

Video as Conditional Graph Hierarchy for Multi-Granular Question Answering

Junbin Xiao, Angela Yao, Zhiyuan Liu, Yicong Li, Wei Ji, Tat-Seng Chua

Department of Computer Science, National University of Singapore

{junbin, ayao, chuats}@comp.nus.edu.sg, acharkq@gmail.com, liyicong@u.nus.edu, jiwei@nus.edu.sg

Abstract

Video question answering requires the models to understand and reason about both the complex video and language data to correctly derive the answers. Existing efforts have been focused on designing sophisticated cross-modal interactions to fuse the information from two modalities, while encoding the video and question holistically as frame and word sequences. Despite their success, these methods are essentially revolving around the sequential nature of video- and question-contents, providing little insight to the problem of question-answering and lacking interpretability as well. In this work, we argue that while video is presented in frame sequence, the visual elements (*e.g.*, objects, actions, activities and events) are not sequential but rather hierarchical in semantic space. To align with the multi-granular essence of linguistic concepts in language queries, we propose to model video as a conditional graph hierarchy which weaves together visual facts of different granularity in a level-wise manner, with the guidance of corresponding textual cues. Despite the simplicity, our extensive experiments demonstrate the superiority of such conditional hierarchical graph architecture, with clear performance improvements over prior methods and also better generalization across different type of questions. Further analyses also demonstrate the model’s reliability as it shows meaningful visual-textual evidences for the predicted answers.

Introduction

The past few years have witnessed a flourish of research in video-language tasks. Video question answering (VideoQA) is one of the most prominent, given its promise to develop interactive AI and communicate with the dynamic visual world via natural language. Despite its popularity, the challenge in VideoQA remains significant; it demands a wide spectrum of recognitive (Ren et al. 2015; He et al. 2016; Carreira and Zisserman 2017), reasoning (Hu et al. 2018) as well as grounding (Hu et al. 2017; Xiao et al. 2020) capabilities to comprehend the questions (Mao et al. 2016) and deduce the correct answers. Existing efforts (Jang et al. 2017; Gao et al. 2018; Yu, Kim, and Kim 2018; Fan et al. 2019; Li et al. 2019b; Jiang et al. 2020) concentrate on capturing the sequential nature of video frames and question words. As the hierarchical, compositional structure of video contents and the multi-granular essence of linguistic concepts

are unaccounted, current models are limited in insight and interpretability, and also their results can be sub-optimal.

In this work, we delve into the problem of video question answering based on a bottom-up and top-down insight¹. As demonstrated in Figure 1, from a bottom-up view, video contents are hierarchical in semantic space. Each atomic action (*e.g.*, pick up) involves a visual subject (*e.g.*, lady) and optionally an object (*e.g.*, lemon). Besides, the atomic actions (*e.g.*, pick up, squeeze, pour) compositionally form a super-action or activity (*e.g.*, cook), while activities (*e.g.*, cook and play) further constitute a global event (*e.g.*, camping). From a top-down view, questions are diverse and answering them demands visual information from different granularity levels. Generally, questions concerning objects and their attributes rely on specific visual entities for answers; questions regarding spatial relations and contact actions (*e.g.*, hug, kiss, hold, carry) are better answered based on object interactions at certain frames; while questions about dynamic actions (*e.g.*, pick up, put down), temporal interactions and global event may require aggregating information from multiple video frames or clips. In addition, a single question may also invoke multiple levels of visual elements, which further demands the awareness of the multi-granularity of both video elements and linguistic concepts.

To capture such insight, we propose to model video as a graph hierarchy with the condition of language query in a level-wise fashion. Concretely, the model weaves together visual facts from low-level entities to higher level video elements through graph aggregation and pooling, during which the local and global textual query are incorporated into different levels to match and pinpoint the relevant video elements at the corresponding granularity levels (*e.g.*, objects, actions, activities and events). In this way, our model can not only identify the query-specific visual objects and actions, but also capture their local interactions as well as infer the constituted activities and events. Such versatility is the first of its kind in VideoQA and is of crucial importance in handling diverse questions concerning different video elements. To validate the effectiveness, we test the model on four VideoQA datasets that challenge the various aspects of

¹Bottom-up is from low-level visual contents (video) to high-level semantics (language query), and vice versa for top down.

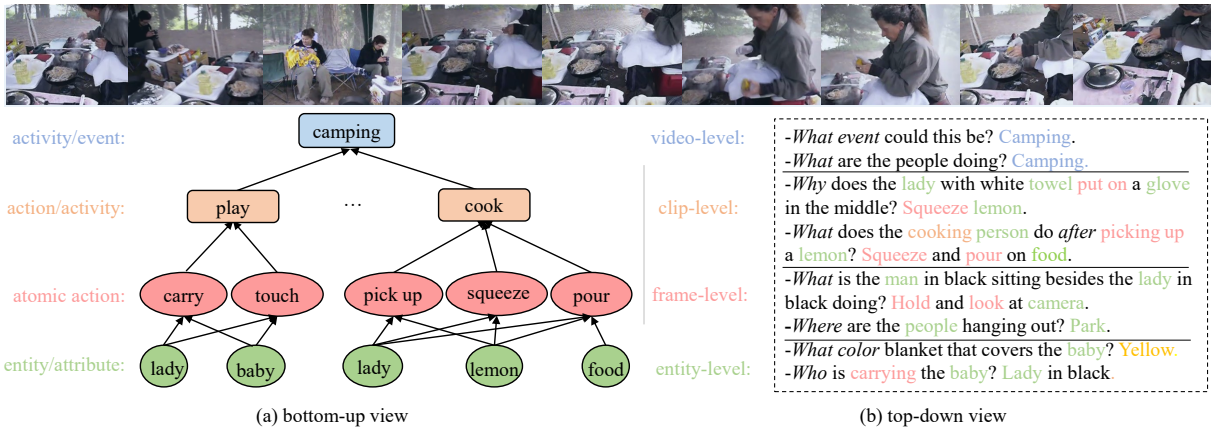


Figure 1: The bottom-up and top-down insights for video question answering. (a) From a bottom-up view, the video contents are hierarchical from low-level visual entities, their local interactions, to high-level activities and global event. (b) From a top-down view, different questions demand video elements at different granularity (e.g., objects, actions, activities and events) for answers, and a single question might also invoke multi-level visual resources to comprehend.

video understanding and achieve consistently strong results.

