Recognition*, 136: 109204.
- Tulving, E.; et al. 1972. Episodic and semantic memory. *Organization of memory*, 1(381-403): 1.
- Vedantam, R.; Lawrence Zitnick, C.; and Parikh, D. 2015. Cider: Consensus-based image description evaluation. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, 4566–4575.
- Venugopalan, S.; Rohrbach, M.; Donahue, J.; Mooney, R.; Darrell, T.; and Saenko, K. 2015. Sequence to sequence-video to text. In *Proceedings of the IEEE international conference on computer vision*, 4534–4542.
- Venugopalan, S.; Xu, H.; Donahue, J.; Rohrbach, M.; Mooney, R.; and Saenko, K. 2014. Translating videos to natural language using deep recurrent neural networks. *arXiv preprint arXiv:1412.4729*.
- Wang, B.; Ma, L.; Zhang, W.; and Liu, W. 2018. Reconstruction network for video captioning. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, 7622–7631.
- Wang, T.; Zhang, R.; Lu, Z.; Zheng, F.; Cheng, R.; and Luo, P. 2021. End-to-End Dense Video Captioning with Parallel Decoding. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 6847–6857.
- Wang, T.; Zheng, H.; and Yu, M. 2020. Dense-Captioning Events in Videos: SYSU Submission to ActivityNet Challenge 2020. *arXiv preprint arXiv:2006.11693*.
- Wu, H.; Liu, H.; Qiao, Y.; and Sun, X. 2024. DIBS: Enhancing Dense Video Captioning with Unlabeled Videos via Pseudo Boundary Enrichment and Online Refinement. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 18699–18708.
- Yang, A.; Nagrani, A.; Laptev, I.; Sivic, J.; and Schmid, C. 2023a. VidChapters-7M: Video Chapters at Scale. In *NeurIPS*.
- Yang, A.; Nagrani, A.; Seo, P. H.; Miech, A.; Pont-Tuset, J.; Laptev, I.; Sivic, J.; and Schmid, C. 2023b. Vid2seq: Large-scale pretraining of a visual language model for dense video captioning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 10714–10726.
- Zhou, L.; Xu, C.; and Corso, J. 2018. Towards automatic learning of procedures from web instructional videos. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 32.
- Zhou, X.; Arnab, A.; Buch, S.; Yan, S.; Myers, A.; Xiong, X.; Nagrani, A.; and Schmid, C. 2024. Streaming dense video captioning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 18243–18252.