

Temporal Working Memory: Query-Guided Segment Refinement for Enhanced Multimodal Understanding

Xingjian Diao*, Chunhui Zhang*, Weiyi Wu, Zhongyu Ouyang,
Peijun Qing, Ming Cheng, Soroush Vosoughi, Jiang Gui

Dartmouth College

{xingjian.diao, chunhui.zhang, weiyi.wu}.gr@dartmouth.edu

{soroush.vosoughi, jiang.gui}@dartmouth.edu

Abstract

Multimodal foundation models (MFMs) have demonstrated significant success in tasks such as visual captioning, question answering, and image-text retrieval. However, these models face inherent limitations due to their finite internal capacity, which restricts their ability to process extended temporal sequences—an essential requirement for comprehensive video and audio analysis. To overcome these challenges, we introduce a specialized cognitive module, temporal working memory (TWM), which aims to enhance the temporal modeling capabilities of MFMs. It selectively retains task-relevant information across temporal dimensions, ensuring that critical details are preserved throughout the processing of video and audio content. The TWM uses a query-guided attention approach to focus on the most informative multimodal segments within temporal sequences. By retaining only the most relevant content, TWM optimizes the use of the model’s limited capacity, enhancing its temporal modeling ability. This plug-and-play module can be easily integrated into existing MFMs. With our TWM, nine state-of-the-art models exhibit significant performance improvements across tasks such as video captioning, question answering, and video-text retrieval. By enhancing temporal modeling, TWM extends the capability of MFMs to handle complex, time-sensitive data effectively. Our code is available at https://github.com/xid32/NAACL_2025_TWM.

1 Introduction

Multimodal foundation models (MFMs) have demonstrated impressive capabilities in tasks such as video captioning, question-answering, image-text retrieval, and broader multimodal understanding (Zhang et al., 2022; Kim et al., 2024; He et al., 2024; Jian et al., 2023, 2024; Zhang et al., 2025;

*contributed equally

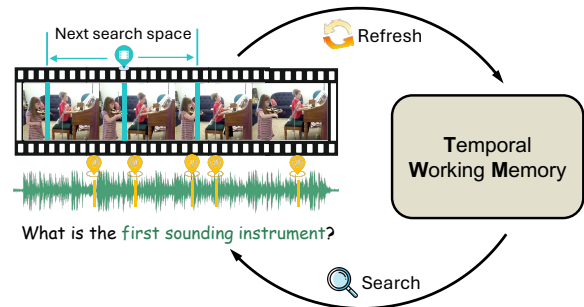


Figure 1: Temporal Working Memory (TWM): TWM employs search engine and memory refresh mechanisms to retain key segments in long multimodal inputs.

Yao et al., 2024a; Xie et al., 2024; Yao et al., 2024b; Xie et al., 2025; Han et al., 2024a; Liu et al., 2024e; Lin et al., 2024; Han et al., 2024b; Liu et al., 2024a). While MFMs excel at processing multimodal inputs, MFMs are often not equipped to explicitly reduce the input context burden, particularly in extracting query-relevant information from the input context for video understanding tasks.

In humans, working memory retains and processes information over short time spans with limited capacity (Baddeley, 2000), and similar constraints apply to MFMs (Liu et al., 2024b). For example, LLaMA has a context length limit of 2048 tokens (Touvron et al., 2023a,b; Dubey et al., 2024), while LLaVA (Liu et al., 2023) and BLIP-2 (Li et al., 2023b) can handle only 256 and 32 tokens per image, respectively. These limited capacities prevent models from effectively retaining sufficient information over extended temporal spans.

While operating under similar constraints, the human cognitive system has evolved effective mechanisms for efficient information processing (Baddeley, 2000), such as selective retention and distraction filtering, to dynamically prioritize relevant information while discarding irrelevant details (Zhang et al., 2024; Gong et al., 2024). In contrast, current MFMs lack such selective mechanisms found in human working memory, prevent-

ing them from effectively filtering and retaining the most relevant temporal segments from multimodal inputs (e.g., video frames or audio clips). Consequently, MFMs tend to process the entire input sequences indiscriminately within their perceptual window, leading to inefficient utilization of the model’s capabilities.

Recent developments in the memory of LLMs have significantly improved their ability to manage temporal contexts. Various approaches have been proposed to address the inherent memory limitations of LLMs and thereby improve their performance on complex tasks. [Li et al. \(2023a\)](#) proposed a knowledge-aware fine-tuning method that instructs LLMs to prioritize relevant external context while minimizing reliance on internal pre-trained knowledge, thereby enhancing their memory by filtering out irrelevant information. Building on this, [Gong et al. \(2024\)](#) demonstrated that ChatGPT’s working memory is similar to that of humans, suggesting that enhancing memory capacity is crucial for advancing the intelligence of AI systems. Further exploring these limitations, [Zhang et al. \(2024\)](#) recommended strategies for more efficient memory utilization, highlighting the importance of improving both memory retention and model autonomy for better reasoning capabilities. In the context of multimodal models, [Wu and Xie \(2024\)](#) integrated visual working memory mechanisms to help models focus on essential features in high-resolution images, significantly improving visual grounding performance.

Despite recent advances in LLM memory management, these models still face fundamental challenges in dynamic multimodal and temporal reasoning: (i) they lack effective mechanisms to filter and retain query-relevant information from multimodal inputs (e.g., audio, video, language), leading to indiscriminate processing of entire sequences regardless of their relevance. (ii) these models lack the ability to effectively capture temporal dependencies. They struggle with short-term changes (e.g., rapid changes in audio or visual content) and long-term relationships (e.g., understanding how earlier events relate to later ones). This reduces their ability to reason about sequential and time-sensitive data, which is crucial for tasks involving events that unfold over time. (iii) the models’ inability to efficiently process large volumes of raw data leads to information overload, straining their limited capacity and degrading performance when dealing with complex and high-volume data.

Therefore, we draw on human working memory to effectively **extract query-relevant multimodal information across temporal dimensions**. We propose a multimodal Temporal Working Memory (TWM) mechanism for MFMs, as shown in [Figure 1](#). This mechanism employs a query-guided attention mechanism to selectively retain only query-relevant audio and visual inputs, focusing on the most informative segments along the temporal axis. The TWM mechanism constructs a temporal memory buffer at the model input stage, enabling MFMs to efficiently store and manage critical information across time. By concentrating on the retention of the most relevant data, TWM significantly improves the model’s ability to reason over extended temporal sequences in a multimodal context. Our contributions are:

- We propose a Temporal Working Memory (TWM) mechanism with an integrated query-guided selection module. This module directs the model to retain key segments in long video and audio sequences, optimizing the use of the model’s limited capacity.
- We employ a multi-scale temporal attention mechanism for both local and global dependencies, enabling accurate identification of relevant video-audio segments across temporal inputs.
- We integrate TWM into nine state-of-the-art MFMs and evaluate our approach on three large-scale multimodal benchmarks, covering tasks including audio-visual question answering, video captioning, and video-text retrieval. Our approach effectively yields significant performance improvements across all tasks.

2 Related Works

Temporal Modeling in MLLMs Multimodal LLMs (MLLMs) for long-video understanding aim to capture long-range temporal patterns. A common strategy is temporal pooling, as used in VideoChatGPT ([Maaz et al., 2024](#)), but this can limit performance due to inadequate temporal modeling. More advanced methods, such as video-LLAMA ([Zhang et al., 2023](#)), incorporate video query transformers to enhance temporal dynamics, but this comes with increased model complexity. To reduce computational demands, some models rely on pre-extracted features, avoiding joint training of backbone architectures ([Hussein et al., 2019](#);

Liu et al., 2024d; Wu and Krahenbuhl, 2021). Techniques like Vis4mer (Islam and Bertasius, 2022) and S5 (Wang et al., 2023) utilize the S4 transformer architecture (Gu et al., 2022) for efficient long-range temporal modeling. Recent developments, such as online video processing (He et al., 2024), employ memory banks to track past content for long-term analysis. In contrast, we propose a TWM mechanism that retains only query-relevant multimodal inputs through search engines within a temporal context.

