

Encoding and Controlling Global Semantics for Long-form Video Question Answering

Thong Thanh Nguyen^{1*} Zhiyuan Hu¹ Xiaobao Wu² Cong-Duy T Nguyen²
See-Kiong Ng¹ Anh Tuan Luu²

¹ National University of Singapore ² Nanyang Technological University

Abstract

Seeking answers effectively for long videos is essential to build video question answering (videoQA) systems. Previous methods adaptively select frames and regions from long videos to save computations. However, this fails to reason over the whole sequence of video, leading to sub-optimal performance. To address this problem, we introduce a state space layer (SSL) into multi-modal Transformer to efficiently integrate global semantics of the video, which mitigates the video information loss caused by frame and region selection modules. Our SSL includes a gating unit to enable controllability over the flow of global semantics into visual representations. To further enhance the controllability, we introduce a cross-modal compositional congruence (C³) objective to encourage global semantics aligned with the question. To rigorously evaluate long-form videoQA capacity, we construct two new benchmarks Ego-QA and MAD-QA featuring videos of considerably long length, *i.e.* 17.5 minutes and 1.9 hours, respectively. Extensive experiments demonstrate the superiority of our framework on these new as well as existing datasets. The code, model, and data have been made available at nguyenthong.github.io/Long_form_VideoQA.

1 Introduction

VideoQA has been extensively studied to develop systems to assist humans in daily activities (Grauman et al., 2022; Lei et al., 2021), *e.g.*, remind users of their past actions, help users locate their belongings, and provide assistance with complex tasks. To implement these functions, we expect videoQA systems to understand and extract relevant information from long-form videos with diverse objects and complex spatial-temporal interactions.

Compared with short clips, long-form videos pose more challenges for videoQA. They consist

of a higher number of objects and events. As such, comprehensively encoding information from them requires expensive computations. Moreover, a high amount of information may be unrelated to the posed question. To address these problems, recent studies (Bain et al., 2021; Wang et al., 2023; Gao et al., 2023) adaptively select a subset of video frames and visual regions associated with the question. Nevertheless, if a question necessitates a reasoning of the entire sequence of events (*e.g.* video 1’s Figure 1), or an understanding of the overall video narration (*e.g.* video 2’s Figure 1), a mere handful of selected frames or regions might not sufficiently encapsulate necessary details.

To tackle these problems, we introduce a state space layer (SSL). Before forwarding video frames to selection modules, SSL fixes long-term dependency patterns for integrating global information into visual representations. Such global information offers the selected frames the global context within the video, so that they can relate to other frames even though those frames are not selected for attention computation. However, a considerable amount of unrelated global information may flow into visual representations. Therefore, we first equip SSL with a gating mechanism to provide more controllability over the flow of global semantics into visual representations, resulting in our Gated State space Multi-modal Transformer (GSMT) architecture. Furthermore, we promote global semantics that is more aligned with the question. In particular, we introduce Cross-modal Compositional Congruence (C³) objective that compares visual attention with its version transitioned to the language basis via cross-modal attention, effectively measuring cross-modal congruence between intra-modal relations. Our rationale behind focusing on intra-modal relations is because videoQA models often need to understand spatial and temporal relationships between entities and events posed by the question (Gandhi et al., 2022),

*Corresponding to: Thong, e0998147@u.nus.edu

Video 1: 3037 IRON MAN2 - Video length: 2.08 hours



person's health deteriorates

replace the ingot of palladium

wearing the armored suit

activate the rt unit

Question: what if the person did not have the ingot of palladium to replace the smoking one in the rt unit?

Answer: The RT unit would degrade, potentially causing the person's health to deteriorate or the armored suit to malfunction.

Video 2: fb7cc35d-3272-44a4-b8f2-15cd24fa345b - Video length: 33.27 minutes



Question: what central theme best describes activities throughout the video?

Answer: routine household management and maintenance

Figure 1: Long-form videoQA examples, with videos taken from MAD (Soldan et al., 2022) and Ego4D (Grauman et al., 2022) datasets, respectively. Question in video 1 requires the model to reason about the relation chain of replacing *ingot of palladium* to activate the *rt unit* that powers the *armored suit* and protects *person's health*. Question in video 2 necessitates an understanding of the overall theme in video 2.

thus we encourage globally informed visual representations to maintain compositional consistency between visual patches and question entities.

Remarkably, we observe that recent long-form videoQA works (Gao et al., 2023; Islam et al., 2024) still mostly evaluate on videos lasting at most one minute or two, and use short-natured questions which necessitate watching only a short period of video, *i.e.*, about 100 seconds (Mangalam et al., 2023), to determine the answer. To more rigorously evaluate long-form videoQA capacity, we introduce a construction procedure which utilizes large language model (LLM) to generate questions and associated answers for egocentric and movie videos whose average lengths are 17.5 minutes and 1.9 hours, respectively. Additionally, we also conduct automatic and manual filtering to obtain high-quality questions which require watching a video up to 1200 seconds to answer, longer than any existing long-form videoQA benchmarks (Xiao et al., 2021; Wu et al., 2021; Mangalam et al., 2023).

To sum up, our contributions are as follows:

- We propose a Gated State space Multi-modal Transformer (GSMT) with state space layer (SSL) to integrate global information into visual representations for long-form videoQA.
- We equip SSL with a gating mechanism to provide controllability over the flow of global video semantics and a Cross-modal Compositional Congruence (C^3) objective to encourage question-aligned visual representations.

- We curate two new datasets with excessively long video lengths and long-natured questions for long-form videoQA. Comprehensive experiments on our curated and five standard datasets substantiate the superiority of our framework over various competitive baselines.

2 Related Work

Video question answering (videoQA). VideoQA datasets (Xu et al., 2017; Jang et al., 2017; Tapaswi et al., 2016; Lei et al., 2018; Nguyen et al., 2024c) mostly focus on short video clips about daily activities, *e.g.* sports or household work. Recent works (Gao et al., 2021; Grunde-McLaughlin et al., 2021; Wu et al., 2021) focus on complex spatial-temporal reasoning ability over longer temporal lengths with causal and transition questions. Regarding methodology, early works (Zhao et al., 2018; Li et al., 2019) present LSTM-based and GNN-based architectures to capture cross-modal and motion-appearance interaction. For example, (Xiao et al., 2023) incorporate graph modeling into Transformer to explicitly encode the object relations in videos. Due to outstanding performance of pre-trained vision-language Transformers, Bain et al. (2021); Fu et al. (2023); Wang et al. (2023) utilize pre-trained Transformer models on downstream videoQA tasks.

Long-form video modeling. To improve the applicability of vision-language systems, various works have focused on long-form video modeling for

videoQA tasks. Wu and Krahenbuhl (2021) introduce short-term feature extraction technique and long-term memory mechanism to alleviate redundant video frame processing. Lin et al. (2022b) propose to compensate sparsely extracted video frames with audio cues. Islam and Bertasius (2022) design structured multi-scale temporal decoder. Gao et al. (2023) utilize question as a guide to select question-related visual segments to mitigate computation. Inspired by (Nguyen et al., 2021; Wu et al., 2024b) in language modeling domains, we introduce a method to encode global semantics of videos to tackle information degradation problem in long-form video modeling.

Cross-modal alignment. Alignment approaches, such as contrastive learning (Nguyen and Luu, 2021; Nguyen et al., 2024e; Wu et al., 2023a, 2024a), optimal transport (Nguyen and Luu, 2022; Wu et al., 2023b, 2024d,c), tree-based fusion (Nguyen et al., 2023c), and energy-based modeling (Nguyen et al., 2023d), have been investigated to compare different input samples. For multimodal learning, recent cross-modal alignment methods focus on establishing a shared latent space in which samples of different modalities can be compared readily (Nguyen et al., 2023b, 2024b,d,a, 2023a, 2022b,a). For example, with a contrastive learning formulation, CLIP (Radford et al., 2021) and ALIGN (Jia et al., 2021) learn generalizable image-text representations from millions of image-text pairs. To exploit relation for cross-modal alignment, Ren et al. (2021) use hard bijective correspondences between words and objects via an argmax over the cross-modal attention matrix to optimize the cross-modal alignment. Pandey et al. (2022) generalize the hard alignment with their soft alignment approach, but they focus on objects within an image rather than a long video.

3 Methodology

The formulation of video question answering (videoQA) is to predict the answer y for a question q about a video \mathcal{V} as follows:

$$\tilde{y} = \arg \max_{y \in \mathcal{A}} \mathcal{F}_\theta(y|q, \mathcal{V}, \mathcal{A}), \quad (1)$$

where \tilde{y} denotes the predicted answer chosen from the set of candidate options \mathcal{A} , and θ denotes the trainable parameters of a videoQA model. With this videoQA task formulation, we explain our proposed GSMT architecture as the model \mathcal{F}_θ and \mathcal{C}^3

objective to support GSMT. The overall framework is illustrated in Figure 2.