To summarize our main contributions: 1) We provide a bottom-up and top-down insight to advance video question answering in a multi-granular fashion. 2) We propose a hierarchical conditional graph model, which serves as an initial prompt, to capture such insight for VideoQA. 3) Extensive experiments evince that our model is effective and is of enhanced generalizability and interpretability; it achieves the state-of-the-art (SOTA) results across different datasets with various type of questions and finds introspective evidences for the predicted answers as well.

Related Work

Canonical approaches use techniques such as cross-model attention (Jang et al. 2017; Zeng et al. 2017; Li et al. 2019b; Jin et al. 2019; Gao et al. 2019; Jiang et al. 2020) and motion-appearance memory (Xu et al. 2017; Gao et al. 2018; Fan et al. 2019) to fuse information from the video and question for answer prediction. These methods focus on designing sophisticated cross-modal interactions, whereas treating video and question as a holistic sequence of frames and words respectively. Sequential modelling neither capture the topological (e.g., hierarchical or compositional) structure of the visual elements nor multi-granularity of linguistic concepts. Consequently, the derived QA models are weak in relation reasoning and handling question diversity.

Graph-structured models (Kipf and Welling 2017; Veličković et al. 2018) are recently more favoured, either for their superior performance in relation reasoning (Li et al. 2019a; Hu et al. 2019), or for the improved interpretability (Norcliffe-Brown, Vafeias, and Parisot 2018). L-GCN (Huang et al. 2020) constructs a fully-connected graph over all the detected regions in space-time and demonstrates the benefits of utilizing object locations. However, the monolithic graph is cumbersome to extend to long videos with multiple objects. More recently, GMIN (Gu et al. 2021) builds a spatio-temporal graph over object trajectories and shows improvements over its attention version (Jin et al.

2019). While L-GCN and GMIN construct query-blind graphs, HGA (Jiang and Han 2020), DualVGR (Wang, Bao, and Xu 2021) and B2A (Park, Lee, and Sohn 2021) design query-specific graphs for better performance. Yet, their graphs are built over coarse video segments in a flat way. On the one hand, these models cannot reason fine-grained object interactions in space-time. On the other hand, they are unable to reflect the hierarchical nature of video contents.

Hierarchical architectures. HCRN (Le et al. 2020) designs and stacks conditional relation blocks to capture temporal relations. Nonetheless, it focuses on temporal reasoning of single object actions and models relations with a simple mean-pooling. As a result, it fails to generalize well to the scenarios with multiple object interacted in space-time (Xiao et al. 2021). A very recent work HOSTR (Dang et al. 2021) follows a similar design philosophy to learn a hierarchical video representation, but introduces a nested graph for spatio-temporal reasoning over object trajectories, and achieves better performance. Yet, both HCRN and HOSTR target at general-purpose visual reasoning. They lack the insights that 1) some video elements (e.g., object, places, spatial relations and contact actions) are easier to identify at frame-level (Gkioxari et al. 2018; Xiao et al. 2020), and 2) different parts of a question may invoke visual information at different granularity levels (Chen et al. 2020). In addition, HOSTR’s good performance relies on accurate object trajectories which are hard to obtain in practice, especially for long video sequences with complex object interactions. In this work, we design and emphasize the hierarchical architecture to realize the bottom-up and top-down insights which enables vision-text matching at multi-granularity levels.

Method

Overview Given a video \mathcal{V} and a question q , VideoQA aims to predict the correct answer a^* that is relevant to the visual content. Currently, the two typical QA formats are multi-choice QA and open-ended QA. In multi-choice QA, each question is presented with several candidate answers




Figure 2: A conditional graph hierarchy built over 4 video clips. G_O , G_F and G_C are graphs defined over objects, their interactions across frames and across clips respectively. The graphs are stacked hierarchically to abstract low-level, locally-related visual components into a high-level global representation. The language query Q is conditioned at each level to guide the graph construction. The final aggregated video representation is fused with global query feature f_Q for answer prediction.

\mathcal{A}_{mc} , and the problem is to pick the correct one:

$$a^* = \arg \max_{a \in \mathcal{A}_{mc}} \mathcal{F}_\theta(a|q, \mathcal{V}, \mathcal{A}_{mc}). \quad (1)$$

In open-ended QA, no answer choices are provided. The task is popularly set as a classification problem to classify the video-question pairs into a globally defined answer set \mathcal{A}_{oe} .

$$a^* = \arg \max_{a \in \mathcal{A}_{oe}} \mathcal{F}_\theta(a|q, \mathcal{V}). \quad (2)$$

To address the problem, we realize the introduced bottom-up and top-down insights by modelling video as a conditional graph hierarchy. As illustrated in Figure 2, we accomplish this by designing a Query-conditioned Graph Attention unit (*i.e.*, QGA, see Figure 3) and further applying it to reason and aggregate video elements of different granularity levels into a global representation for question answering.

We next elaborate on our model design by first introducing the data representation, then the QGA unit, and finally the hierarchical architecture and answer decoder.

Data Representation

Video. We extract a video at p frames per second and then partition it into K clips of length L . For each clip, we maintain a dense stream of L frames to obtain the clip-level motion feature and a sparse stream of γL frames ($\gamma \in (0, 1)$) to obtain the region and frame appearance features. In our implementation, the motion features $F_m = \{f_{m_k}\}_{k=1}^K$ and frame appearance features $F_a = \{f_{a_t}\}_{t=1}^T$ ($T = \gamma L K$) are extracted from pre-trained CNNs, specifically 3D version ResNeXt-101 (Hara, Kataoka, and Satoh 2018) for motion and ResNet-101 (He et al. 2016) for frame appearance. Importantly, we also extract N RoIs’ (region of interest) appearance features $F_r = \{f_{r_n}\}_{n=1}^N$ along with their bounding boxes from each frame in the sparse stream, using a pre-trained object detector (Anderson et al. 2018).