Video Understanding Video understanding tasks evaluate a model’s ability to process multimodal content, focusing on both temporal and semantic aspects. Key tasks for long-video understanding include audio-visual question answering (AVQA), video captioning, and video-text retrieval, supported by extensive research and large-scale datasets (Liu et al., 2024c; Xu et al., 2016; Bain et al., 2020). Prior AVQA methods finetuned pretrained visual models with adapters (Liu et al., 2024c; Lin et al., 2023; Duan et al., 2023; Diao et al., 2024), while recent approaches use unified multimodal encoders with LLMs (Han et al., 2024a). Video captioning models employ graph neural networks (GNNs) (Hendria et al., 2023), simplified image-text architectures (Wang et al., 2022), or causal effect networks (CEN) to enhance temporal coherence (Nadeem et al., 2024). In video-text retrieval, adaptive frame aggregation reduces visual tokens to accelerate encoding (Ren et al., 2023). In contrast to previous work focusing on specific multimodal applications, this study emphasizes the role of TWM in enhancing fundamental temporal grounding across audio, video, and language.

3 Temporal Working Memory

Our temporal working memory (TWM) framework is a complementary architecture for multimodal large language models, integrating visual and auditory buffer components. If the model does not involve audio, the auditory component can be omitted. The working pipeline of the TWM framework is outlined in Algorithm 1, while the search and memory update processes are depicted in Figure 2.

3.1 Visual Memory Management

3.1.1 Query-Relevant Video-Frame Search

Our mechanism emulates human cognitive strategies by identifying and retaining critical visual

Algorithm 1 TWM for MFMs

```
# search query-related segments from video&audio
def neural_search(video, audio, qry):
    # Step 1: Segment and encode video and audio
    v_seg, a_seg = segment(video, audio)
    v_embs, a_embs = encode(v_seg, a_seg)

    # Step 2: Calculate relevance scores
    v_scores, a_scores = sim(v_embs, a_embs, qry)

    # Step 3: Iterate to select the relevant segments
    v_buffer, a_buffer = select(v_embs, a_embs,
                               v_scores, a_scores)

    return v_buffer, a_buffer

# temporal working memory for video and audio
v_buffer, a_buffer = neural_search(video, audio,
                                   qry)

# MFM with temporal working memory
output = MFM(v_buffer, a_buffer, qry)
```

information from long video sequences. It alternates between two key operations—**search** and **update**—to dynamically adjust memory and focus on the most relevant segments, as shown in Figure 2.

Initially, k frames are selected from the full sequence of N frames and processed through a visual encoder to generate embeddings. Each frame v_i is assigned a **Similarity Score** ($S(v_i)$), defined as:

$$S(v_i) = \alpha_1 D(v_i) + \alpha_2 R(v_i, q), \quad (1)$$

where $D(v_i)$ is function representing the **distinctiveness** of frame v_i , $R(v_i, q)$ is function representing the **relevance** of frame v_i to the query q . α_1 and α_2 are adaptive weights that vary based on the nature of the downstream task and the duration of the video samples. The selection process is iterative. In each iteration, the frame with the highest similarity score ($S(v_i)$) is chosen as the midpoint, and k frames are searched uniformly within a range of $\frac{N}{k}$ around it. These frames are added to visual memory, excluding frames already present to maintain diversity. The process concludes upon convergence.

3.1.2 Training Neural Search Engine

To identify frames relevant to a given query, we use a **cross-modal alignment strategy** (Figure 3). Pre-trained visual and language encoders are employed, with a linear projection layer that maps visual embeddings into the textual embedding space. The **InfoNCE loss** (Oord et al., 2018) is used to opti-

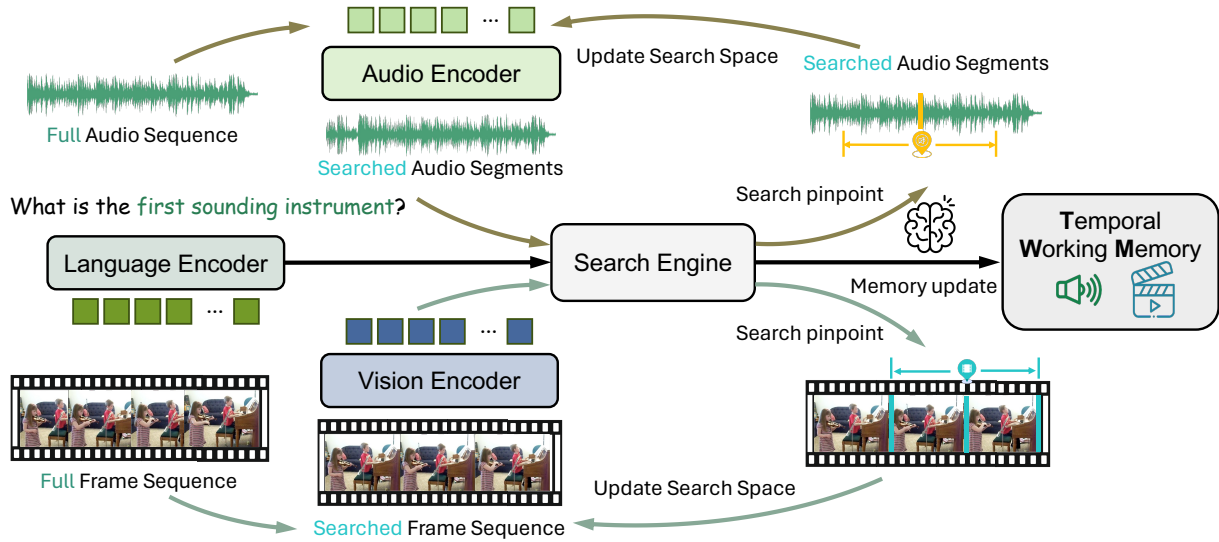


Figure 2: The temporal working memory (TWM) pipeline retains the most relevant segments from video and audio inputs based on a language query. The Language Encoder processes the query, guiding the Search Engine to identify and select key video and audio segments. TWM ensures the retention of only the most informative data, enabling the efficient utilization of multimodal foundation models’ capabilities.

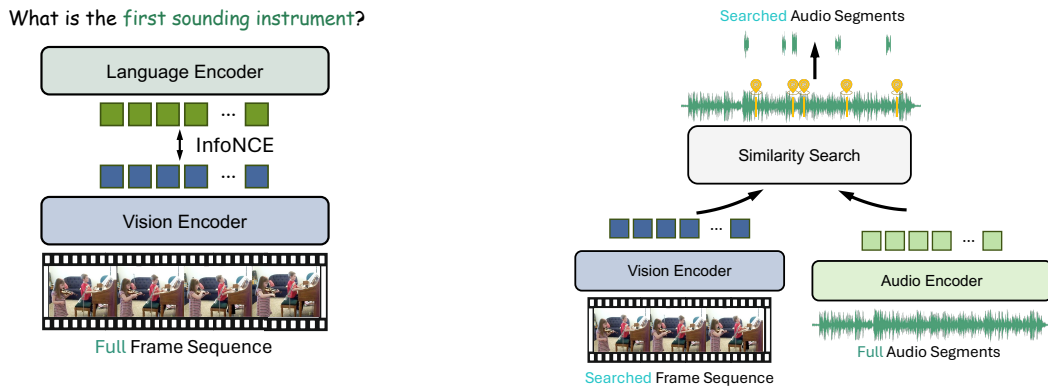


Figure 3: Aligning frames with language query. A linear projection layer trained with InfoNCE loss aligns visual embeddings with query-based anchors.

mize this alignment:

$$\mathcal{L}_{\text{InfoNCE}} = -\log \frac{\exp\left(\frac{\text{sim}(\mathbf{e}_v, \mathbf{e}_{t_i})}{\tau}\right)}{\sum_{j=1}^N \exp\left(\frac{\text{sim}(\mathbf{e}_v, \mathbf{e}_{t_j})}{\tau}\right)}, \quad (2)$$

where \mathbf{e}_v is the embedding of video frame v , \mathbf{e}_{t_i} is the embedding of text description t_i , $\text{sim}(x, y)$ denotes the cosine similarity function, τ is the temperature parameter controlling the sharpness of the distribution, and N represents the number of negative samples. The InfoNCE loss maximizes the similarity between the corresponding video and text embeddings while minimizing the similarity to unrelated samples. This ensures that the model effectively aligns the most relevant frames with their corresponding text, thereby optimizing its ability to retain meaningful visual information.

Figure 4: Similarity search for query-relevant audio segments. The audio encoder utilizes visual embeddings as a query to search for the most relevant audio segments, updating the auditory buffer to retain only the essential audio information.

3.2 Auditory Memory Management

3.2.1 Query-Relevant Audio-Segment Search

The search process for identifying key audio segments mirrors the methodology used for video frames. The audio sequence is divided into predefined segments, typically 5-6 segments depending on video length, to model adequate temporal dependencies for tasks like AVQA and video captioning. This segmentation allows the model to focus efficiently on relevant audio intervals, improving attention allocation across extended sequences.