3.1 Gated State Space Multi-Modal Transformer (GSMT)

Our GSMT takes videos and questions as input, divides each video frame into visual patches and a question into textual words, and then forwards visual patches and textual words into video and text embedder to extract initial representations.

Input Embedder. For video embedder, we utilize frozen pre-trained vision-language Transformer to extract patch-level features $X = \{\mathbf{x}_0, \mathbf{x}_1, \dots, \mathbf{x}_{L-1}\}$, of all T frames, where t -th frame consists of N patch-level features $\{\mathbf{x}_{t,j}\}_{j=0}^{N-1}$, hence $L = NT$, and $\mathbf{x} \in \mathbb{R}^d$. For text embedder, a similar frozen pre-trained vision-language Transformer is used to obtain word-level features $W = \{\mathbf{w}_0, \mathbf{w}_1, \dots, \mathbf{w}_{M-1}\}$, where \mathbf{w}_0 corresponds to the [CLS] token and $\mathbf{w}_1, \dots, \mathbf{w}_{M-1}$ correspond to words in the question.

Gated State Space Layer (Gated SSL). Inspired by (Gu et al., 2021), we define a sequence-to-sequence map $SSL(X)$ from a sequence of patch-level features to d_S -dim hidden states, parameterized by learnable state matrices $A \in \mathbb{R}^{d_S \times d_S}$, $B \in \mathbb{R}^{d_S \times d}$, $C \in \mathbb{R}^{d \times d_S}$, and step size Δ as:

$$\mathbf{g}_{t+1} = \bar{A} \cdot \mathbf{g}_t + \bar{B} \cdot \mathbf{x}_{t+1}, \quad \mathbf{o}_{t+1} = \bar{C} \cdot \mathbf{g}_{t+1}, \quad (2)$$

$$\bar{A} = e^{A\Delta}, \quad \bar{B} = (\bar{A} - I)A^{-1}B, \quad \bar{C} = C. \quad (3)$$

We unroll the mapping to obtain:

$$\mathbf{o}_t = \sum_{j=0}^t \bar{C} \bar{A}^j \bar{B} \cdot \mathbf{x}_{t-j}. \quad (4)$$

This can be written as a convolutional representation $O = \bar{\Gamma} * X$, where $\bar{\Gamma} = (\bar{C}\bar{B}, \bar{C}\bar{A}\bar{B}, \dots, \bar{C}\bar{A}^{L-1}\bar{B})$ denotes the convolutional kernel, $*$ the discrete convolution operator, X the input sequence, O the corresponding output sequence. This convolution denotes the fixed global dependency pattern that facilitates the computation of global information among visual patches. We use Fast Fourier Transformer (FFT) (Cooley and Tukey, 1965) to compute the convolution in parallel provided that $\bar{\Gamma}$ has been obtained.

Computing the kernel $\bar{\Gamma}$ is non-trivial since it requires L distinct matrix powers. Instead, inspired by (Gupta et al., 2022), we initialize A to be a

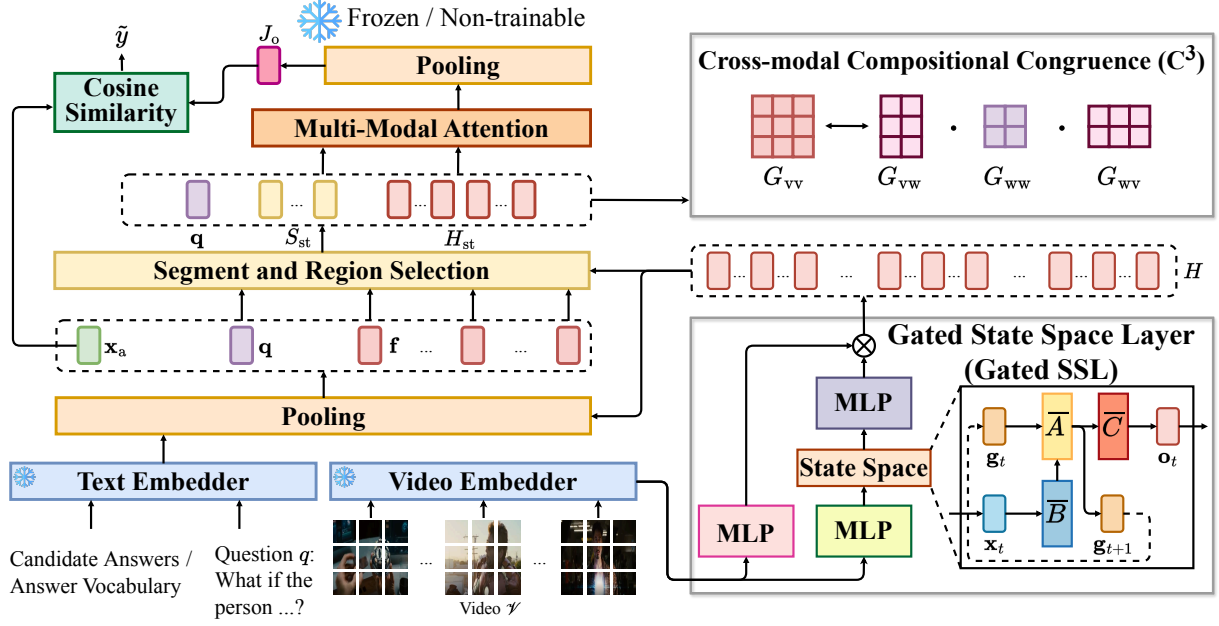


Figure 2: Illustration of the GSMT architecture empowered by gated SSL and C^3 training objective.

diagonal matrix $\text{diag}(\lambda_1, \lambda_2, \dots, \lambda_{d_S})$ and B to all-one matrix $\mathbb{1}_{d_S \times d}$. Due to this initialization, the kernel can be computed as:

$$\bar{\Gamma} = (C \odot E) \cdot P, \quad (5)$$

where $E = \left(\frac{e^{\lambda_1 \Delta - 1}}{\lambda_1}, \frac{e^{\lambda_2 \Delta - 1}}{\lambda_2}, \dots, \frac{e^{\lambda_{d_S} \Delta - 1}}{\lambda_{d_S}} \right) \in \mathbb{R}^{d_S}$, in which $P_{i,j} = e^{\lambda_i \cdot j \cdot \Delta} \in \mathbb{R}^{d_S \times L}$, and \odot denotes the element-wise multiplication.

To equip SSL with the control over which global semantics to integrate into visual representations, we construct a gating unit as:

$$U = \phi(\text{Linear}(X)), \quad V = \phi(\text{Linear}(X)), \quad (6)$$

$$O = \text{Linear}(\text{SSM}(U)), \quad (7)$$

$$H = \text{Linear}(O \odot V), \quad (8)$$

where $U, V, O \in \mathbb{R}^{L \times d_{\text{gating}}}$ denote the intermediate gating representations, $H \in \mathbb{R}^{L \times d_h}$ the visual representations output by gated SSL, ϕ the non-linear activation, and $d_{\text{gating}} < d, d_h, d_s$.

Visual Segment and Region Selection. After obtaining visual patch-level hidden representations H , we proceed to obtain frame features by pooling each frame t 's visual patches:

$$\mathbf{f}_t = \text{pool}(\mathbf{h}_{t,0}, \mathbf{h}_{t,1}, \dots, \mathbf{h}_{t,N-1}). \quad (9)$$

Subsequently, we group non-overlapping consecutive frames into a segment to obtain I segments, each of which contains $N_p = \lceil \frac{L}{T} \rceil$ patches and

$N_f = \lceil \frac{T}{I} \rceil$ frames. We proceed to compute segment features through pooling frame features corresponding to the segment:

$$\mathbf{s}_i = \text{pool}(\mathbf{f}_{i,0}, \mathbf{f}_{i,1}, \dots, \mathbf{f}_{i,N_f-1}). \quad (10)$$

Similarly, we also pool word features to obtain the question representation:

$$\mathbf{q} = \text{pool}(\mathbf{w}_0, \mathbf{w}_1, \dots, \mathbf{w}_{M-1}). \quad (11)$$

Given the segment features $S = \{\mathbf{s}_i\}_{i=0}^{I-1}$ and question feature \mathbf{q} , we conduct top- k segment selection:

$$Q = \text{Linear}(\mathbf{q}), \quad K = \text{Linear}(S), \quad (12)$$

$$B = \text{selector}_{\text{Top}_k} \left(\text{softmax} \left(\frac{QK^\top}{\sqrt{d_k}} \right) \right). \quad (13)$$

We implement *selector* as a differentiable Gumbel-softmax selection function. The output of the *selector* is a sequence of segment index \mathcal{B} . Thereby, we extract respective segment features with respect to the selected segment indices, *i.e.* $S_{\text{st}} = \{\mathbf{s}_b \mid b \in \mathcal{B}\}$.