After the extraction, all three types of features are projected into a d -dimension space. Specifically, for motion and frame appearance features, the projections are achieved by applying two respective 1-D convolution operations along

the time dimension, in which the window sizes are set to 3 to consider the previous and next neighbors as contexts. For clarity, we denote the projected features as $F_m^{K \times d} = \{f_{m_k}^d\}_{k=1}^K$ for motion and $F_a^{T \times d} = \{f_{a_t}^d\}_{t=1}^T$ for frame appearance. For each object, to retain both semantic and geometric information, we jointly represent its RoI appearance f_r , bounding box location f_s , and temporal position f_t similar to (Huang et al. 2020). The final object feature comes from concatenating the three components and projecting them into d -dimensions with a linear transformation followed by an ELU activation: $f_o = \text{ELU}(W_o[f_r; f_s; f_t])$, in which $[\cdot]$ denotes concatenation and W_o are the parameters of the linear projection. Again, for clarity, we denote the projected object features in a frame as $F_o^{N \times d} = \{f_{o_n}^d\}_{n=1}^N$.

Question. To obtain a well-contextualized word representation, we extract the token-wise sentence embeddings from the penultimate layer of a fine-tuned BERT model (Devlin et al. 2018) (refer to Appendix B for details). Similar to (Xiao et al. 2021), we further apply a Bi-GRU (Cho et al. 2014) to project the word representations into the d -dimension space as the visual part for convenience of cross-model interaction. Consequently, a language query of length M is represented by $Q^{M \times d} = \{q_m^d | q_m = [\vec{q}_i; \overleftarrow{q}_i]\}_{m=1}^M$, where $\vec{q}_i^{\frac{d}{2}}$ and $\overleftarrow{q}_i^{\frac{d}{2}}$ denoted the forward and backward hidden states respectively. Particularly, the last hidden state q_M^d is treated as the global query representation f_Q .

Conditional Graph Reasoning and Pooling

In this section, we introduce QGA - the key component in our model architecture. As illustrated in Figure 3, QGA first contextualizes a set of input visual nodes X_{in} in relation to their neighbors in both the semantic and geometric (realized by the geometric embeddings of the nodes) space, under the condition of a language query $Q^{M \times d}$, and then aggregates the contextualized output nodes X_{out} into a single global descriptor x_p . The input nodes $X_{in}^{* \times d} = \{x_{in_i}\}_{i=1}^*$ depend on the hierarchy level; at the bottom, X_{in} are the object features F_o , while at higher levels, the inputs are the outputs of QGAs at the preceding level. The dimension varies accordingly and




Figure 3: Illustration of QGA unit. Color nodes in Q denote the words referring to video elements at different levels.

we use $*$ as a placeholder for the number of input nodes.

Query Condition. By condition, we pay attention to the video elements that are invoked in the questions, which is achieved by augmenting the corresponding nodes' representations with Q :

$$\hat{x} = x_{in} + \sum_{j=1}^M \alpha_m q_m, \quad \text{where } \alpha = \sigma(x_{in} Q'), \quad (3)$$

where $\sigma(\cdot)$ is the softmax normalization function and $'$ indicates matrix transpose. The indices of x are omitted for brevity. We expect that, through the cross-attention aggregation of the word representations with respect to the visual node x_{in} , the visual node's textual correspondence will have a stronger response in the aggregation if the represented video element is mentioned in the question. As a result, the corresponding video element (represented by x_{in}) will be highlighted and contribute more to subsequent operations.

Graph Attention. After obtaining the augmented node representations $\hat{X}^{* \times d} = \{\hat{x}\}_{i=1}^*$, the edges (including self-loops) represented by the values of the adjacency matrix $A^{* \times *}$ are dynamically computed as the similarities between the node pairs:

$$A = \sigma(\phi_{W_a}(\hat{X})\phi_{W_a}(\hat{X}')), \quad (4)$$

in which the function ϕ_{W_a} denotes linear transformation with learnable parameter $W_a \in \mathbb{R}^{d \times \frac{d}{2}}$. The softmax operation normalizes each row, so that the i^{th} row A_i denotes the attention values (*i.e.*, values of normalized dot-product) of node i with regard to all the other nodes. We then apply a H -layer graph attention aggregation with skip-connections to refine the nodes in relation to their neighbors based on the adjacency matrix A :

$$\hat{X}^{(h)} = \text{ReLU}((A + I)\hat{X}^{(h-1)}W^{(h)}), \quad h \in [1, H], \quad (5)$$

where $W^{(h)} \in \mathbb{R}^{d \times d}$ and $\hat{X}^{(h)}$ are the parameters and outputs of the h^{th} -layer graph attention respectively. $\hat{X}^{(0)}$ is initialized with the query-attended node \hat{X} . I is the identity matrix for skip connections. The final output is obtained by a last skip-connection: $X_{out} = \hat{X} + \hat{X}^{(H)}$.

Node Aggregation. To get an aggregated representation x_p for a QGA unit, we apply self-attention pooling (Lee, Lee, and Kang 2019) over the set of output nodes:

$$x_p = \sum_{i=1}^* \beta_i x_{out_i}, \quad \text{where } \beta = \sigma(\phi_{W_p}(X_{out})), \quad (6)$$

where $W_p \in \mathbb{R}^{d \times 1}$ are the learnable linear mapping weights.