Building upon Pathformer’s dual-attention mechanism (Chen et al., 2024), we extend their approach to enhance the correlation between audio and video

data. Specifically, visual embeddings are used as queries in both attentions (Figure 4). To enable multimodal audio-visual synchronization, we integrate audio patches derived from Mel-spectrograms for temporal segmentation. This facilitates the alignment between the audio and visual inputs. Our audio encoder employs two key attention mechanisms:

- Inter-segment attention is designed to model *global dependencies* across audio segments, enabling the model to capture broader relationships such as shifts in tone, mood, or overall sound context. Specifically, inter-segment attention calculates attention scores between the query Q , which is derived from visual features, and the keys K_i from the audio segments $Att_{inter} = \text{softmax}\left(\frac{QK_i^T}{\sqrt{d_K}}\right)V_i$, $i \in [1, n]$. Here, Q represents the visual embeddings, while K_i and V_i are the audio embeddings from segment i . By using visual features as the query, this attention aligns audio information with relevant visual cues, effectively capturing how the audio context evolves to the video over time.
- Intra-segment attention aims to capture *local dependencies* within individual audio segments, thereby modeling fine-grained temporal patterns such as audio variations or sudden changes in sound effects. The result of intra-segment attention for each segment is computed as: $Att_{intra} = \text{Concat}(Att_{intra_i} \mid i = 1, \dots, n)$. Att_{intra_i} represents the computed attention within each segment i . The concatenation operation aggregates these intra-segment attention results across all segments, ensuring that short-term changes are captured and preserved for subsequent processing. This aggregation allows the model to retain a detailed representation of the short-term dynamics within each segment.

In the fusion layer, we apply a cross-modal attention mechanism to synchronize features from both audio and visual streams. Additionally, multi-kernel pooling aggregates audio patches across different-scale temporal dependencies, enhancing the alignment and understanding of temporal multimodal inputs.

As shown in Figure 5, a pretrained visual feature extractor is used to align the audio segments with their corresponding visual frames to establish cross-modal coherence. To identify query-relevant audio patches, we apply cosine similarity between audio

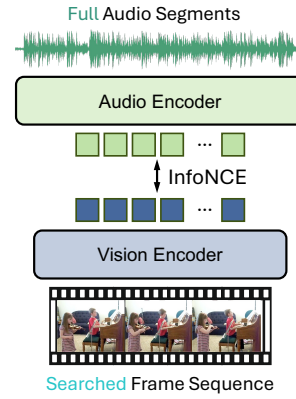


Figure 5: Audio segments aligned with query-relevant frames. An audio encoder learns temporal distance and resolution for audio-visual embedding alignment.

embeddings (e_{a_i}) and visual embeddings (e_v) as $\text{sim}(e_{a_i}, e_v) = \frac{e_{a_i} \cdot e_v}{\|e_{a_i}\| \|e_v\|}$. This similarity score is used to select the audio patches that are most relevant to the corresponding visual frames. The selected audio segments are then updated in the auditory buffer, ensuring that the most important audio information is retained. This iterative refinement enhances the synchronization and complementarity between audio and visual content.

3.2.2 Training Audio Search Engine

To identify query-relevant audio segments, we also use InfoNCE loss to achieve cross-modal alignment (see Figure 5). Let e_{a_i} denote the embedding of the audio patch a_i , and let e_v denote the embedding of the video frame v . The alignment loss is:

$$\mathcal{L}_{\text{InfoNCE}} = -\log \frac{\exp\left(\frac{\text{sim}(e_v, e_{a_i})}{\tau}\right)}{\sum_{j=1}^N \exp\left(\frac{\text{sim}(e_v, e_{a_j})}{\tau}\right)}, \quad (3)$$

where τ is a temperature parameter controlling the distribution’s sharpness, and N is the number of negative samples. This approach ensures effective alignment between audio and visual embeddings, allowing the model to identify cross-modal relationships effectively and refine the auditory buffer.

4 Experiments

We perform three major experiments to validate the effectiveness of the Temporal Working Memory mechanism discussed in the previous section. We evaluated the performance of state-of-the-art baseline models and the same models augmented with our Temporal Working Memory mechanism on the following downstream tasks: (1) **audio-visual question answering (AVQA)**, (2) **video captioning**, and (3) **video-text retrieval**.

4.1 Setup

Datasets We conduct experiments on AVQA, video captioning, and video-text retrieval:

- **MUSIC-AVQA v2.0** (Liu et al., 2024c): MUSIC-AVQA v2.0 introduces 1,230 additional videos and 8,100 new question-answer (QA) pairs to further mitigate data bias. It builds on the original MUSIC-AVQA dataset (Li et al., 2022), which contains 9,288 videos of 22 musical instruments (over 150 hours) with 45,867 question-answer (QA) pairs across 33 templates in 9 categories.
- **MSR-VTT** (Xu et al., 2016): Comprises 10,000 video clips (over 41 hours) from various online sources. Each video has 20 human-annotated captions, totaling 200,000 video-text pairs across diverse categories like music, sports, and gaming.
- **CMD** (Bain et al., 2020): The Condensed Movies Dataset includes over 33,000 clips from 3,600 movies, averaging 2 minutes each, with annotations such as captions, dialogues, and action labels, ideal for video-text retrieval tasks.

Baselines Details of the baseline models can be found in Appendix A. We evaluate nine state-of-the-art MFMs reproduced with open-source code and pretrained weights.

Evaluation Metrics Below standard metrics reflect the accuracy, quality, and retrieval capabilities:

- **Audio-Visual Question Answering:** Accuracy is measured for Audio (Counting, Comparative), Visual (Counting, Location), Audio-Visual (Existential, Counting, Location, Comparative, Temporal) question types, along with average accuracies for Audio, Visual, Audio-Visual, and overall.
- **Video Captioning:** Metrics include ROUGE-L, CIDEr, and SPICE, assessing overlap with ground truth, consensus with human annotations, and diversity of scene descriptions, respectively.
- **Video-Text Retrieval:** Metrics are Recall@1, Recall@5, and Recall@10, measuring retrieval performance within top 1, 5, and 10 predictions.

Implementations The settings of each dataset:

- **MUSIC-AVQA v2.0:** Videos are 60 seconds long, with questions targeting specific portions of the video. We set $k = 11$ and ran 6 iterations, using $\alpha_1 = 0.2$ and $\alpha_2 = 0.8$, resulting in frame sampling rates consistent with 1 frame per second (fps). Audio segments are extracted every 5 seconds, selecting the highest-scoring segment from a total of 12 segments.

- **MSR-VTT:** With a frame rate of 21 fps, we set $k = 3$ and ran 3 iterations, yielding 8–9 searched frames, with $\alpha_1 = 0.5$ and $\alpha_2 = 0.5$ for balanced frame selection.
- **CMD:** With a frame rate of 30 fps, we used $k = 5$ and ran 7 iterations, producing 30–35 frames in total, using $\alpha_1 = 0.6$ and $\alpha_2 = 0.4$ to prioritize frame diversity.

As depicted in Figure 3, the dimensions of the linear mapping layer, and similarly in Figure 5 for the fusion layer, correspond to the output embedding dimensions of the text encoder used by each model. The embedding dimensions typically span 768-D, 4096-D, or 16384-D, as detailed in Section 4.1 where the baseline models are referenced.

Training is conducted on the NVIDIA H100 80GB GPUs using PyTorch. The Adam optimizer with a learning rate of $1e^{-4}$ is used, and each model is trained for 10 epochs. On the MUSIC-AVQA dataset, each model requires an average training time of 65.2 hours. For the MSR-VTT dataset, the average training time per model is 14.6 hours, while for the CMD dataset, each model takes approximately 146.6 GPU-hours to train.