For every frame τ in the selected segments, we then perform cross-modal attention between its patch-level hidden representations $H_\tau = \{\mathbf{h}_{\tau,j}\}_{j=0}^{N-1}$, $\tau \in \{\lfloor \frac{b}{N_f} \rfloor \mid b \in \mathcal{B}\}$ and the question to select top- j question-related patches:

$$Q = \text{Linear}(\mathbf{q}), \quad K = \text{Linear}(H_\tau), \quad (14)$$

$$\mathcal{T} = \text{selector}_{\text{Top}_j} \left(\text{softmax} \left(\frac{QK^\top}{\sqrt{d_k}} \right) \right). \quad (15)$$

Lastly, we stack the selected patches of all selected frames to obtain $H_{\text{st}} = \{\mathbf{h}_{\tau,j} | \tau \in \mathcal{T}\}_{j=0}^{N-1}$.

Multi-Modal Attention. At present, we employ self-attention to produce multi-modal hidden representations that fuse the information of question and video. In particular, we concatenate the question word-level features $Q = \{\mathbf{w}_i\}_{i=0}^{M-1}$, selected segment features $S_{\text{st}} = \{\mathbf{s}_b | b \in \mathcal{B}\}$, and selected patch features $H_{\text{st}} = \{\mathbf{h}_{\tau,j} | \tau \in \mathcal{T}\}_{j=0}^{N-1}$:

$$J = [\text{Linear}(S_{\text{st}}), \text{Linear}(H_{\text{st}}), \text{Linear}(\mathbf{q})], \quad (16)$$

where $[\cdot]$ denotes the concatenation operator. Thereafter, we iteratively conduct self-attention over the concatenated features for N_L layers to accomplish multi-modal contextual representations:

$$J^{(i+1)} = \text{SelfAttention}\left(J^{(i)}\right), \quad 0 \leq i \leq N_L - 1. \quad (17)$$

Answer Prediction. Afterwards, we pool the features of all multi-modal self-attention layers:

$$J_o = \text{pool}\left(J^{(0)}, J^{(1)}, \dots, J^{(N_L-1)}\right). \quad (18)$$

Then, we calculate the cosine similarity between J_o and the feature of all candidate answers $X_{\text{ans}} = \{\mathbf{x}_a | a \in \mathcal{A}\}$ obtained by utilizing the pre-trained model. We choose the candidate answer of the highest similarity as the final prediction \tilde{y} :

$$\tilde{y} = \arg \max_{y \in \mathcal{A}} \left(J_o \cdot (X_{\text{ans}})^\top \right). \quad (19)$$

3.2 Cross-modal Compositional Congruence (\mathcal{C}^3) Objective

Based on Eq. (17), we denote the output representations of visual patches as J_v and those of question words as J_w . We calculate cross-modal attention:

$$G_{\text{vw}} = J_v \cdot (J_w)^\top, \quad G_{\text{ww}} = J_w \cdot (J_v)^\top, \quad (20)$$

and intra-modal visual and textual attention as:

$$G_{\text{vv}} = J_v (J_v)^\top, \quad G_{\text{ww}} = J_w (J_w)^\top. \quad (21)$$

Given the intra-modal and cross-modal attention, we perform the change of basis to compute the intra-modal visual attention in the language space:

$$R_{\text{vv}} = G_{\text{vw}} G_{\text{ww}} (G_{\text{vw}})^\top. \quad (22)$$

As such, we define the loss objective to align the original G_{vv} with the change-of-basis version R_{vv} :

$$\mathcal{L}_{\mathcal{C}^3} = \text{m-KL}(\text{softmax}(R_{\text{vv}}), \text{softmax}(G_{\text{vv}})), \quad (23)$$

where m-KL denotes the symmetric Kullback-Leibler Divergence (KL) between R and G :

$$\text{m-KL}(R, G) = \text{KL}(R||G) + \text{KL}(G||R). \quad (24)$$

3.3 Overall Objective

We jointly optimize the softmax cross-entropy loss \mathcal{L}_{CE} between the log-likelihood of the model prediction \tilde{y} and the groundtruth answer y , with our proposed cross-modal alignment objective $\mathcal{L}_{\mathcal{C}^3}$:

$$\mathcal{L} = \mathcal{L}_{\text{CE}} + \gamma \cdot \mathcal{L}_{\mathcal{C}^3}, \quad (25)$$

where γ denotes the hyperparameter to control the effect of the cross-modal alignment objective.

4 Ego-QA and MAD-QA Benchmarks

To more rigorously evaluate long-form videoQA models, we construct two new datasets, *i.e.* Ego-QA and MAD-QA.

4.1 Ego-QA

We inherit 3k hours of 8640 egocentric videos from the Ego4D dataset (Grauman et al., 2022). Each video is associated with about 280 dense captions of consecutive moments. Based on these captions, we create our dataset in 2 stages, *i.e.* question-answer generation and data filtering.

Question-answer generation. In this first stage, we concatenate a video’s dense captions following the time order to construct its language description. We utilize GPT-4 (Achiam et al., 2023) to generate 20 questions per video. In our prompt, we encourage GPT-4 to avoid questions that are visually biased and can be answered by a short video moment. Then, we present the generated questions to GPT-4 to generate the correct answer along with 4 wrong answer choices.

Data filtering. In the second stage, we filter out questions that include clue words, *e.g.* “*passage*”, “*text*”, and “*description*”. Moreover, we also remove questions that GPT-4 can answer without looking at the concatenated narration or the question. Then, we adopt manual filtering by asking ten graduate students who are native English speakers to ensure the veracity and temporal certificate length for every question-answer sample. Particularly, annotators are instructed to verify that 1) questions are valid and the correct answer is indeed correct, 2) all distractor answers are incorrect, and 3) the video length to watch to determine the correct answer is at least 2 minutes.

The filtering stage reduces the number of admissible questions by a factor of $4\times$ to $5\times$. We accomplish 18.8K questions for 992 videos, which we split into 80% train, 10% val, and 10% test.

4.2 MAD-QA

We follow the same process for Ego-QA to obtain MAD-QA by utilizing 1.2K hours of 650 videos from the MAD dataset (Soldan et al., 2022). Since video lengths and the number of dense captions are larger than Ego-QA, we ask GPT-4 to generate 60 instead of 20 questions per video. Because GPT-4 might store external knowledge about the movie, we replace the name of characters in the caption with *person_1*, *person_2*, etc. Afterwards, we obtain 15.7K questions for 650 videos, and we split them into 80% train, 10% val, and 10% test.

The average video lengths in Ego-QA and MAD-QA are 17.5 minutes and 1.9 hours, respectively. Moreover, the average necessary video lengths humans need to watch to determine the answer for the two datasets are respectively 1204.4 and 396.07 seconds, longer than the average 100-second length of the recent very long-form videoQA dataset EgoSchema (Mangalam et al., 2023). We visualize the statistics of our datasets in Appendix E, and language prompts to generate question-answer samples in Appendix A, along with the instructions for human annotators in Appendix B.

5 Experiments

5.1 Standard Benchmarks

In addition to our constructed Ego-QA and MAD-QA datasets, we follow previous works (Gao et al., 2023; Xiao et al., 2023; Wang et al., 2023) to evaluate GSMT on four additional publicly available datasets for long-form videoQA: AGQA (Grunde-McLaughlin et al., 2021), NExT-QA (Xiao et al., 2021), STAR (Wu et al., 2021), Env-QA (Gao et al., 2021), and EgoSchema (Mangalam et al., 2023).

AGQA (Grunde-McLaughlin et al., 2021) is a videoQA dataset for compositional spatio-temporal reasoning. As recommended by the dataset creator, we employ its v2 version, which possesses more balanced distributions. AGQA consists of 2.27M QA pairs for 9.7K videos.

NExT-QA (Xiao et al., 2021) focuses on causal and temporal reasoning. The dataset comprises 5,440 videos associated with 52K questions.

STAR (Wu et al., 2021) concentrates on situated reasoning questions. The dataset provides 60K questions related to 22K videos clips.