Hierarchical Architecture

In this section, we explain how to apply the QGA units to achieve the hierarchical architecture to reflect the bottom-up and top-down insights for question answering. As shown in Figure 2, the QGA units at the bottom level (G_O) take as inputs a set of object features frame-wisely and capture a static picture of object interactions:

$$F_{G_O} = G_O(F_O) = \text{QGA}(F_O), \quad (7)$$

The output features $F_{G_O}^{T \times d} = \{f_{G_{O_t}}\}_{t=1}^T$ are then combined with the frame appearance features $F_a^{T \times d}$ (serve as global context) by concatenation:

$$\bar{f}_{G_{O_t}} = \text{ELU}(W_{ao}[f_{a_t}; f_{G_{O_t}}]), \quad t \in [1, T], \quad (8)$$

where $W_{ao} \in \mathbb{R}^{2d \times d}$ are linear parameters. Then, the QGA units at the second level (G_F) are applied to $\bar{F}_{G_O}^{T \times d} = \{\bar{f}_{G_{O_t}}\}_{t=1}^T$ clip-wisely to model a short-term interaction dynamics, as well as to reason and aggregate the low-level visual components to a higher granularity level (*e.g.*, from actions to activities):

$$F_{G_F} = G_F(\bar{F}_{G_O}) = \text{QGA}(\bar{F}_{G_O}). \quad (9)$$

The outputs features $F_{G_F} \in \mathbb{R}^{K \times d}$ are then combined with the global motion feature $F_m^{K \times d}$ to obtain \bar{F}_{G_F} in a way analogous to Equ. 8. Furthermore, the QGA unit at the top level (G_C) operates over \bar{F}_{G_F} to reason about the local, short-term interactions, and aggregate them into a single global representation f_V over the entire video:

$$f_V = G_C(\bar{F}_{G_F}) = \text{QGA}(\bar{F}_{G_F}). \quad (10)$$

Overall, the hierarchical architecture is achieved by nesting the conditional graph operations that can be conceptually represented as

$$f_V = G_C(G_F(F_a, G_O(F_o)), F_m). \quad (11)$$

Finally, $f_V \in \mathbb{R}^d$ is passed to the answer-decoder, along with the global query representation f_Q , to jointly determine the correct answers.

By the hierarchical graph structure, the video elements of different granularity can be level-wisely inferred. Besides, by the multi-level token-level query conditions, the model is capable of pinpointing the referred video elements at different granularity, and it is also flexible in handling different textual queries. In addition, by introducing the context features (*i.e.*, F_a and F_m), the model can make up the downside of lacking the respective global information at each level. Importantly, the model is of enhanced flexibility and interpretability, from a perspective of query-instantiated neural modular networks (Hu et al. 2018) and from a introspective analysis of the learned attention weights as our model are purely attention-based, respectively.

Answer Decoder

For multi-choice QA, we concatenate each candidate answer with the corresponding question to form a holistic query.

The resulting global query feature f_Q is fused with the final video feature f_V via Hadamard product (*a.k.a.*, element-wise product) \odot . Then, a full-connected layer with softmax is applied as classifier:

$$s = \sigma(W_c'(f_Q \odot f_V) + b), \quad (12)$$

where $W_c \in \mathbb{R}^{d \times 1}$ and b are learnable parameters. s is the prediction score. During training, we maximize the margin between the positive and negative QA-pairs (*i.e.*, s^p and s^n respectively) with the hinge loss: $\mathcal{L} = \sum_{i=1}^{|\mathcal{A}_{mc}|} \max(0, 1 + s_i^n - s_i^p)$, where $|\mathcal{A}_{mc}|$ is the number of choices in a question.

For open-ended QA, as the number of categories (answers) are large, we empirically find that it is better to concatenate the global question feature f_Q with the video feature f_V before the classifier.

$$s = \sigma(W_c'(\text{ELU}(W_{qv}[f_Q; f_V])) + b), \quad (13)$$

in which $W_{qv} \in \mathbb{R}^{2d \times d}$, $W_c \in \mathbb{R}^{d \times |\mathcal{A}_{oe}|}$, and $b^{|\mathcal{A}_{oe}|}$ are learnable parameters. Besides, $|\mathcal{A}_{oe}|$ is the size of the predefined answer set. During training, the optimization is achieved by minimizing the cross-entropy loss: $\mathcal{L} = -\sum_{i=1}^{|\mathcal{A}_{oe}|} y_i \log s_i$, where s_i is the prediction score for the i^{th} sample. $y_i = 1$ if the answer index corresponds to the i^{th} sample's ground-truth answer and 0 otherwise.

Experiments

Datasets

We experiment on four VideoQA datasets that challenge the various aspects of video understanding: **TGIF-QA** (Jang et al. 2019) features questions about action repetition, state transition and frame QA. Action repetition includes the sub-tasks of repetition counting and repeating action recognition. In this work, we experiment on the latter since the prediction of numbers are hard to explain. **MSRVTT-QA** and **MSVD-QA** mainly challenge a recognition of video elements. Their question-answer pairs are automatically generated from the respective video descriptions by (Xu et al. 2017). **NExT-QA** (Xiao et al. 2021) is a challenging benchmark that goes beyond superficial video description to emphasize causal and temporal multi-object interactions. It is rich in object relations in space-time (Shang et al. 2019). The dataset has both multi-choice QA and generation-based QA. In this work, we focus on the former and leave the generation-based QA for future exploration. For all datasets, we report accuracy (percentage of correctly answered questions) as the evaluation metric. Other statistical details are given in Appendix A.

Implementation Details

We extract each video in NExT-QA, MSVD-QA and MSRVTT-QA at $p = \{6, 15, 15\}$ frames per second respectively. For TGIF-QA, all the frames are used. Based on the average video lengths, we uniformly sample $K = \{16, 8, 8, 4\}$ clips for each video in the four datasets respectively, while fixing the clip length $L = 16$. The sparse stream is obtained by evenly sampling with $\gamma = 0.25$. For each frame in the sparse stream, we detect $N = 20$ regions for

NExT-QA and 10 for the others. The dimension of the models' hidden states is $d = 512$ and the default number of graph layers in QGA is $H = 2$. For training, we adopt a two-stage scheme by firstly training the model with learning rate $lr = 10^{-4}$ and then fine-tune the best model obtained in the 1st stage with a smaller lr , *e.g.*, 5×10^{-5} . For both stages, we train the models by using Adam optimizer with batch size of 64 and maximum epoch of 25. Other details are presented in Appendix B.