4.2 Overall Comparisons

4.2.1 Audio-Visual Question Answering

TWM captures fine-grained multimodal dependencies for comparative reasoning In AVQA (Table 1), TWM excels in identifying and preserving fine-grained dependencies between audio and visual inputs, especially in complex comparative tasks. In audio-related comparative QA, LAVisH+TWM improves by 12.40%, DG-SCT+TWM gains 10.12% and LAST-Att+TWM shows an increase of 11.73%. Similarly, significant improvements are observed in the audiovisual comparative QA: LAVisH+TWM improves by 10.08%, DG-SCT+TWM by 13.56% and LAST-Att+TWM by 10.92%. TWM also leads to significant increases in overall average accuracy across all audiovisual tasks, with LAVisH+TWM improving by 6.35%, DG-SCT+TWM by 8.58% and LAST-Att+TWM by 5.01%. By focusing on query-relevant segments and filtering out irrelevant content, TWM ensures that the model attends to the most informative parts of each modality, thereby improving cross-modal reasoning accuracy. This selective attention mechanism allows the model to better isolate critical elements within audiovisual streams, enabling it to reason about context-dependent relationships be-

Method	Audio-related QA			Visual-related QA			Audio&Visual-related QA						Avg
	Count	Comp	Avg	Count	Local	Avg	Exist	Count	Local	Comp	Temp	Avg	
LAVisH (Lin et al., 2023)	83.82	58.19	75.72	82.81	81.73	82.30	73.26	73.45	65.64	64.26	60.82	67.75	73.28
LAVisH + TWM	79.22	70.59	76.91	84.52	84.10	84.31	78.66	65.21	73.05	74.34	57.24	74.10	74.42
Gain (Δ)	-4.60	+12.40	+1.19	+1.71	+2.37	+2.01	+5.40	-8.24	+7.41	+10.08	-3.58	+6.35	+1.14
DG-SCT (Duan et al., 2023)	83.13	62.54	76.62	81.61	82.76	82.19	83.43	72.70	64.65	64.78	67.34	70.38	74.53
DG-SCT + TWM	80.05	72.66	77.20	87.77	85.24	86.35	88.69	87.21	76.06	78.34	59.82	78.96	79.22
Gain (Δ)	-3.08	+10.12	+0.58	+6.16	+2.48	+4.16	+5.26	+14.51	+11.41	+13.56	-7.52	+8.58	+4.69
LAST-Att (Liu et al., 2024c)	86.03	62.52	79.44	84.12	84.01	84.05	76.21	75.23	68.91	65.60	60.60	69.04	75.44
LAST-Att + TWM	79.52	74.25	76.43	88.52	85.98	87.01	80.12	82.40	76.06	76.52	55.82	74.05	77.96
Gain (Δ)	-6.51	+11.73	-3.01	+4.40	+1.97	+2.96	+3.91	+7.17	+7.15	+10.92	-4.78	+5.01	+2.52

Table 1: Results of different models on the test set of MUSIC-AVQA 2.0 (Liu et al., 2024c). **Bold** results indicate the better performance.

tween different inputs. The fine-tuned multimodal focus highlights TWM’s effectiveness in tasks requiring fine-grained comparisons, as it actively suppresses noise and enhances relevant signals.

4.2.2 Video Captioning

Method	ROUGE-L ↑	CIDEr ↑	SPICE ↑
Git (Wang et al., 2022)	24.51	32.43	13.70
Git + TWM	26.10	39.25	14.31
Gain (Δ)	+1.59	+6.82	+0.61
AKGNN (Hendria et al., 2023)	21.42	25.90	11.99
AKGNN + TWM	21.33	27.46	11.02
Gain (Δ)	-0.09	+1.56	-0.97
CEN (Nadeem et al., 2024)	27.90	49.87	15.76
CEN + TWM	28.10	52.01	15.90
Gain (Δ)	+0.20	+2.14	+0.14

Table 2: Test results of different models on MSR-VTT.

TWM enhances temporal coherence through selective attention to key events In video captioning tasks (Table 2), TWM’s selective attention mechanism significantly improves temporal coherence by focusing on key events and discarding irrelevant content. For example, Git+TWM achieves a 6.82% improvement in CIDEr and a 1.59% increase in ROUGE-L, highlighting the model’s enhanced ability to generate coherent narratives that follow the flow of events. By retaining only the most contextually relevant audio-visual segments, TWM helps to avoid disjointed or fragmented scene descriptions, which is critical for accurately representing long or complex narratives.

TWM captures fine-grained scene transitions, enhancing descriptive richness TWM also excels at capturing fine-grained details within scenes, allowing the model to generate richer and more informative descriptions. For example, CEN+TWM shows a 0.14% improvement in SPICE, reflecting the model’s enhanced ability to capture var-

ied and accurate content in captions. In addition, Git+TWM shows a 0.61% increase in SPICE, indicating an improved ability to capture changes in the audiovisual content of the video, such as changes in action or context. By dynamically updating memory with the most relevant visual and auditory elements, TWM ensures that critical details are highlighted, resulting in more contextually accurate and detailed output. This enhanced descriptive richness is essential in scenarios involving complex or rapidly changing scenes, where capturing semantic shifts in the narrative is key to generating informative captions.

4.2.3 Video-Text Retrieval

Method	Recall@1 ↑	Recall@5 ↑	Recall@10 ↑
VINDLU (Cheng et al., 2023)	18.4	36.4	41.8
VINDLU + TWM	20.5	38.2	44.6
Gain (Δ)	+2.1	+1.8	+2.8
TESTA (Ren et al., 2023)	21.5	42.4	50.7
TESTA + TWM	22.1	45.3	52.1
Gain (Δ)	+0.6	+2.9	+1.4
MovieSeq (Lin et al., 2024)	25.8	45.3	50.3
MovieSeq + TWM	27.5	47.0	51.1
Gain (Δ)	+1.7	+1.7	+0.8

Table 3: Test results of different models on CMD.

**TWM maintains retrieval performance across broader scopes through adaptive segment re-
tention** TWM’s ability to enhance cross-modal alignment extends beyond immediate retrieval precision to broader retrieval tasks (Table 3). For instance, VINDLU + TWM achieves improvements of 2.1% in Recall@1, 1.8% in Recall@5, and 2.8% in Recall@10. TESTA + TWM demonstrates gains of 2.9% in Recall@5 and 1.4% in Recall@10, showcasing TWM’s capacity to retain relevant segments even in complex or diverse video datasets. Similarly, MovieSeq + TWM shows consistent improvements with a 1.7% increase in Recall@1 and

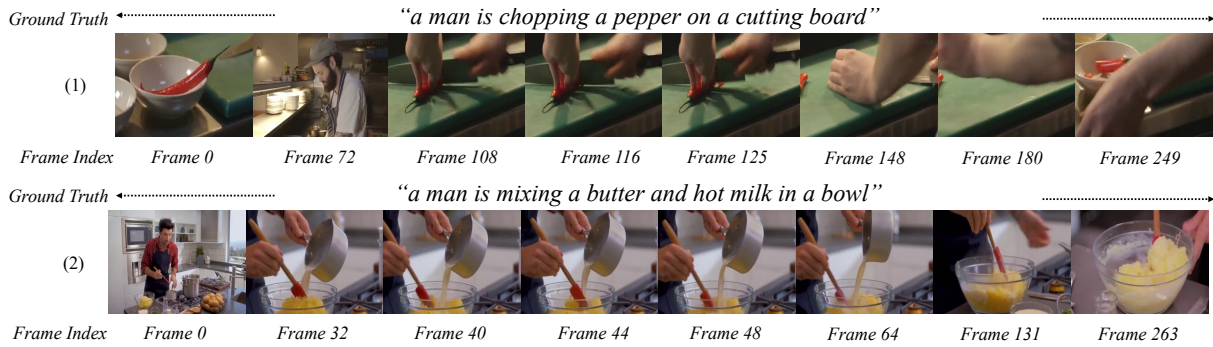


Figure 6: TWM-searched frames for the video captioning task, integrated with AKGNN (Hendria et al., 2023). Two examples of frames searched by TWM are presented alongside their ground truth captions. The total number of frames per clip in MSR-VTT (Xu et al., 2016) ranges from 210 to 630. TWM-searched 8 frames effectively encapsulate the key visual information required to generate the ground truth captions.

Recall@5, and a 0.8% gain in Recall@10. These results indicate that TWM’s memory update mechanism is flexible enough to adapt to a wide range of retrieval tasks. By selectively focusing on the most important audio-visual elements, TWM improves the model’s ability to retrieve relevant content across larger candidate sets. This adaptive retention mechanism allows the model to effectively balance precision and scope, ensuring that both specific and broad retrieval queries benefit from TWM’s selective attention and memory update strategies.

4.3 The Impacts of Temporal Sequence Selection

To illustrate the effectiveness of TWM, we present a case study where TWM’s search engine selects highly relevant frames based on the input of the language query (Figure 6). Below, we provide a detailed analysis of how TWM ensures completeness of action representation, eliminates irrelevant noise, and optimizes model performance through selective frame reduction.