Env-QA (Gao et al., 2021) is curated for dynamic environment understanding. Env-QA contains 23K egocentric videos collected on virtual environment

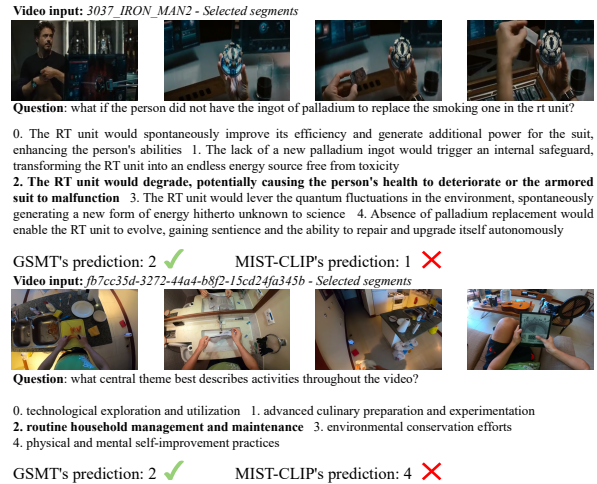


Figure 3: Qualitative results on the constructed MAD-QA and Ego-QA datasets.




Figure 4: GPU memory cost with respect to visual sequence length L of attention, convolution, and gated SSL mechanism.

AI2THOR (Kolve et al., 2017), which are used to generate 85K questions.

EgoSchema (Mangalam et al., 2023) consists of egocentric videos of 3-minute length. Questions in EgoSchema require humans on average 100 seconds watching the video to answer.

5.2 Implementation Details

Our framework can be implemented on most multi-modal Transformers. To fairly compare with previous works (Wang et al., 2023; Gao et al., 2023), we evaluate upon two popular types of pre-trained models, *i.e.* CLIP (ViT/B-32) (Radford et al., 2021) and All-in-One-Base (Wang et al., 2023). We divide a video frame into 4×4 patches and send them to video embedder. For our gated SSL, we use $d_S = d = d_h = 512$, $d_{\text{gating}} = 128$. For selection modules, we use $\text{top-}k = 4$, $\text{top-}j = 12$. For the multimodal attention module, we use $N_L = 2$. Based on validation, we employ max-pooling for all pooling operations, ReLU activation for ϕ , and $N_L = 2$.

Question Types	Object-relation	Relation-action	Object-action	Superlative	Sequencing	Exists	Duration comparison	Activity recognition	All
AIO	48.34	48.99	49.66	37.53	49.61	50.81	45.36	18.97	48.59
ATP	50.15	49.76	46.25	39.78	48.25	51.79	49.59	18.96	49.79
MIST-AIO	51.43	54.67	55.37	41.34	53.14	53.49	47.48	20.18	50.96
MIST-CLIP	51.68	67.18	68.99	42.05	67.24	60.33	54.62	19.69	54.39
GSMT-AIO	53.67	56.10	56.61	43.44	53.84	56.26	49.83	21.73	52.61
GSMT-CLIP	53.94	69.84	72.53	44.19	69.12	61.45	57.32	21.20	56.16

Table 1: Results of videoQA on AGQA-v2.

Method	Attribute	State	Event	Order	Number	All
STAGE	39.49	49.93	34.52	55.32	38.54	41.97
AIO	41.78	52.98	37.57	55.16	38.50	44.86
ATP	42.87	53.49	38.35	55.25	38.65	45.43
MIST-AIO	43.63	55.17	40.99	55.44	39.54	47.19
MIST-CLIP	44.05	58.13	42.54	56.83	40.32	48.97
GSMT-AIO	48.76	57.99	44.96	57.39	43.18	50.81
GSMT-CLIP	49.23	61.10	46.66	58.83	44.03	52.73

Table 2: Results of videoQA on Env-QA.

Method	Interaction	Sequence	Prediction	Feasibility	Mean
CLIP	39.80	40.50	35.50	36.00	38.00
RESERVE-B	44.80	42.40	38.80	36.20	40.50
Flamingo-9B	-	-	-	-	43.40
AIO	47.53	50.81	47.75	44.08	47.54
ATP	50.63	52.87	49.36	40.61	48.37
CoVGT	-	-	-	-	46.23
MIST-AIO	53.00	52.37	49.52	43.87	49.69
MIST-CLIP	55.59	54.23	54.24	44.48	51.13
GSMT-AIO	56.59	55.55	52.23	46.04	51.36
GSMT-CLIP	59.36	57.52	57.21	46.68	52.85

Table 3: Results of videoQA on STAR.

We apply \mathcal{L}_{C^3} upon $J^{(2)}$, and observe no difference between applying on $J^{(1)}$ and $J^{(2)}$. For fair comparison on AGQA, NExT-QA, STAR, Env-QA, and EgoSchema datasets, we sample 32 frames per video, and split them into $K = 8$ segments. Since video lengths are longer in our EgoQA and MAD-QA datasets, we sample 128 and 8192 frames per video, respectively, and split into $K = 8$ segments. For language modality, we embed the question with the same pre-trained model as the video embedder, and embed the answer with the pre-trained BERT-base model. We apply $\lambda = 0.005$ to balance the scale of \mathcal{L}_{C^3} and \mathcal{L}_{CE} .

5.3 Baselines

We compare our proposed framework against the following baseline:

- **HQGA** (Xiao et al., 2022): models video as a conditional graph hierarchy which aggregates visual facts in a level-wise manner with the guidance of textual modality,
- **CoVGT** (Xiao et al., 2023): a model which exploits a graph transformer to encode video by capturing visual objects and their relations for spatio-temporal reasoning, along with con-

Method	Causal	Temporal	Descriptive	All
HQGA	48.48	51.24	61.65	51.42
CLIP	46.30	39.00	53.10	43.70
VQA-T	49.60	51.49	63.19	52.32
AIO	48.04	48.63	63.24	50.60
ATP	53.10	50.20	66.80	54.30
CoVGT	59.69	58.00	69.88	60.73
MIST-AIO	51.54	51.63	64.16	53.54
MIST-CLIP	54.62	56.64	66.92	57.18
GSMT-AIO	59.72	59.04	69.91	60.76
GSMT-CLIP	60.87	61.16	70.26	62.49

Table 4: Results of videoQA on NExT-QA.

Method	Accuracy
EgoVLP	34.86
VideoReCap	50.23
MIST-AIO	56.27
MIST-CLIP	56.42
GSMT-AIO	58.28
GSMT-CLIP	58.55

Table 5: Results of videoQA on EgoSchema.

trastive learning to align its graph transformer with the textual encoder,

- **ATP** (Buch et al., 2022): a model which learns to select one salient video frame to perform the videoQA task,
- **AIO** (Wang et al., 2023): an end-to-end videoQA model that releases the need of unimodal encoders and leverages non-parametric token rolling operation,
- **STAGE** (Lei et al., 2020): a model that jointly carries out videoQA with moment localization and object grounding,
- **VQA-T** (Yang et al., 2021): a cross-modal Transformer which is pre-trained with supervised contrastive learning on a large-scale videoQA dataset,
- **RESERVE-B** (Zellers et al., 2022): jointly represents video frames, texts, and audio, and learns to predict the masked textual and audio tokens given video frames,
- **Flamingo-9B** (Alayrac et al., 2022): a model

Method	Ego-QA	MAD-QA
AIO	24.19	16.14
CoVGT	26.72	15.71
MIST-AIO	27.71	14.19
MIST-CLIP	29.73	17.15
GSMT-AIO	28.72	15.69
GSMT-CLIP	32.40	19.11
Human	80.29	73.21

Table 6: Results on constructed Ego-QA and MAD-QA.

Method	NExT-QA	STAR	Ego-QA	MAD-QA
SSL	61.91	52.40	31.30	18.95
Non-diag SSL	60.44	51.63	30.51	18.86
Attention	59.74	51.07	30.06	18.41
Convolution	59.34	50.85	30.17	18.33
Gated SSL	62.49	52.85	32.40	19.11

Table 7: Ablation results of gated SSL.

which is able to process multi-modal prompt and cast videoQA as text prediction,

- **EgoVLP** (Lin et al., 2022a): pre-trains multi-modal Transformer on egocentric videos for egocentric videoQA.
- **VideoReCap** (Islam et al., 2024): a recursive model that synergizes different video hierarchies to process hour-long videos,
- **MIST** (Gao et al., 2023): a multi-modal Transformer which decomposes video into frames, patches, and segments to efficiently process with its self-attention mechanism.