The State of the Art Comparison

In Table 1 and Table 2, we compare our model with some established VideoQA techniques covering 4 major categories: **1)** cross-attention (*e.g.*, ST-VQA (Jang et al. 2017), PSAC (Li et al. 2019b), STA (Gao et al. 2019), MIN (Jin et al. 2019) and QueST (Jiang et al. 2020)), **2)** motion-appearance memory (*e.g.*, AMU (Xu et al. 2017), Co-Mem (Gao et al. 2018) and HME (Fan et al. 2019)), **3)** graph-structured models (*e.g.*, L-GCN² (Huang et al. 2020), HGA (Jiang and Han 2020), DualVGR (Wang, Bao, and Xu 2021), GMIN (Gu et al. 2021) and B2A (Park, Lee, and Sohn 2021)) and **4)** hierarchical models (*e.g.*, HCRN (Le et al. 2020) and HOSTR (Dang et al. 2021)). The results show that our Hierarchical QGA (HQGA) model performs consistently better than the others on all the experimented datasets.

Particularly, both L-GCN and GMIN are graph-based methods that focus on leveraging object-level information (similar to us) for question-answering. However, they model the object either in a monolithic way in space-time or in trajectories. This neither reflect the hierarchical nor the compositional nature of the video elements. Furthermore, their graphs are constructed without the guidance of language queries. By filling such gaps, our model shows clear superiority to both methods on the experimented datasets.

HCRN and HOSTR are similar to us in designing hierarchical conditional architectures. Nonetheless, HCRN is limited to hierarchical temporal relations between frames, in which the relations are modeled by simple average pooling. It is helpful for identifying repeated actions and state transition of single object (see results on TGIF-QA), but it is insufficient to understand more complicated object interactions in space-time. As a result, it performs even worse than L-GCN on NExT-QA. HOSTR advances HCRN by building the hierarchy over object trajectories and adopting graph operation for relation reasoning. Yet still, it focuses merely on designing general-purpose neural building blocks and lacks the bottom-up and top-down insight (Figure 1) for VideoQA, which results in its sub-optimal model design and results.

While HCRN and HOSTR use global query representations, QueST breaks down the question into spatial and temporal components, and designs separated attention modules to aggregate information from video for question answering. It obtains good results on TGIF-QA but does not generalize well to MSRVTT-QA and MSVD-QA where the questions do not have such spatial and temporal syntactic structure. Finally, the methods HGA, DualVGR and B2A, like us, try to

²L-GCN's results on NExT-QA and MSRVTT-QA are reproduced by us with the official code.

Models	NExT-QA Val				NExT-QA Test			
	Causal	Temporal	Descriptive	Overall	Causal	Temporal	Descriptive	Overall
ST-VQA	44.76	49.26	55.86	47.94	45.51	47.57	54.59	47.64
Co-Mem	45.22	49.07	55.34	48.04	45.85	<u>50.02</u>	54.38	48.54
HME	46.18	48.20	58.30	48.72	46.76	48.89	57.37	49.16
L-GCN	45.15	50.37	55.98	48.52	47.85	48.74	56.51	49.54
HGA	<u>46.26</u>	<u>50.74</u>	<u>59.33</u>	<u>49.74</u>	<u>48.13</u>	49.08	<u>57.79</u>	<u>50.01</u>
HCRN	45.91	49.26	53.67	48.20	47.07	49.27	54.02	48.89
HQGA (Ours)	48.48	51.24	61.65	51.42	49.04	52.28	59.43	51.75

Table 1: Comparison of accuracy. The best and second-best results are highlighted in bold and underline respectively.

Models	TGIF-QA			MSRV	MSVD
	Action	Transition	FrameQA	TT-QA	-QA
ST-VQA	62.9	69.4	49.50	30.9	31.3
PSAC	70.4	76.9	55.7	-	-
STA	72.3	79.0	56.6	-	-
MIN	72.7	80.9	57.1	35.4	35.0
QueST	<u>75.9</u>	81.0	<u>59.7</u>	34.6	36.1
AMU	-	-	-	32.5	32.0
Co-Mem	68.2	74.3	51.5	31.9	31.7
HME	73.9	77.8	53.8	33.0	33.7
L-GCN	74.3	81.1	56.3	33.7	34.3
HGA	75.4	81.0	55.1	35.5	34.7
DualVGR	-	-	-	35.5	39.0
GMIN	73.0	81.7	57.5	36.1	35.4
B2A	<u>75.9</u>	82.6	57.5	<u>36.9</u>	37.2
HCRN	75.0	81.4	55.9	35.6	36.1
HOSTR	75.0	<u>83.0</u>	58.0	35.9	<u>39.4</u>
HQGA	76.9	85.6	61.3	38.6	41.2

Table 2: Comparison of accuracy.

align words in the language query with their visual correspondences in the video. However, all of them adopt a flat way of alignment at segment level and lack the hierarchical structure, which most likely accounts for their inferiority.

Model Analysis

We analyze our model on the validation sets of NExT-QA and MSRVTT-QA. We first conduct an ablation study on the number of graph attention layers H , sampled video clips K and regions N to find the optimal settings on the two datasets (refer to Appendix C for more details). Then, we fix the optimal settings for subsequent experiments.

Hierarchy. The top section of Table 3 shows that accuracy drops by 0.9% on both datasets when removing the QGA units at the bottom (*w/o* G_O) (F_a and F_m are used as respective inputs for G_F and G_C). The results demonstrate the importance of modeling the static picture of object interactions. Taking away the 2nd level graph units (*w/o* G_F) by directly applying G_C over G_O 's outputs corresponding to the respective middle frames of the clips, we can observe even more drastic accuracy drops (over 1.2%) on both datasets. This results demonstrate the critical role of G_F in modeling the short-term dynamics of object interactions. Finally, when we remove both G_O and G_F and apply a monolithic graph over the clips (directly using F_m as inputs for G_C), the accuracy drops sharply by 1.5% and 2.6% on NExT-QA and MSRVTT-QA respectively. This clearly shows that only modeling clip-level information is unsatisfactory. The comparisons confirm the significance of hierarchically weaving together video elements of different levels.