Completeness of action representation TWM captures all the core stages of primary actions like chopping or mixing. For the chopping example, it captures the pepper being placed on the cutting board (Frames 0 and 72), the initiation of slicing (Frames 108 and 116), intermediate chopping stages (Frames 125 and 148), and the completion of the task (Frames 180 and 249). In the mixing example, it includes frames that depict the pouring of ingredients (Frames 40 and 44), the stirring process (Frames 48, 64, and 131), and the final state of the mixture. By covering all key moments, TWM pro-

vides the captioning model with a comprehensive understanding of the actions.

Elimination of irrelevant noise In cooking videos, various elements can distract from the main actions, such as the kitchen background, other utensils, or idle moments unrelated to the primary tasks like chopping. By selecting only the frames that focus on essential actions such as chopping or mixing in the examples, the TWM efficiently filters out these distractions, providing cleaner visual information for the captioning model. By minimizing distracting frames in the captioning inputs, the resulting descriptions of the events in the video become more accurate and useful.

Selective frame reduction for efficient model capacity utilization By selecting a limited number of informative frames that encompass all core stages, TWM optimizes memory usage and computational resources. It directs computational resources to the most relevant parts of the sequence by eliminating redundant or unrelated frames, thereby enhancing performance without overburdening the model. This selective retention not only allows for the precise capture of essential actions, but also facilitates efficient processing, significantly saving memory and speeding up the captioning model. Additional case studies on video-text retrieval and audio-visual question answering tasks are available in the Appendix C.

5 Conclusion

We introduce Temporal Working Memory (TWM), a cognitive module designed to enhance multi-modal foundation models’ capabilities on complex

video-audio-language tasks. Multimodal foundation models often struggle with dynamic multimodal input due to inefficient utilization of limited internal capacity. TWM addresses these limitations by (i) maintaining memory states to process temporal multimodal sequences through selective segment retention; (ii) modeling multi-scale temporal dependencies between video and audio inputs; and (iii) extracting query-relevant information from rich multimodal data for efficient memory utilization. Our approach effectively improves the performance of nine state-of-the-art multimodal models on three temporal reasoning tasks: audio-visual question answering, video captioning, and video-text retrieval.

Limitations

The effectiveness of TWM has been demonstrated on specific multimodal tasks and benchmarks, such as video captioning, question answering, and video-text retrieval. However, its generalizability to other in-the-wild domains remains unexplored. Extending TWM's application to other more practical multimodal tasks would be future direction for real-world applications.

Ethical Considerations

We examined the study describing the publicly available datasets used in this research and identified no ethical issues regarding the datasets.

Acknowledgment

This study is supported by the Department of Defense grant HT9425-23-1-0267.

References

Alan Baddeley. 2000. The episodic buffer: a new component of working memory? *Trends in cognitive sciences*.

Max Bain, Arsha Nagrani, Andrew Brown, and Andrew Zisserman. 2020. Condensed movies: Story based retrieval with contextual embeddings. In *Asian Conference on Computer Vision*.

Peng Chen, Yingying Zhang, Yunyao Cheng, Yang Shu, Yihang Wang, Qingsong Wen, Bin Yang, and Chenjuan Guo. 2024. Pathformer: Multi-scale transformers with adaptive pathways for time series forecasting. In *International Conference on Learning Representations*.

Feng Cheng, Xizi Wang, Jie Lei, David Crandall, Mohit Bansal, and Gedas Bertasius. 2023. Vindlu: A recipe for effective video-and-language pretraining. In *Conference on Computer Vision and Pattern Recognition*.

Xingjian Diao, Chunhui Zhang, Tingxuan Wu, Ming Cheng, Zhongyu Ouyang, Weiyi Wu, and Jiang Gui. 2024. Learning musical representations for music performance question answering. In *Findings of the Association for Computational Linguistics: EMNLP 2024*.

Haoyi Duan, Yan Xia, Zhou Mingze, Li Tang, Jieming Zhu, and Zhou Zhao. 2023. Cross-modal prompts: Adapting large pre-trained models for audio-visual downstream tasks. In *Advances in Neural Information Processing Systems*.

Abhimanyu Dubey, Abhinav Jauhri, Abhinav Pandey, Abhishek Kadian, Ahmad Al-Dahle, Aiesha Letman, Akhil Mathur, Alan Schelten, Amy Yang, Angela Fan, et al. 2024. The llama 3 herd of models. *arXiv preprint arXiv:2407.21783*.

Dongyu Gong, Xingchen Wan, and Dingmin Wang. 2024. Working memory capacity of chatgpt: An empirical study. In *AAAI Conference on Artificial Intelligence*.

Albert Gu, Karan Goel, and Christopher Ré. 2022. Efficiently modeling long sequences with structured state spaces. In *International Conference on Learning Representations*.

Jiaming Han, Kaixiong Gong, Yiyuan Zhang, Jiaqi Wang, Kaipeng Zhang, Dahua Lin, Yu Qiao, Peng Gao, and Xiangyu Yue. 2024a. Onellm: One framework to align all modalities with language. In *Conference on Computer Vision and Pattern Recognition*.

Xiaotian Han, Yiren Jian, Xuefeng Hu, Haogeng Liu, Yiqi Wang, Qihang Fan, Yuang Ai, Huaibo Huang, Ran He, Zhenheng Yang, and Quanzeng You. 2024b. InfiMM-webmath-40b: Advancing multimodal pre-training for enhanced mathematical reasoning. In *The 4th Workshop on Mathematical Reasoning and AI at NeurIPS'24*.

Bo He, Hengduo Li, Young Kyun Jang, Menglin Jia, Xuefei Cao, Ashish Shah, Abhinav Shrivastava, and Ser-Nam Lim. 2024. Ma-lmm: Memory-augmented large multimodal model for long-term video understanding. In *Conference on Computer Vision and Pattern Recognition*.

Willy Fitra Hendria, Vania Velda, Bahy Helmi Hartoyo Putra, Fikriansyah Adzaka, and Cheol Jeong. 2023. Action knowledge for video captioning with graph neural networks. *Journal of King Saud University-Computer and Information Sciences*.

Noureddien Hussein, Efstratios Gavves, and Arnold WM Smeulders. 2019. Timeception for complex action recognition. In *Conference on Computer Vision and Pattern Recognition*.

Md Mohaiminul Islam and Gedas Bertasius. 2022. Long movie clip classification with state-space video models. In *European Conference on Computer Vision*.