5.4 Quantitative Results

We show our experiments on AGQA-v2, Env-QA, STAR, NExT-QA, and EgoSchema in Table 1, 2, 3, 4, and 5, respectively, while revealing results on our constructed Ego-QA and MAD-QA in Table 6. We can observe that our method achieves superior performance over the latest methods on all datasets. In terms of the overall accuracy, we outperform the second-best method on AGQA-v2, Env-QA, STAR, and EgoSchema, *i.e.* MIST-CLIP, by 1.77%, 3.76%, 1.72%, and 2.13%, respectively. Equivalently, we outperform CoVGT, which is the second-best method on NExT-QA, by 1.20%.

Inspecting more closely, we note that our framework obtains more significant performance increase on questions that require the capacity of reasoning among visual concepts, *i.e.* improving 2.66% and 3.54% respectively for *relation-action* and *object-action* on AGQA-v2, 6.25% and 4.52% respectively for causal and temporal on NExT-QA,

Method	NExT-QA	STAR	Ego-QA	MAD-QA
GSMT	60.78	49.67	29.73	18.00
GSMT w/ OT	60.89	50.30	30.14	18.22
GSMT w/ POT	60.93	50.59	30.55	18.26
GSMT w/ C ³	62.49	52.85	32.40	19.11

Table 8: Ablation results of cross-modal alignment.

Dataset / d_{gating}	32	64	128	256	512	No gating
NExT-QA	60.74	61.09	62.49	61.90	61.33	60.61
STAR	51.82	52.63	52.85	52.70	51.86	51.27
Ego-QA	29.96	31.37	32.40	31.55	31.25	29.88

Table 9: Effect of various d_{gating} dimensions and no gating on NExT-QA, STAR, and Ego-QA datasets.

Method	NExT-QA	STAR	Ego-QA	MAD-QA
Multi-modal SSL	60.69	50.87	29.74	18.61
GSMT	62.49	52.85	32.40	19.11

Table 10: Effect of the position of SSL.

than those that require the ability to extract information within one frame, *i.e.* improving 2.26% for *object-relation* on AGQA-v2 and 3.34% for *descriptive* on NExT-QA. These results demonstrate our global semantics signal can address the challenging long-range temporal reasoning problems of long-form videoQA.

Remarkably, existing methods demonstrate significantly low performance on our curated datasets. For example, MIST-CLIP only achieves 29.73% on Ego-QA, and 17.15% accuracy on MAD-QA, which is less than random chance. In contrast, humans obtain 80.29% and 73.21% accuracy on Ego-QA and MAD-QA, respectively. These results suggest that previous methods might not encompass sufficient information in their selected segments and visual regions. Conversely, with the integrated global information, our framework can enhance videoQA performance on these challenging datasets. However, the accuracy remains substantially below human performance. Future research should focus more on genuine long-form videoQA, where videos can extend to several hours.

5.5 Ablation Study

Gated SSL implementation. We explore the effect of our gated SSL in Table 7 on NExT-QA, STAR, Ego-QA, and MAD-QA datasets. As can be observed, removing the gating unit, *i.e.* the SSL approach, results in performance drops, since redundant and noisy information might be passed to the visual representations. Additionally, not initializing state space parameters as diagonal matrices, *i.e.* non-diag SSL, does not remarkably impact

the performance. However, the time and memory complexity would become $O(L^2)$, which is significantly more costly than our initialization approach.

Choices for Global Semantics. We compare gated SSL with other choices to extract global semantics among visual elements, *i.e.* self-attention and convolution, in terms of videoQA performance in Table 7 and GPU memory cost in Figure 4. As can be observed, our gated SSL not only brings less computational cost than self-attention but also higher accuracy, validating the effectiveness of its global information signal. Moreover, whereas convolution pays equivalent computational cost to our gated SSL, its local pattern does not provide productive contextual information among visual elements, resulting in lower accuracy than gated SSL.

Effect of Gating Unit. We ablate the gating unit and vary the gating dimension d_{gating} in our gated SSL. As shown in Table 9, increasing the gating dimension leads to higher videoQA accuracy, as model has more controllability towards global information into visual representations. However, when the gating dimension becomes larger, the performance saturates and deteriorates. We posit that the model might become more constrained to allow the encoding of global semantics, degenerating to the architecture with limited global semantics. Apparently, removing gating unit results in performance drops, because irrelevant global information for the question could flow into visual hidden states without model controllability.

Effect of C^3 objective. We compare our C^3 alignment objective with alternative approaches, *i.e.* optimal transport (OT) (Pramanick et al., 2022) and its partial variant (POT) (Chapel et al., 2020). As shown in Table 8, both OT and POT can polish videoQA performance, while C^3 yields the highest performance. This shows that compositional consistency is important for long-form videoQA, since the model needs to grasp the relations among entities, specifically those specified by the question.

Position of State Space Layer. We replace the penultimate multi-modal attention with SSL. As shown in Table 10, such design choice deteriorates the videoQA performance. The reason might be that since the frames and regions have already been selected, limited global information can be extracted from the video. Moreover, SSL does not explicitly calculate dependency between tokens, thus producing little refined representations to compute the final answer, which has been observed by previous work (Zuo et al., 2022). This substanti-

ates our decision to adopt SSL to integrate global semantics of video in an earlier stage.

5.6 Qualitative Results

We visualize videoQA cases in Figure 1. As can be observed, our model can choose the correct answer for questions that require information over the video with a limited number of video segments. We posit that due to our integrated global semantics, visual representations not only encode information of the selected segments but also the video context, thus furnish the model with sufficient cues to ascertain the correct answer. In contrast, constrained to the selected segments, previous state-of-the-art, *i.e.* MIST-CLIP (Gao et al., 2023), struggles in such questions and produces the incorrect output.

6 Conclusion

We introduce a Gated State space Multi-modal Transformer (GSMT) with a state space layer (SSL) to integrate global semantics of video into visual representations to tackle long-form videoQA. We further incorporate a gating unit to provide more controllability over the integrated global semantics and a cross-modal compositional congruence (C^3) objective to encourage the semantics aligned with the question. To comprehensively evaluate long-form videoQA, we curate two long-form videoQA datasets with excessively long video lengths and long-natured questions. Extensive experiments on these and standard datasets validate the superiority of our framework.

7 Limitations

Our proposed framework has achieved promising improvement by integrating productive global context of long videos for long-form videoQA, but we consider the following limitations as future work:

- **Improve the generalizability of videoQA systems:** Long videos exhibit diverse content, which make it unlikely that the model will encounter inputs of the same distribution during inference. Since training different models for different settings is computationally costly, it is desirable to construct a long-form videoQA model that can generalize to myriad content.
- **Extend our datasets to multicultural settings:** Many long videos from different nations with distinct backgrounds are prevalent in the Internet. Therefore, to increase the usefulness of long-form videoQA model, there is a need to construct

benchmarks for different cultures and societies to facilitate the development of future multi-cultural long-form videoQA systems.

Acknowledgement

This research/project is supported by the National Research Foundation, Singapore under its AI Singapore Programme (AISG Award No: AISG3-PhD-2023-08-051T). Thong Nguyen is supported by a Google Ph.D. Fellowship in Natural Language Processing.