Graph. The middle section of Table 3 shows that substituting the top QGA (G_C) with a sum-pooling over the corresponding input nodes (*w/o* $G_C(s)$), leads to accuracy drops from HQGA of 0.68% and 0.54% on NExT-QA and MSRVTT-QA respectively. The result validates the importance of G_C in reasoning over local, short-term dynamic interactions. Replacing the middle QGAs (G_F) with sum-poolings (*ss*) further degrades the performance on both datasets. The result indicates that a sum-pooling is neither sufficient in capturing the short-term dynamics of object interactions nor capable of relation reasoning over them. Thus, it evinces the strengths of both G_C and G_F . Finally, by replacing the last (bottom) QGAs (G_O) with sum-poolings (*sss*), we can observe a further exacerbation of performances on the basis of the previous two ablations. The results demonstrate G_O 's superiority in capturing object relations in static frames. This experiment demonstrates the advantage of graph attention over a simple sum-pooling.

Multi-level Condition. As shown in Table 3 (bottom part), the results in the first three rows show that removing the language condition at any single level jeopardize the overall performance, indicating the necessity of injecting the query cues at multiple levels. Specially, we investigate replacing the token-wise query representations Q with a global one f_Q , *i.e.*, by concatenating f_Q with the respective graph nodes at all levels. From the results shown in the row *w/f_Q*, we find that a global condition can slightly boost the performance on MSRVTT-QA compared with the model variant without any conditioning (*e.g.*, 37.52% vs. 37.03%). However, such benefit disappears when the model is extrapolated to the scenario where the videos and questions are much more complex (*e.g.*, NExT-QA). Finally, we conduct additional ablation studies on F_a and F_m that serve as global contexts to enhance the representations of graph nodes at the corresponding levels. The results show that both features help a bit to the overall performance.

Discussion. By jointly considering the ablation results of graph hierarchy and multi-level condition in Table 3, we can see that, compared with the graph hierarchy, the multi-level query condition has relatively smaller influence on the performance. We speculate that some questions are too simple to provide meaningful referring clues to the video contents (*e.g.*, 'what is happening'). Such questions purely rely the model to reason the video contents for answers.

Another finding is that, both the graph hierarchy and multi-level condition have relatively smaller performance gains on NExT-QA than on MSRVTT-QA. A possible reason could be that NExT-QA emphasizes causal and tem-

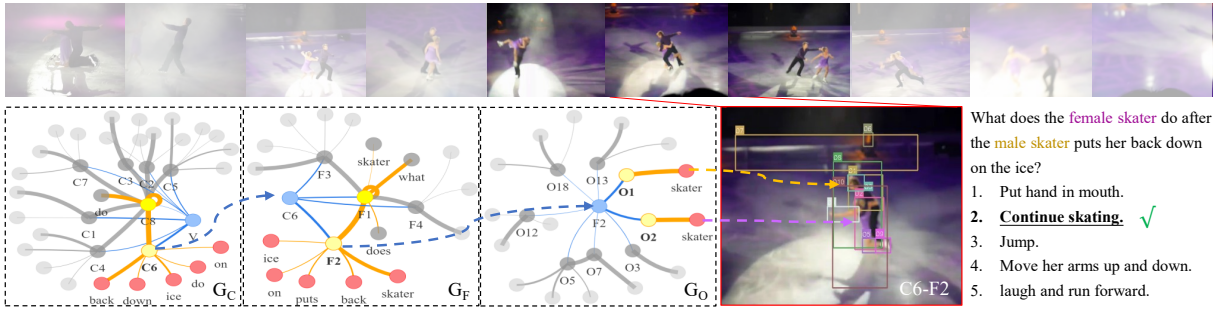


Figure 4: A correct prediction case from NExT-QA (Xiao et al. 2021). G_C , G_F , G_O show the learned graph and conditional attention weights at different levels. The nodes and edges that response relatively stronger to the aggregated nodes are highlighted. (Blue: weights of self-attention pooling β . Orange: weights of adjacency matrix A and query condition α .)

Model Variants	NExT-QA	MSRVTT-QA
HQGA	51.42	38.23
w/o G_O	50.50	37.26
w/o G_F	50.00	37.05
w/o G_O & G_F	49.96	35.66
w/o G_C (s)	50.74	37.69
w/o G_C & G_F (ss)	50.44	36.94
w/o G_C & G_F & G_O (sss)	50.32	35.88
w/o Q_C	51.30	38.17
w/o Q_C & Q_F	51.08	37.62
w/o Q_C & Q_F & Q_O	50.62	37.03
w/ f_Q	50.16	37.52
w/o F_m	50.90	37.94
w/o F_a & F_m	50.34	37.86

Table 3: Model ablation results on the validation sets.

poral action relations; it requires high-quality action recognition. Yet recognizing actions in such complex multiple-object scenarios remains a significant challenge in video understanding (Gu et al. 2018; Feichtenhofer et al. 2019). Our model can reason on the relations between video elements at multi-granularity levels. However, such capability seriously relies on object appearance features in its current version; the absence of region-level motion is a limitation. Our further analyses of model performances per question type on MSRVTT-QA and MSVD-QA (see Table 4) show that there are still clear gaps between action/activity recognition and object/attribute recognition in videos; the gaps remain 4.6% on MSRVTT-QA and 7.5% on MSVD-QA for ‘what’ questions. A promising solution would be to jointly model both the appearance and motion for each object. As it is not the focus of this work, we leave it for future exploration.