- Yiren Jian, Chongyang Gao, and Soroush Vosoughi. 2023. Bootstrapping vision-language learning with decoupled language pre-training. In *Advances in Neural Information Processing Systems*.
- Yiren Jian, Tingkai Liu, Yunzhe Tao, Chunhui Zhang, Soroush Vosoughi, and Hongxia Yang. 2024. Expedited training of visual conditioned language generation via redundancy reduction. In *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics*.
- Byoungjip Kim, Dasol Hwang, Sungjun Cho, Youngsoo Jang, Honglak Lee, and Moontae Lee. 2024. Show think and tell: Thought-augmented fine-tuning of large language models for video captioning. In *Conference on Computer Vision and Pattern Recognition*.
- Daliang Li, Ankit Singh Rawat, Manzil Zaheer, Xin Wang, Michal Lukasik, Andreas Veit, Felix Yu, and Sanjiv Kumar. 2023a. Large language models with controllable working memory. In *Findings of the Association for Computational Linguistics: ACL*.
- Guangyao Li, Yake Wei, Yapeng Tian, Chenliang Xu, Ji-Rong Wen, and Di Hu. 2022. Learning to answer questions in dynamic audio-visual scenarios. In *Conference on Computer Vision and Pattern Recognition*.
- Junnan Li, Dongxu Li, Silvio Savarese, and Steven Hoi. 2023b. BLIP-2: Bootstrapping language-image pre-training with frozen image encoders and large language models. In *International Conference on Machine Learning*.
- Kevin Qinghong Lin, Pengchuan Zhang, Difei Gao, Xide Xia, Joya Chen, Ziteng Gao, Jinheng Xie, Xuhong Xiao, and Mike Zheng Shou. 2024. Learning video context as interleaved multimodal sequences. In *European Conference on Computer Vision*.
- Yan-Bo Lin, Yi-Lin Sung, Jie Lei, Mohit Bansal, and Gedas Bertasius. 2023. Vision transformers are parameter-efficient audio-visual learners. In *Conference on Computer Vision and Pattern Recognition*.
- Haogeng Liu, Quanzeng You, Yiqi Wang, Xiaotian Han, Bohan Zhai, Yongfei Liu, Wentao Chen, Yiren Jian, Yunzhe Tao, Jianbo Yuan, Ran He, and Hongxia Yang. 2024a. InfiMM: Advancing multimodal understanding with an open-sourced visual language model. In *Findings of the Association for Computational Linguistics: ACL 2024*.
- Haotian Liu, Chunyuan Li, Qingyang Wu, and Yong Jae Lee. 2023. Visual instruction tuning. In *Advances in Neural Information Processing Systems*.
- Nelson F Liu, Kevin Lin, John Hewitt, Ashwin Paranjape, Michele Bevilacqua, Fabio Petroni, and Percy Liang. 2024b. Lost in the middle: How language models use long contexts. *Transactions of the Association for Computational Linguistics*.
- Xiulong Liu, Zhikang Dong, and Peng Zhang. 2024c. Tackling data bias in music-avqa: Crafting a balanced dataset for unbiased question-answering. In *Winter Conference on Applications of Computer Vision*.
- Zheyuan Liu, Guangyao Dou, Eli Chien, Chunhui Zhang, Yijun Tian, and Ziwei Zhu. 2024d. Breaking the trilemma of privacy, utility, and efficiency via controllable machine unlearning. In *International World Wide Web Conference*.
- Zheyuan Liu, Guangyao Dou, Mengzhao Jia, Zhaoxuan Tan, Qingkai Zeng, Yongle Yuan, and Meng Jiang. 2024e. Protecting privacy in multimodal large language models with mllmu-bench. *arXiv preprint arXiv:2410.22108*.
- Muhammad Maaz, Hanoona Rasheed, Salman Khan, and Fahad Khan. 2024. Video-ChatGPT: Towards detailed video understanding via large vision and language models. In *Annual Meeting of the Association for Computational Linguistics*.
- Asmar Nadeem, Faegheh Sardari, Robert Dawes, Syed Sameed Husain, Adrian Hilton, and Armin Mustafa. 2024. Narrativebridge: Enhancing video captioning with causal-temporal narrative. *arXiv preprint arXiv:2406.06499*.
- Aaron van den Oord, Yazhe Li, and Oriol Vinyals. 2018. Representation learning with contrastive predictive coding. *arXiv preprint arXiv:1807.03748*.
- Shuhuai Ren, Sishuo Chen, Shicheng Li, Xu Sun, and Lu Hou. 2023. TESTA: Temporal-spatial token aggregation for long-form video-language understanding. In *Findings of the Association for Computational Linguistics: EMNLP*.
- Hugo Touvron, Thibaut Lavril, Gautier Izacard, Xavier Martinet, Marie-Anne Lachaux, Timothée Lacroix, Baptiste Rozière, Naman Goyal, Eric Hambro, Faisal Azhar, et al. 2023a. Llama 2: Open foundation and fine-tuned chat models. *arXiv preprint arXiv:2302.13971*.
- Hugo Touvron, Louis Martin, Kevin Stone, Peter Albert, Amjad Almahairi, Yasmine Babaei, Nikolay Bashlykov, Soumya Batra, Prajjwal Bhargava, Shrutu Bhosale, et al. 2023b. Llama 2: Open foundation and fine-tuned chat models. *arXiv preprint arXiv:2307.09288*.
- Jianfeng Wang, Zhengyuan Yang, Xiaowei Hu, Linjie Li, Kevin Lin, Zhe Gan, Zicheng Liu, Ce Liu, and Lijuan Wang. 2022. Git: A generative image-to-text transformer for vision and language. *Transactions on Machine Learning Research*.
- Jue Wang, Wentao Zhu, Pichao Wang, Xiang Yu, Linda Liu, Mohamed Omar, and Raffay Hamid. 2023. Selective structured state-spaces for long-form video understanding. In *Conference on Computer Vision and Pattern Recognition*.

Chao-Yuan Wu and Philipp Krahenbuhl. 2021. Towards long-form video understanding. In *Conference on Computer Vision and Pattern Recognition*.

Penghao Wu and Saining Xie. 2024. V*: Guided visual search as a core mechanism in multimodal llms. In *Conference on Computer Vision and Pattern Recognition*.

Wulin Xie, Xiaohuan Lu, Yadong Liu, Jiang Long, Bob Zhang, Shuping Zhao, and Jie Wen. 2024. Uncertainty-aware pseudo-labeling and dual graph driven network for incomplete multi-view multi-label classification. In *International Conference on Multimedia*.

Wulin Xie, Lian Zhao, Jiang Long, Xiaohuan Lu, and Bingyan Nie. 2025. Multi-view factorizing and disentangling: A novel framework for incomplete multi-view multi-label classification. *arXiv preprint arXiv:2501.06524*.

Jun Xu, Tao Mei, Ting Yao, and Yong Rui. 2016. Msr-vtt: A large video description dataset for bridging video and language. In *Conference on Computer Vision and Pattern Recognition*.

Jiawei Yao, Qi Qian, and Juhua Hu. 2024a. Customized multiple clustering via multi-modal subspace proxy learning. In *Advances in Neural Information Processing Systems*.

Jiawei Yao, Qi Qian, and Juhua Hu. 2024b. Multi-modal proxy learning towards personalized visual multiple clustering. In *Conference on Computer Vision and Pattern Recognition*.

Chunhui Zhang, Chao Huang, Youhuan Li, Xiangliang Zhang, Yanfang Ye, and Chuxu Zhang. 2022. Look twice as much as you say: Scene graph contrastive learning for self-supervised image caption generation. In *International Conference on Information & Knowledge Management*.

Chunhui Zhang, Yiren Jian, Zhongyu Ouyang, and Soroush Vosoughi. 2024. Scaling cognitive limits: Identifying working memory limits in llms. In *Conference on Empirical Methods in Natural Language Processing*.

Chunhui Zhang, Yiren Jian, Zhongyu Ouyang, and Soroush Vosoughi. 2025. Pretrained image-text models are secretly video captioners. In *Annual Conference of the North American Chapter of the Association for Computational Linguistics*.

Hang Zhang, Xin Li, and Lidong Bing. 2023. Video-llama: An instruction-tuned audio-visual language model for video understanding. In *Conference on Empirical Methods in Natural Language Processing*.

A Baselines

We evaluate nine MFMs reproduced with open-source code and pretrained weights:

- **LAVisH** (Lin et al., 2023): LAVisH adapter uses a small set of latent tokens, forming an attention bottleneck that reduces the quadratic cost of standard cross-attention. The trained LAVisH module alongside Swin-v2 serves as TWM visual encoder.
- **DG-SCT** (Duan et al., 2023): Adds cross-modal interaction layers to pretrained audio-visual encoders for adaptive extraction across spatial, channel, and temporal dimensions. All DG-SCT and Swin-T blocks serve as TWM visual encoder.
- **LAST-Att** (Liu et al., 2024c): Explores the inter-relationships between audio-visual-text modalities. Swin-v2 serves as TWM visual encoder.
- **Git** (Wang et al., 2022): Simplifies the architecture to a single image encoder and a text decoder under a unified language modeling task. The pretrained image encoder serves as TWM visual encoder.
- **AKGNN** (Hendria et al., 2023): Introduces a grid-based node representation, where nodes are represented by features extracted from a grid of video frames. The trained graph neural network from AKGNN serve as TWM visual encoder.
- **CEN** (Nadeem et al., 2024): Utilizes independent encoders to capture causal dynamics and generate time-sequenced captions. The pretrained CLIP-ViT from CEN serves as TWM visual encoder.
- **VINDLU** (Cheng et al., 2023): Develops a step-wise approach for efficient VidL pretraining. The trained video encoder from VINDLU serves as TWM visual encoder.
- **TESTA** (Ren et al., 2023): Compresses video semantics by adaptively aggregating similar frames and similar blocks within each frame. The full video encoder blocks from TESTA serve as TWM visual encoder.
- **MovieSeq** (Lin et al., 2024): Through instruction tuning, MovieSeq enables a language model to interact with videos using cross-modal instructions. CLIP vision encoder from MovieSeq serve as TWM visual encoder.




Figure 7: Ablation studies on different components of TWM using MUSIC-AVQA v2.0 dataset (Liu et al., 2024c). We evaluate the effectiveness of two key components in TWM: Visual Working Memory (VWM) which refines temporal features in visual modality, and Audio Working Memory (AWM) which refines temporal features in audio modality. TWM w/o VWM indicates TWM is only applied to auditory modality while visual features retain the baseline model’s distribution, and TWM w/o AWM indicates TWM is only applied to visual modality while maintaining the original auditory distribution.