References

- Josh Achiam, Steven Adler, Sandhini Agarwal, Lama Ahmad, Ilge Akkaya, Florencia Leoni Aleman, Diogo Almeida, Janko Altenschmidt, Sam Altman, Shyamal Anadkat, et al. 2023. Gpt-4 technical report. *arXiv preprint arXiv:2303.08774*.
- Jean-Baptiste Alayrac, Jeff Donahue, Pauline Luc, Antoine Miech, Iain Barr, Yana Hasson, Karel Lenc, Arthur Mensch, Katherine Millican, Malcolm Reynolds, et al. 2022. Flamingo: a visual language model for few-shot learning. *Advances in neural information processing systems*, 35:23716–23736.
- Max Bain, Arsha Nagrani, Gül Varol, and Andrew Zisserman. 2021. Frozen in time: A joint video and image encoder for end-to-end retrieval. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 1728–1738.
- Shyamal Buch, Cristóbal Eyzaguirre, Adrien Gaidon, Jiajun Wu, Li Fei-Fei, and Juan Carlos Niebles. 2022. Revisiting the "video" in video-language understanding. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 2917–2927.
- Laetitia Chapel, Mokhtar Z Alaya, and Gilles Gasso. 2020. Partial optimal transport with applications on positive-unlabeled learning. *Advances in Neural Information Processing Systems*, 33:2903–2913.
- James W Cooley and John W Tukey. 1965. An algorithm for the machine calculation of complex fourier series. *Mathematics of computation*, 19(90):297–301.
- Tsu-Jui Fu, Linjie Li, Zhe Gan, Kevin Lin, William Yang Wang, Lijuan Wang, and Zicheng Liu. 2023. An empirical study of end-to-end video-language transformers with masked visual modeling. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 22898–22909.
- Mona Gandhi, Mustafa Omer Gul, Eva Prakash, Madeleine Grunde-McLaughlin, Ranjay Krishna, and Maneesh Agrawala. 2022. Measuring compositional consistency for video question answering. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 5046–5055.
- Difei Gao, Ruiping Wang, Ziyi Bai, and Xilin Chen. 2021. Env-qa: A video question answering benchmark for comprehensive understanding of dynamic environments. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 1675–1685.
- Difei Gao, Luowei Zhou, Lei Ji, Linchao Zhu, Yi Yang, and Mike Zheng Shou. 2023. Mist: Multi-modal iterative spatial-temporal transformer for long-form video question answering. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 14773–14783.
- Kristen Grauman, Andrew Westbury, Eugene Byrne, Zachary Chavis, Antonino Furnari, Rohit Girdhar, Jackson Hamburger, Hao Jiang, Miao Liu, Xingyu Liu, et al. 2022. Ego4d: Around the world in 3,000 hours of egocentric video. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 18995–19012.
- Madeleine Grunde-McLaughlin, Ranjay Krishna, and Maneesh Agrawala. 2021. Agqa: A benchmark for compositional spatio-temporal reasoning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 11287–11297.
- Albert Gu, Karan Goel, and Christopher Ré. 2021. Efficiently modeling long sequences with structured state spaces. *arXiv preprint arXiv:2111.00396*.
- Ankit Gupta, Albert Gu, and Jonathan Berant. 2022. Diagonal state spaces are as effective as structured state spaces. *Advances in Neural Information Processing Systems*, 35:22982–22994.
- Md Mohaiminul Islam and Gedas Bertasius. 2022. Long movie clip classification with state-space video models. In *European Conference on Computer Vision*, pages 87–104. Springer.
- Md Mohaiminul Islam, Ngan Ho, Xitong Yang, Tushar Nagarajan, Lorenzo Torresani, and Gedas Bertasius. 2024. Video recap: Recursive captioning of hour-long videos. *arXiv preprint arXiv:2402.13250*.
- Yunseok Jang, Yale Song, Youngjae Yu, Youngjin Kim, and Gunhee Kim. 2017. Tgif-qa: Toward spatio-temporal reasoning in visual question answering. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 2758–2766.
- Chao Jia, Yinfei Yang, Ye Xia, Yi-Ting Chen, Zarana Parekh, Hieu Pham, Quoc Le, Yun-Hsuan Sung, Zhen Li, and Tom Duerig. 2021. Scaling up visual and vision-language representation learning with noisy text supervision. In *International conference on machine learning*, pages 4904–4916. PMLR.

- Eric Kolve, Roozbeh Mottaghi, Winson Han, Eli VanderBilt, Luca Weihs, Alvaro Herrasti, Matt Deitke, Kiana Ehsani, Daniel Gordon, Yuke Zhu, et al. 2017. Ai2-thor: An interactive 3d environment for visual ai. *arXiv preprint arXiv:1712.05474*.
- Jie Lei, Licheng Yu, Mohit Bansal, and Tamara Berg. 2018. Tvqa: Localized, compositional video question answering. In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 1369–1379.
- Jie Lei, Licheng Yu, Tamara Berg, and Mohit Bansal. 2020. Tvqa+: Spatio-temporal grounding for video question answering. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 8211–8225.
- Stan Weixian Lei, Yuxuan Wang, Dongxing Mao, Difei Gao, and Mike Zheng Shou. 2021. Assistsr: Affordance-centric question-driven video segment retrieval. *arXiv preprint*.
- Xiangpeng Li, Jingkuan Song, Lianli Gao, Xianglong Liu, et al. 2019. Beyond rnns: Positional self-attention with co-attention for video question answering. In *AAAI*, pages 8658–8665.
- Kevin Qinghong Lin, Jinpeng Wang, Mattia Soldan, Michael Wray, Rui Yan, Eric Z Xu, Difei Gao, Rong-Cheng Tu, Wenzhe Zhao, Weijie Kong, et al. 2022a. Egocentric video-language pretraining. *Advances in Neural Information Processing Systems*, 35:7575–7586.
- Yan-Bo Lin, Jie Lei, Mohit Bansal, and Gedas Bertasius. 2022b. Eclipse: Efficient long-range video retrieval using sight and sound. In *European Conference on Computer Vision*, pages 413–430. Springer.
- Karttikeya Mangalam, Raiymbek Akshulakov, and Jitendra Malik. 2023. Egoschema: A diagnostic benchmark for very long-form video language understanding. *arXiv preprint arXiv:2308.09126*.
- Cong-Duy Nguyen, Thong Nguyen, Duc Vu, and Anh Luu. 2023a. Improving multimodal sentiment analysis: Supervised angular margin-based contrastive learning for enhanced fusion representation. In *Findings of the Association for Computational Linguistics: EMNLP 2023*, pages 14714–14724.
- Cong-Duy Nguyen, Thong Nguyen, Xiaobao Wu, and Anh Tuan Luu. 2024a. Kdmcse: Knowledge distillation multimodal sentence embeddings with adaptive angular margin contrastive learning. *arXiv preprint arXiv:2403.17486*.
- Cong-Duy Nguyen, The-Anh Vu-Le, Thong Nguyen, Tho Quan, and Anh-Tuan Luu. 2023b. Expand bert representation with visual information via grounded language learning with multimodal partial alignment. In *Proceedings of the 31st ACM International Conference on Multimedia*, pages 5665–5673.
- Thong Nguyen, Yi Bin, Xiaobao Wu, Xinshuai Dong, Zhiyuan Hu, Khoi Le, Cong-Duy Nguyen, See-Kiong Ng, and Luu Anh Tuan. 2024b. Meta-optimized angular margin contrastive framework for video-language representation learning. *arXiv preprint arXiv:2407.03788*.
- Thong Nguyen, Yi Bin, Junbin Xiao, Leigang Qu, Yicong Li, Jay Zhangjie Wu, Cong-Duy Nguyen, See-Kiong Ng, and Luu Anh Tuan. 2024c. Video-language understanding: A survey from model architecture, model training, and data perspectives. *arXiv preprint arXiv:2406.05615*.
- Thong Nguyen and Anh Tuan Luu. 2021. Contrastive learning for neural topic model. *Advances in neural information processing systems*, 34:11974–11986.
- Thong Nguyen, Anh Tuan Luu, Truc Lu, and Tho Quan. 2021. Enriching and controlling global semantics for text summarization. *arXiv preprint arXiv:2109.10616*.
- Thong Nguyen, Cong-Duy Nguyen, Xiaobao Wu, See-Kiong Ng, and Anh Tuan Luu. 2022a. Vision-and-language pretraining. *arXiv preprint arXiv:2207.01772*.
- Thong Nguyen, Xiaobao Wu, Xinshuai Dong, Khoi M Le, Zhiyuan Hu, Cong-Duy Nguyen, See-Kiong Ng, and Anh Tuan Luu. 2024d. Read-pvla: Recurrent adapter with partial video-language alignment for parameter-efficient transfer learning in low-resource video-language modeling. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 38, pages 18824–18832.
- Thong Nguyen, Xiaobao Wu, Xinshuai Dong, Anh Tuan Luu, Cong-Duy Nguyen, Zhen Hai, and Lidong Bing. 2023c. Gradient-boosted decision tree for listwise context model in multimodal review helpfulness prediction. *arXiv preprint arXiv:2305.12678*.
- Thong Nguyen, Xiaobao Wu, Xinshuai Dong, Cong-Duy Nguyen, See Kiong Ng, and Anh Luu. 2023d. Demaformer: Damped exponential moving average transformer with energy-based modeling for temporal language grounding. In *Findings of the Association for Computational Linguistics: EMNLP 2023*, pages 3635–3649.
- Thong Nguyen, Xiaobao Wu, Anh-Tuan Luu, Cong-Duy Nguyen, Zhen Hai, and Lidong Bing. 2022b. Adaptive contrastive learning on multimodal transformer for review helpfulness predictions. *arXiv preprint arXiv:2211.03524*.
- Thong Thanh Nguyen and Anh Tuan Luu. 2022. Improving neural cross-lingual abstractive summarization via employing optimal transport distance for knowledge distillation. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 36, pages 11103–11111.
- Thong Thanh Nguyen, Xiaobao Wu, Xinshuai Dong, Cong-Duy T Nguyen, See-Kiong Ng, and Anh Tuan