Intriguingly, the results of our simplest model variant (w/o G_O & G_F) are still on par with some previous SOTAs. Such strong results can be attributed to our better data representation. Here to verify BERT, we additionally explore substituting BERT with GloVe (Pennington, Socher, and Manning 2014) and achieve 37.2% on MSRVTT-QA test set. This result confirms the advantages of BERT (38.6% vs. 37.2%) as a good contextualized representation to fulfill the multi-granular condition, e.g., the disambiguation between the male and female skaters denoted as O_1 and O_2 respectively in Figure 4 (G_O). Also, the result prompts future exploration of finetuning pre-trained vision-text architectures (Lei et al. 2021) for potential improvement of performance.

Datasets	what _a	what _o	what	who	how	when	where	all
MSRVTT	30.1	34.7	32.5	48.9	81.5	78.3	38.4	38.6
MSVD	25.4	32.9	30.4	57.2	76.2	75.9	32.1	41.2

Table 4: Test accuracy per question type. Roughly $\frac{1}{2}$ ($\frac{1}{3}$) of the ‘what’ questions in MSRVTT-QA (MSVD-QA) ask actions/activities; others are about objects/attributes. We distinguish them via the pattern ‘what ... doing’ in questions.

Qualitative Analysis. We show a prediction case in Figure 4 (find more examples in Appendix C). Firstly, by tracing down the self-attention pooling weights β , our model precisely finds the video contents that are relevant to the textual query, e.g., from the 6th video clip C_6 to its 2nd frame F_2 , and further to the male and female skaters (denoted as O_1 and O_2 respectively in F_2). Secondly, according to the query-conditional weights α , our model successfully differentiates the information of different granularity levels for both the video and question contents, and discriminatively match them at the corresponding levels. For example, the nodes at high level G_C show stronger responses to the action-related words (‘do’ and ‘back down’), whereas those nodes at the lower levels (G_F and G_O) respond strongly to the visual objects (‘skater’). Finally, according to the learned adjacency matrices, while we construct fully-connected graphs in QGA, the learned connections are quite sparse, suggesting that our model can learn to filter out meaningless relations with respect to the query.

Conclusion

This work delves into video question answering and uncovers the insights of bottom-up and top-down for video-language alignment. To capture the insight, we propose to build video as a conditional graph hierarchy which level wisely reasons and aggregates low level visual resources into high level video elements, in which the language queries are injected into different levels to match and pinpoint the video elements at multi-granularity. To accomplish this, we design a reusable query-conditioned graph attention unit and stack it to achieve the hierarchical architecture. Our extensive experiments and analyses have validated the effectiveness of the proposed method. Future attempts can be made on incorporating object-level motion information, or exploiting pretraining techniques to boost the performance.

Acknowledgements

This research is supported by the Sea-NExT Joint Lab.

References

- Anderson, P.; He, X.; Buehler, C.; Teney, D.; Johnson, M.; Gould, S.; and Zhang, L. 2018. Bottom-up and top-down attention for image captioning and visual question answering. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 6077–6086.
- Carreira, J.; and Zisserman, A. 2017. Quo vadis, action recognition? a new model and the kinetics dataset. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 6299–6308.
- Chen, S.; Zhao, Y.; Jin, Q.; and Wu, Q. 2020. Fine-grained video-text retrieval with hierarchical graph reasoning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 10638–10647.
- Cho, K.; Van Merriënboer, B.; Gulcehre, C.; Bahdanau, D.; Bougares, F.; Schwenk, H.; and Bengio, Y. 2014. Learning phrase representations using RNN encoder-decoder for statistical machine translation. *arXiv preprint arXiv:1406.1078*.
- Dang, L. H.; Le, T. M.; Le, V.; and Tran, T. 2021. Hierarchical Object-oriented Spatio-Temporal Reasoning for Video Question Answering. In *Proceedings of the 30th International Joint Conference on Artificial Intelligence (IJCAI)*.
- Devlin, J.; Chang, M.-W.; Lee, K.; and Toutanova, K. 2018. Bert: Pre-training of deep bidirectional transformers for language understanding. *arXiv preprint arXiv:1810.04805*.
- Fan, C.; Zhang, X.; Zhang, S.; Wang, W.; Zhang, C.; and Huang, H. 2019. Heterogeneous memory enhanced multimodal attention model for video question answering. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 1999–2007.
- Feichtenhofer, C.; Fan, H.; Malik, J.; and He, K. 2019. Slowfast networks for video recognition. In *Proceedings of the IEEE/CVF international conference on computer vision (ICCV)*, 6202–6211.
- Gao, J.; Ge, R.; Chen, K.; and Nevatia, R. 2018. Motion-appearance co-memory networks for video question answering. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 6576–6585.
- Gao, L.; Zeng, P.; Song, J.; Li, Y.-F.; Liu, W.; Mei, T.; and Shen, H. T. 2019. Structured two-stream attention network for video question answering. In *Proceedings of the AAAI Conference on Artificial Intelligence (AAAI)*, volume 33, 6391–6398.
- Gkioxari, G.; Girshick, R.; fDollár, P.; and He, K. 2018. Detecting and recognizing human-object interactions. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 8359–8367.
- Gu, C.; Sun, C.; Ross, D. A.; Vondrick, C.; Pantofaru, C.; Li, Y.; Vijayanarasimhan, S.; Toderici, G.; Ricco, S.; Sukthankar, R.; et al. 2018. Ava: A video dataset of spatio-temporally localized atomic visual actions. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 6047–6056.
- Gu, M.; Zhao, Z.; Jin, W.; Hong, R.; and Wu, F. 2021. Graph-Based Multi-Interaction Network for Video Question Answering. *IEEE Transaction on Image Processing* 30: 2758–2770.
- Hara, K.; Kataoka, H.; and Satoh, Y. 2018. Can spatiotemporal 3d cnns retrace the history of 2d cnns and imagenet? In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 6546–6555.
- He, K.; Zhang, X.; Ren, S.; and Sun, J. 2016. Deep residual learning for image recognition. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 770–778.
- Hu, R.; Andreas, J.; Darrell, T.; and Saenko, K. 2018. Explainable neural computation via stack neural module networks. In *European Conference on Computer Vision (ECCV)*, 53–69.
- Hu, R.; Rohrbach, A.; Darrell, T.; and Saenko, K. 2019. Language-conditioned graph networks for relational reasoning. In *International Conference on Computer Vision and Pattern Recognition (CVPR)*, 10294–10303.
- Hu, R.; Rohrbach, M.; Andreas, J.; Darrell, T.; and Saenko, K. 2017. Modeling relationships in referential expressions with compositional modular networks. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 1115–1124.
- Huang, D.; Chen, P.; Zeng, R.; Du, Q.; Tan, M.; and Gan, C. 2020. Location-aware graph convolutional networks for video question answering. In *Proceedings of the AAAI Conference on Artificial Intelligence (AAAI)*, volume 34, 11021–11028.
- Jang, Y.; Song, Y.; Kim, C. D.; Yu, Y.; Kim, Y.; and Kim, G. 2019. Video question answering with spatio-temporal reasoning. *International Journal of Computer Vision (IJCV)* 127(10): 1385–1412.
- Jang, Y.; Song, Y.; Yu, Y.; Kim, Y.; and Kim, G. 2017. TGIF-QA: Toward Spatio-Temporal Reasoning in Visual Question Answering. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*.
- Jiang, J.; Chen, Z.; Lin, H.; Zhao, X.; and Gao, Y. 2020. Divide and conquer: Question-guided spatio-temporal contextual attention for video question answering. In *Proceedings of the AAAI Conference on Artificial Intelligence (AAAI)*, volume 34, 11101–11108.
- Jiang, P.; and Han, Y. 2020. Reasoning with heterogeneous graph alignment for video question answering. In *Proceedings of the AAAI Conference on Artificial Intelligence (AAAI)*, volume 34, 11109–11116.
- Jin, W.; Zhao, Z.; Gu, M.; Yu, J.; Xiao, J.; and Zhuang, Y. 2019. Multi-interaction network with object relation for