B Ablation Studies

To thoroughly investigate the effectiveness of two key components in TWM and analyze their relative contributions to model performance, we conduct comprehensive ablation experiments on the MUSIC-AVQA v2.0 dataset (Liu et al., 2024c). Specifically, we examine two critical variants: (1) **TWM w/o VWM**, where TWM is exclusively applied to the auditory modality while retaining the baseline model’s original distribution for visual features, and (2) **TWM w/o AWM**, where TWM is solely applied to visual modality while maintaining the original auditory distribution.

As shown in Figure 7, the ablation studies reveal several notable findings. First, Visual Working Memory (VWM) demonstrates substantial impact on both visual perception and cross-modal understanding. Taking DG-SCT+TWM as an example, the removal of VWM leads to significant performance degradation in VQA (-3.15%) and AVQA (-8.39%) average scores, indicating VWM’s crucial role in capturing and refining temporal dependencies in visual features. Second, Auditory Working Memory (AWM) exhibits particularly strong effects on audio-related tasks, evidenced by a notable 4.67% decrease in AQA performance when removing AWM from LAVisH+TWM, while its impact on visual tasks remains relatively minimal (e.g., -0.35% in LAST-Att+TWM’s VQA performance). Most notably, the complete TWM framework (AWM+VWM) consistently achieves promising performance across all evaluation metrics, with particularly impressive gains in cross-modal scenarios. For example, DG-SCT+TWM shows substantial improvements over its ablated variants, achiev-

ing an 8.39% boost in AVQA average performance. These results empirically validate our hypothesis that AWM and VWM serve complementary functions in enhancing multimodal feature representations, which is essential for robust multimodal temporal reasoning.




Figure 8: TWM on video-text retrieval models using CMD dataset (Bain et al., 2020). Lighter bars indicate baseline models while darker bars represent baseline models enhanced with TWM.

To further validate the generalizability of these findings across different tasks and datasets, we conduct additional quantitative evaluations illustrated in Figure 8 and Figure 9. Figure 8 presents a systematic performance evaluation of video-text retrieval models on the CMD dataset, where the comparative analysis demonstrates consistent performance improvements across multiple retrieval metrics after incorporating TWM. The empirical results from Figure 9 further validate TWM’s contribution through a comprehensive assessment of video captioning quality on the MSR-VTT dataset, where the evaluation spans multiple standard cap-




Figure 9: TWM on video captioning models using MSR-VTT dataset (Xu et al., 2016). Lighter bars indicate baseline models while darker bars represent baseline models enhanced with TWM.

tioning metrics. These visualizations collectively substantiate TWM’s capability to enhance model performance across different architectures and tasks, with particularly pronounced improvements observed in certain evaluation criteria.

C Additional Case Studies

In addition to Section 4.3, we examine TWM’s integration with different state-of-the-art models through four detailed case studies: crowd scene understanding, indoor performance analysis, dynamic action sequences, and social interaction comprehension. These cases highlight TWM’s capability in frame selection, temporal coherence maintenance, and visual information preservation across diverse scenarios. TWM consistently achieves significant frame reduction while preserving essential visual information across different models and datasets. Its selection strategy ensures temporal coherence across selected frames, whether handling complex social scenes, dynamic actions, crowded environments, or indoor performances. This efficiency in frame selection, combined with maintenance of semantic continuity, highlights TWM’s effectiveness as a visual memory management system.

C.1 Crowd Scene Understanding

Figure 10 shows TWM’s integration with DG-SCT (Duan et al., 2023) on the MUSIC-AVQA v2.0 dataset (Liu et al., 2024c), demonstrating its capabilities in handling complex crowd scenes in a street performance setting. The frame selection reveals TWM’s systematic approach to managing dynamic crowd environments. In early frames, TWM establishes the overall scene context, capturing the spatial layout and crowd distribution. During the main sequence, it maintains tracking of individual

performers while preserving their spatial relationships within the crowd. This strategic frame selection enhances DG-SCT’s ability to understand complex crowd interactions while filtering out redundant or less informative frames.

C.2 Indoor Scene Understanding

Figure 11 demonstrates TWM’s effectiveness when integrated with LAVisH (Lin et al., 2023) on MUSIC-AVQA v2.0 (Liu et al., 2024c) for analyzing indoor performance scenes. Within this controlled environment, TWM exhibits precise frame selection that captures the key moments of performance. The selection strategy begins with establishing shots of the performance space, followed by carefully chosen frames that track the performer’s movements. TWM’s enhancement enables LAVisH to better handle indoor lighting conditions and maintain focus on the main subject despite varying camera angles and performer positions, demonstrating strong visual coherence in indoor settings.

C.3 Dynamic Action Sequences

Figure 12 illustrates TWM’s enhancement of TESTA (Ren et al., 2023) on the CMD dataset (Bain et al., 2020) for handling fast-paced action in "Kid’s bike spins out in the middle of his race with Dogg". The frame selection shows TWM’s adaptive sampling strategy for dynamic events. The mechanism increases the sampling density during crucial moments of the action, particularly during the spin-out sequence. It maintains broader context through strategic selection of frames before and after the key event, enabling TESTA to better balance between capturing fine-grained action details and maintaining overall scene comprehension.

C.4 Social Interaction Analysis

The case in Figure 13 showcases TWM’s integration with MovieSeq (Lin et al., 2024) on the CMD dataset (Bain et al., 2020) for capturing meaningful social interactions. Through the sequence "Walter introduces his employees to their new boss, The Beaver", TWM’s frame selection effectively traces the progression of this social event. The initial frames establish the gathering context, showing the assembly of employees. The middle sequence frames capture key moments of the introduction, while later frames document audience reactions and interaction outcomes. This selection pattern enhances MovieSeq’s ability to maintain narrative coherence through selective frame retention.