- Luu. 2024e. Topic modeling as multi-objective optimization with setwise contrastive learning. In *The Twelfth International Conference on Learning Representations*.
- Rohan Pandey, Rulin Shao, Paul Pu Liang, Ruslan Salakhutdinov, and Louis-Philippe Morency. 2022. Cross-modal attention congruence regularization for vision-language relation alignment. *arXiv preprint arXiv:2212.10549*.
- Shraman Pramanick, Aniket Roy, and Vishal M Patel. 2022. Multimodal learning using optimal transport for sarcasm and humor detection. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision*, pages 3930–3940.
- Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal, Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, et al. 2021. Learning transferable visual models from natural language supervision. In *International Conference on Machine Learning*, pages 8748–8763.
- Shuhuai Ren, Junyang Lin, Guangxiang Zhao, Rui Men, An Yang, Jingren Zhou, Xu Sun, and Hongxia Yang. 2021. Learning relation alignment for calibrated cross-modal retrieval. *arXiv preprint arXiv:2105.13868*.
- Mattia Soldan, Alejandro Pardo, Juan León Alcázar, Fabian Caba, Chen Zhao, Silvio Giancola, and Bernard Ghanem. 2022. Mad: A scalable dataset for language grounding in videos from movie audio descriptions. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 5026–5035.
- Makarand Tapaswi, Yukun Zhu, Rainer Stiefelwagen, Antonio Torralba, Raquel Urtasun, and Sanja Fidler. 2016. Movieqa: Understanding stories in movies through question-answering. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 4631–4640.
- Jinpeng Wang, Yixiao Ge, Rui Yan, Yuying Ge, Kevin Qinghong Lin, Satoshi Tsutsui, Xudong Lin, Guanyu Cai, Jianping Wu, Ying Shan, et al. 2023. All in one: Exploring unified video-language pre-training. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 6598–6608.
- Bo Wu, Shoubin Yu, Zhenfang Chen, Joshua B Tenenbaum, and Chuang Gan. 2021. Star: A benchmark for situated reasoning in real-world videos. In *Thirty-fifth Conference on Neural Information Processing Systems Datasets and Benchmarks Track (Round 2)*.
- Chao-Yuan Wu and Philipp Krahenbuhl. 2021. Towards long-form video understanding. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 1884–1894.
- Xiaobao Wu, Xinshuai Dong, Thong Nguyen, Chaoqun Liu, Liang-Ming Pan, and Anh Tuan Luu. 2023a. Infotm: A mutual information maximization perspective of cross-lingual topic modeling. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 37, pages 13763–13771.
- Xiaobao Wu, Xinshuai Dong, Thong Thanh Nguyen, and Anh Tuan Luu. 2023b. Effective neural topic modeling with embedding clustering regularization. In *International Conference on Machine Learning*, pages 37335–37357. PMLR.
- Xiaobao Wu, Xinshuai Dong, Liangming Pan, Thong Nguyen, and Anh Tuan Luu. 2024a. Modeling dynamic topics in chain-free fashion by evolution-tracking contrastive learning and unassociated word exclusion. *arXiv preprint arXiv:2405.17957*.
- Xiaobao Wu, Thong Nguyen, and Anh Tuan Luu. 2024b. A survey on neural topic models: methods, applications, and challenges. *Artificial Intelligence Review*, 57(2):18.
- Xiaobao Wu, Thong Nguyen, Delvin Ce Zhang, William Yang Wang, and Anh Tuan Luu. 2024c. Fastopic: A fast, adaptive, stable, and transferable topic modeling paradigm. *arXiv preprint arXiv:2405.17978*.
- Xiaobao Wu, Fengjun Pan, Thong Nguyen, Yichao Feng, Chaoqun Liu, Cong-Duy Nguyen, and Anh Tuan Luu. 2024d. On the affinity, rationality, and diversity of hierarchical topic modeling. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 38, pages 19261–19269.
- Junbin Xiao, Xindi Shang, Angela Yao, and Tat-Seng Chua. 2021. Next-qa: Next phase of question-answering to explaining temporal actions. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 9777–9786.
- Junbin Xiao, Angela Yao, Zhiyuan Liu, Yicong Li, Wei Ji, and Tat-Seng Chua. 2022. Video as conditional graph hierarchy for multi-granular question answering. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 36, pages 2804–2812.
- Junbin Xiao, Pan Zhou, Angela Yao, Yicong Li, Richang Hong, Shuicheng Yan, and Tat-Seng Chua. 2023. Contrastive video question answering via video graph transformer. *arXiv preprint arXiv:2302.13668*.
- Dejing Xu, Zhou Zhao, Jun Xiao, Fei Wu, Hanwang Zhang, Xiangnan He, and Yueting Zhuang. 2017. Video question answering via gradually refined attention over appearance and motion. In *Proceedings of the 25th ACM international conference on Multimedia*, pages 1645–1653.
- Antoine Yang, Antoine Miech, Josef Sivic, Ivan Laptev, and Cordelia Schmid. 2021. Just ask: Learning to answer questions from millions of narrated videos. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 1686–1697.

Rowan Zellers, Jiasen Lu, Ximing Lu, Youngjae Yu, Yanpeng Zhao, Mohammadreza Salehi, Aditya Kusupati, Jack Hessel, Ali Farhadi, and Yejin Choi. 2022. Merlot reserve: Neural script knowledge through vision and language and sound. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 16375–16387.

Zhou Zhao, Zhu Zhang, Shuwen Xiao, Zhou Yu, Jun Yu, Deng Cai, Fei Wu, and Yueting Zhuang. 2018. Open-ended long-form video question answering via adaptive hierarchical reinforced networks. In *IJCAI*, volume 2, page 8.

Simiao Zuo, Xiaodong Liu, Jian Jiao, Denis Charles, Eren Manavoglu, Tuo Zhao, and Jianfeng Gao. 2022. Efficient long sequence modeling via state space augmented transformer. *arXiv preprint arXiv:2212.08136*.

A Prompt for generating Ego-QA and MAD-QA datasets

A.1 Question prompt

I want you to act as a teacher in the class called ‘‘Long-term video understanding’’. I will provide video narrations describing events in the time order of the video and you will generate highly difficult and diverse questions for your students about the high-level details in the video. You want to test students’ following abilities:

Ability 1: Students’ ability to summarize and compare long parts of the video

Ability 2: Students’ ability to compress information from the video rather than just listing the actions that happened in the video.

Ability 3: Students’ ability to identify the most important parts of the video.

Your questions should not mention any particular timestamps or narrations. Remember to make sure the correct answers to your questions do not list information from the narrations but compress them in a concise conclusion.

Examples of good and difficult questions:

‘‘What if A happened instead of ...?’’

‘‘Why did A do action ...?’’

‘‘What did A do after/before ...?’’

AVOID the following types of questions:

‘‘When ...?’’

‘‘How many ...?’’

‘‘How much ...?’’

When announcing the question please label each question as ‘‘Question 1,2,3: [full question]’’

Video narrations: [video narrations]

A.2 Answer prompt

I want you to act as a teacher in the class called ‘‘Long-term video understanding.’’ I will provide video action narrations and highly difficult and diverse questions for your students about the high-level details in the video. I want you to test students’ following abilities:

Ability 1: Students’ ability to summarize and compare long parts of the video

Ability 2: Students’ ability to compress information from the video rather than just listing the actions that happened in the video.

Ability 3: Students’ ability to identify the most important parts of the video.

I want you to create a difficult multiple-choice exam that tests above student abilities based on the questions I just provided. Each question should have five similar open-ended but short answers, but only one should be correct. Make it very difficult for students to find the correct answer among all the wrong answers. All answers should be closely related to what happens in the video. Make wrong answers significantly longer than correct answers. Ensure all of the correct answers compress information from narrations them into a concise conclusion. Your answers should not mention any particular timestamps or narrations.

Do not use letters for the answer choices

Print each correct answer exactly as ‘‘Correct answer: [full answer]’’

Please print each wrong answer on a new line and print each wrong answer as ‘‘Wrong answer 1,2,3,4: [full answer]’’

Video narrations: [video narrations]

Questions: [question list]

B Manual annotation

We utilize ten professional English speakers to ensure the quality of our datasets. Our manual annotation consists of two procedures: question filtering and human accuracy testing.

B.1 Question filtering

Human annotators are responsible for ensuring that the question-answer samples in two datasets are of the highest quality possible. We provide the instruction that we show to the annotators:

The annotation we need is to say that the Question-correct answer-wrong answer set (the whole set) is good if all these three conditions pass:

(Condition A) Question is Answerable: The question can be answered from the video and requires more than just a second of video to answer (so, if the answer is not present in the video or, if the answer can be formed with just a few frames (less than say, a second) then it fails this condition).

(Condition B) The Marked Correct Answer is correct: The ""correct answer"" is more the correct answer to the question

(Condition C) The Marked Wrong Answers are wrong: All 4 ""wrong answers"" are less correct than the ""correct answer"" (So for example, if a wrong answer is not completely false, but simply does not contain all the information that the "" correct answer"" does, then it is still a fine "" wrong answer"") IF even one of the marked answer is correct, the set should be labeled as bad.

(Condition D) The question is actually long-term: This is a very very important condition. We want the certificate for the question to be at least 30 seconds minimum. If the certificate is non-contiguous (I.e. 5 seconds at one place, 20 seconds at another, and 15 more seconds at a third place) the sum of lengths of all the sub-certificates together should be more than 30 seconds. Another example is, if a question can be answered simply from a few frames of the video, the certificate is small (and less than 30 seconds) and hence would fail this condition. Additional details on how to handle certificate edge cases are provided in the annotator training through examples.

(Condition E) Avoid Boring Questions: Questions that ask about the frequency of something ("How many times..") fail this condition.

If any of these five conditions fail we want the whole set (Question / Correct Answer / Wrong Answer) marked bad.

Optional:

Since GOOD sets are so rare, in cases where it seems that a set is good but a small part of the above five conditions is not being met or, if one/two words were different this can be a good set, please label as MAYBE and we will fix it in the second round.

Extended notes:

1. In our experience, the wrong answers are made such that they differ from the correct answer in small but crucially meaningful ways. There are many cases where a careless reading of the wrong answer might make it seem that it is correct but upon careful inspection, it will become clear that something about the wrong answer indeed makes it wrong. While this is common, there are indeed cases where the wrong answer is also just as correct as the correct answer. In such cases, the wrong answer fails condition C, and the set is bad.
2. Roughly speaking, we expect about 20-25% of the questions that we have provided to be found as good . However, this is not necessary and the percentage can be smaller or larger depending on each three-minute clip.
3. Edge Cases:
 1. If the asked question has multiple answers and at least one of them aligns with the correct answer while none of them align with any of the other wrong answers, then provided that the top 5 conditions are met, we can mark the set as good.
 2. If two questions are very similar (even within different clips) and both are GOOD, only choose one as GOOD and reject the other one with a comment mentioning this. We do not expect this to happen more than 1 or 2 times in a 100.
 3. There might be more such edge cases, please feel free to contact me in such cases and we can explain.

B.2 Human accuracy testing

To benchmark human, we recruit another team of ten human annotators to answer the questions in the curated datasets. The answers are randomly shuffled and presented in the form of a test.

C Examples of Ego-QA dataset

In this appendix, we provide more examples of our constructed Ego-QA dataset. Our videos are taken from the Ego4D dataset (Grauman et al., 2022).

Video input: 1246d6ec-5620-4f71-8b4b-d823775f58c2 - Video length: 30.6 minutes



Selected segments:



Question: From the descriptions of interactions, what conclusion can be drawn about the primary focus of the video?

- 0. Exploring the dynamics of household relationships
- 1. The efficacy of communication technology in a domestic setting
- 2. The impact of routine activities on personal development
- 3. The importance of meal preparation and dietary habits
- 4. Highlighting the isolation experienced by individuals in a family setting


GSMT's prediction: 0 ✓ MIST-CLIP's prediction: 3 ✗

Video length to watch to answer the question: 656.94 seconds

Video input: 3d751dba-7897-4ad2-8baa-2569f32dbe6d - Video length: 27.5 minutes



Selected segments:



Question: Considering the sequence of activities performed, what can be inferred about the primary purpose of the video?

- 0. to showcase a comprehensive home workout routine that incorporates various exercises
- 1. to demonstrate the correct usage and maintenance of workout equipment at home
- 2. to compare the effectiveness of different exercise techniques for physical fitness
- 3. to highlight the role of technology and mobile devices in enhancing workout efficiency
- 4. to emphasize the importance of rest and recovery periods during a workout session


GSMT's prediction: 0 ✓ MIST-CLIP's prediction: 3 ✗

Video length to watch to answer the question: 824.83 seconds

Video input: d2985e7b-6dc8-4c7e-9aef-40d633895aa0 - Video length: 24.4 minutes



Selected segments:



Question: Considering the characters' various activities with electronic devices, books, and games, what can be inferred about the video's message on leisure and entertainment?

- 0. it critiques the overwhelming presence of electronic devices in daily life
- 1. it celebrates the diversity of activities available for leisure and entertainment
- 2. it underscores the necessity of balancing technology use with traditional activities
- 3. it contrasts the solitary activities against group interactions as forms of leisure
- 4. it shows a progression from individual to group activities, highlighting the importance of community

GSMT's prediction: 1 ✓ MIST-CLIP's prediction: 4 ✗

Video length to watch to answer the question: 733.89 seconds

D Examples of MAD-QA dataset

In this appendix, we provide more examples of our constructed MAD-QA dataset. Our videos are taken from the MAD dataset (Soldan et al., 2022).

Video input: 3076 *THE SOCIAL NETWORK* - Video length: 2.0 hours



Selected segments:



Question: why is someone exiting the thirsty scholar pub at this moment?

0. get some fresh air
1. use the restroom
2. make a phone call
3. **return home**
4. buy cigarettes

GSMT's prediction: 3 ✓ MIST-CLIP's prediction: 0 ✗

Video length to watch to answer the question: 134.4 seconds

Video input: 3042 *KARATE KID* - Video length: 2.1 hours



Selected segments:



Question: why are there mostly asian travelers on the plane?

0. airline has a special deal for Asian passengers
1. **plane is likely flying to or from an asian country**
2. plane's altitude is more comfortable for Asian people
3. passengers won a free trip to Asia
4. plane is carrying a large group of Asian tourists


GSMT's prediction: 1 ✓ MIST-CLIP's prediction: 4 ✗

Video length to watch to answer the question: 235.83 seconds

Video id input: 3077 *THE VOW* - Video length: 1.7 hours



Selected segments:



Question: what if the brunette woman did not share a smile with the man at the government building?

0. man would have still rushed to the hospital, but for a different reason
1. man and woman would not have started a passionate romance at the government building
2. brunette woman would have been perfectly content without sharing a smile with the man
3. **couple might not have formed a relationship, and the man would not have rushed to the hospital after the accident**
4. lack of a shared smile would have had no impact on their future interactions

GSMT's prediction: 3 ✓ MIST-CLIP's prediction: 1 ✗

Video length to watch to answer the question: 206.78 seconds

E Dataset Statistics

In this appendix, we provide the total number of questions and videos, and the average number of narration sentences per video in Table 11. We also visualize the distribution of video lengths of datasets widely used for long-form videoQA, *i.e.* NExT-QA (Xiao et al., 2021), STAR (Wu et al., 2021), EgoSchema (Mangalam et al., 2023), and our curated datasets MAD-QA and Ego-QA in Figure 5. In addition, we show the distribution of temporal certificate lengths, *i.e.* video lengths humans need to watch to determine the answer, of standard long-form videoQA and our datasets in Figure 6. As can be observed, our datasets exhibit longer video input length and also certificate length to verify the correct answer, validating their effectiveness to evaluate long-form videoQA performance.

Dataset	# Questions	# Videos	# Narration sentences per video (Avg)
Ego-QA	18838	992	279.9
MAD-QA	15674	650	641.3

Table 11: Statistics of the curated Ego-QA and MAD-QA datasets.

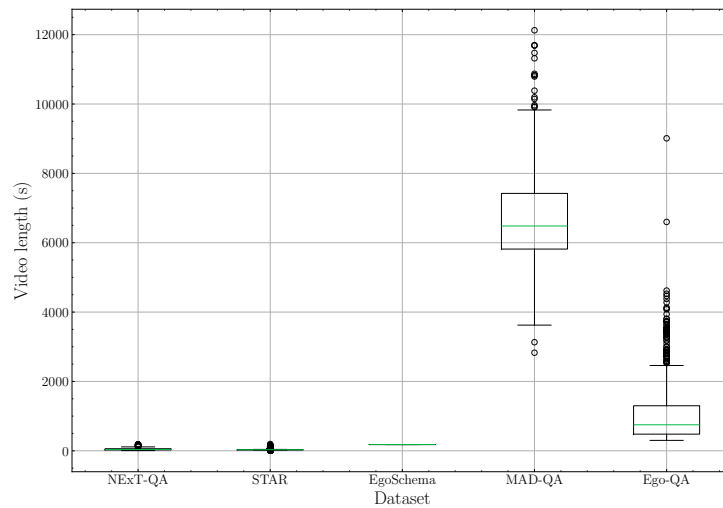


Figure 5: Distribution of video input duration of long-form videoQA datasets.




Figure 6: Distribution of temporal certificate lengths of long-form videoQA datasets.