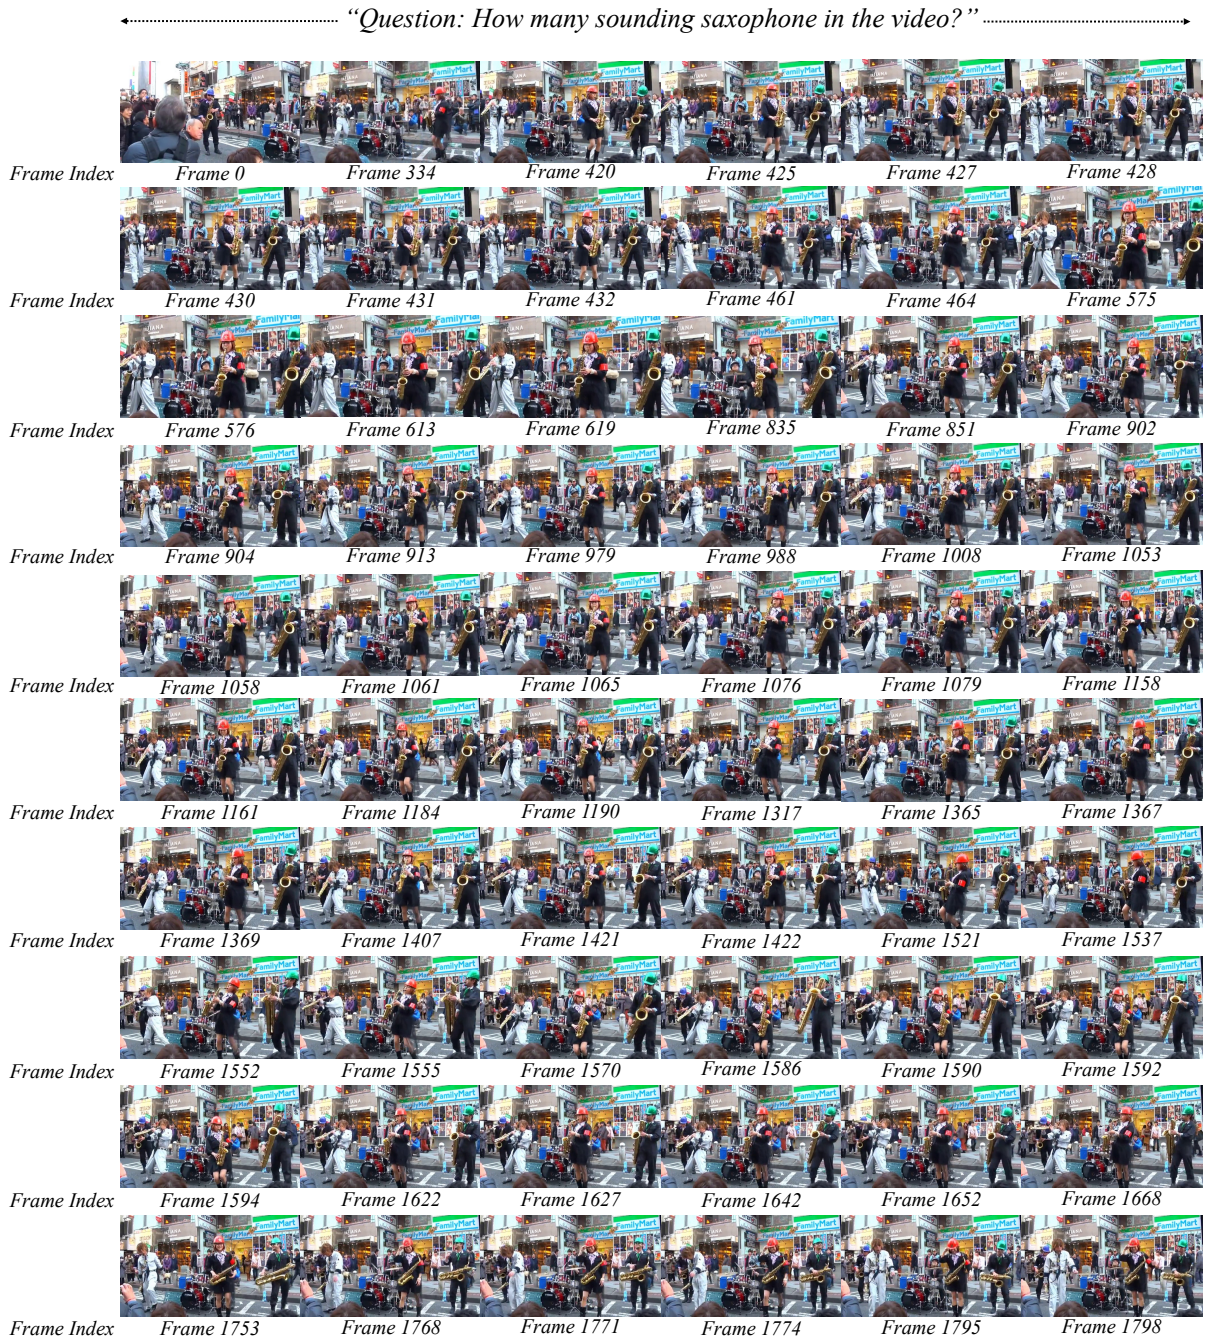


Figure 10: TWM-searched frames with corresponding questions for audio-visual question answering, integrated with DG-SCT (Duan et al., 2023) on the MUSIC-AVQA v2.0 dataset (Liu et al., 2024c). The visualization demonstrates TWM’s capability to search and select key frames that capture comprehensive audio-visual information necessary for answering the posed questions.



Figure 11: TWM-searched frames with corresponding questions for audio-visual question answering, integrated with LAVisH (Lin et al., 2023) on the MUSIC-AVQA v2.0 dataset (Liu et al., 2024c). This visualization showcases TWM’s capability in extracting frames that effectively capture the temporal-semantic information required to answer the given questions.

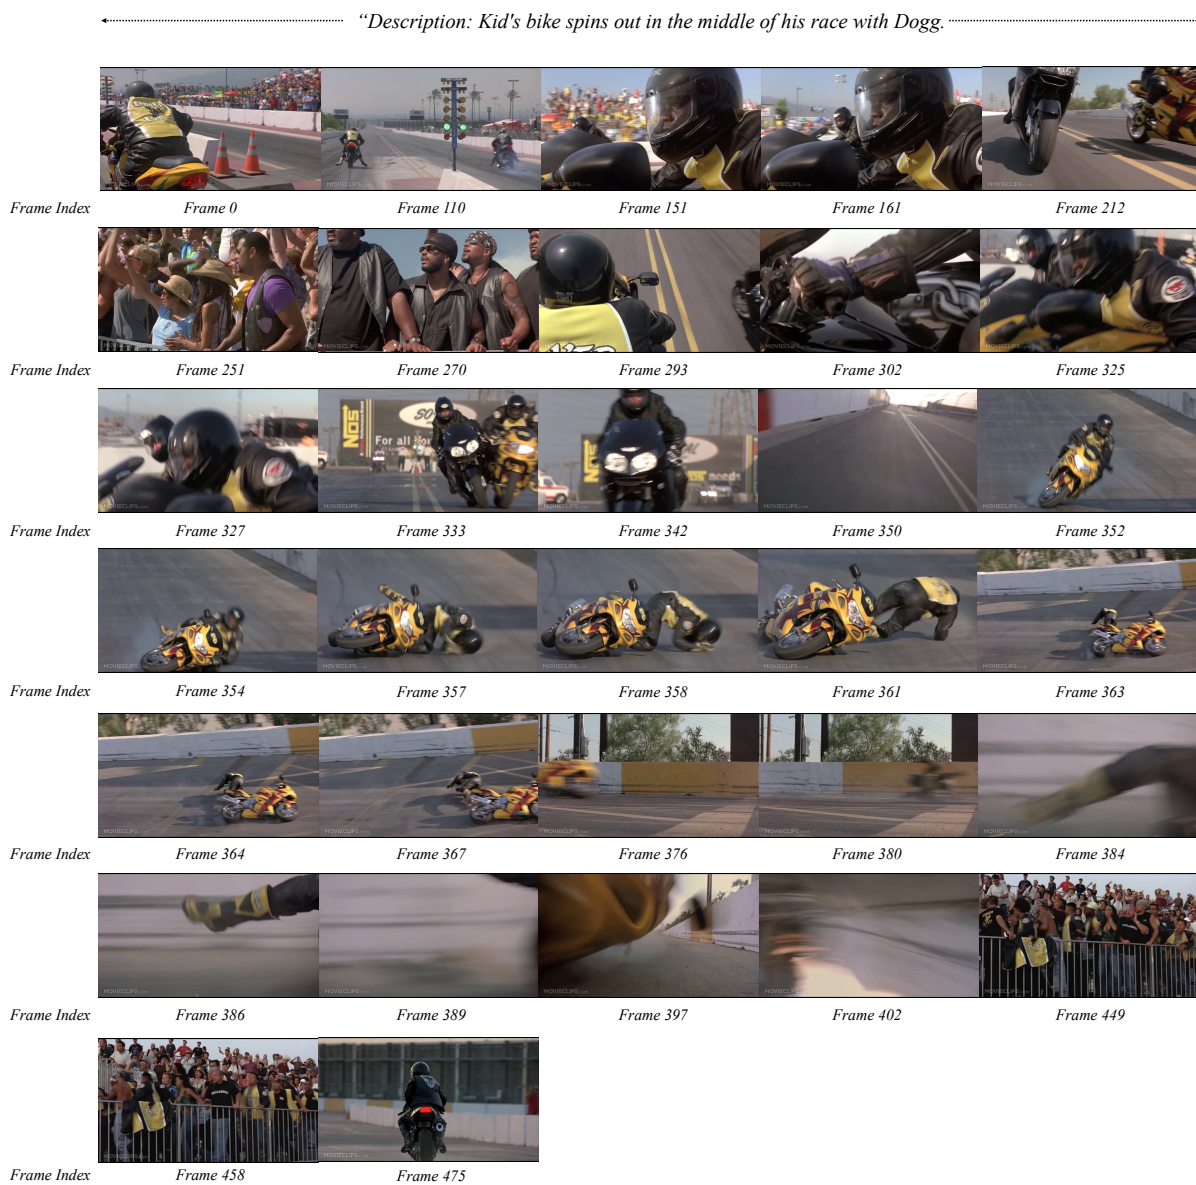


Figure 12: TWM-searched frames with corresponding text description for video-text retrieval, integrated with TESTA (Ren et al., 2023) on the CMD dataset (Bain et al., 2020). The visualization demonstrates TWM’s effectiveness in searching and selecting key frames that preserve the temporal-semantic information from the video sequence.



Figure 13: TWM-searched frames with corresponding text description for video-text retrieval, integrated with MovieSeq (Lin et al., 2024) on the CMD dataset (Bain et al., 2020). This example illustrates TWM’s capability to select visual frames that effectively capture the essential temporal dynamics and semantic information.