- video question answering. In *Proceedings of the 27th ACM International Conference on Multimedia*, 1193–1201. ACM.
- Kipf, T. N.; and Welling, M. 2017. Semi-supervised classification with graph convolutional networks. In *International Conference on Learning Representations (ICLR)*.
- Le, T. M.; Le, V.; Venkatesh, S.; and Tran, T. 2020. Hierarchical conditional relation networks for video question answering. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 9972–9981.
- Lee, J.; Lee, I.; and Kang, J. 2019. Self-attention graph pooling. In *International Conference on Machine Learning (ICML)*, 3734–3743.
- Lei, J.; Li, L.; Zhou, L.; Gan, Z.; Berg, T. L.; Bansal, M.; and Liu, J. 2021. Less is more: Clipbert for video-and-language learning via sparse sampling. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 7331–7341.
- Li, L.; Gan, Z.; Cheng, Y.; and Liu, J. 2019a. Relation-aware graph attention network for visual question answering. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, 10313–10322.
- Li, X.; Song, J.; Gao, L.; Liu, X.; Huang, W.; He, X.; and Gan, C. 2019b. Beyond rnns: Positional self-attention with co-attention for video question answering. In *Proceedings of the AAAI Conference on Artificial Intelligence (AAAI)*, volume 33, 8658–8665.
- Mao, J.; Huang, J.; Toshev, A.; Camburu, O.; Yuille, A. L.; and Murphy, K. 2016. Generation and comprehension of unambiguous object descriptions. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 11–20.
- Norcliffe-Brown, W.; Vafeias, E.; and Parisot, S. 2018. Learning conditioned graph structures for interpretable visual question answering. In *Advances in Neural Information Processing Systems (NeurIPS)*, 8344–8353.
- Park, J.; Lee, J.; and Sohn, K. 2021. Bridge to Answer: Structure-aware Graph Interaction Network for Video Question Answering. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 15526–15535.
- Pennington, J.; Socher, R.; and Manning, C. D. 2014. Glove: Global vectors for word representation. In *Proceedings of the 2014 conference on empirical methods in natural language processing (EMNLP)*, 1532–1543.
- Ren, S.; He, K.; Girshick, R.; and Sun, J. 2015. Faster r-cnn: Towards real-time object detection with region proposal networks. In *Advances in Neural Information Processing Systems (NeurIPS)*, volume 28, 91–99.
- Shang, X.; Di, D.; Xiao, J.; Cao, Y.; Yang, X.; and Chua, T.-S. 2019. Annotating objects and relations in user-generated videos. In *Proceedings of the 2019 on International Conference on Multimedia Retrieval (ICMR)*, 279–287.
- Veličković, P.; Cucurull, G.; Casanova, A.; Romero, A.; Liò, P.; and Bengio, Y. 2018. Graph Attention Networks. In *International Conference on Learning Representations (ICLR)*.
- Wang, J.; Bao, B.; and Xu, C. 2021. DualVGR: A Dual-Visual Graph Reasoning Unit for Video Question Answering. *IEEE Transaction on Multimedia*.
- Xiao, J.; Shang, X.; Yang, X.; Tang, S.; and Chua, T.-S. 2020. Visual relation grounding in videos. In *European Conference on Computer Vision (ECCV)*, 447–464. Springer.
- Xiao, J.; Shang, X.; Yao, A.; and Chua, T.-S. 2021. NEX-TQA: Next Phase of Question-Answering to Explaining Temporal Actions. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 9777–9786.
- Xu, D.; Zhao, Z.; Xiao, J.; Wu, F.; Zhang, H.; He, X.; and Zhuang, Y. 2017. Video question answering via gradually refined attention over appearance and motion. In *International Conference on Multimedia*, 1645–1653. ACM.
- Yu, Y.; Kim, J.; and Kim, G. 2018. A joint sequence fusion model for video question answering and retrieval. In *Proceedings of the European Conference on Computer Vision (ECCV)*, 471–487.
- Zeng, K.-H.; Chen, T.-H.; Chuang, C.-Y.; Liao, Y.-H.; Niebles, J. C.; and Sun, M. 2017. Leveraging video descriptions to learn video question answering. In *Proceedings of the AAAI Conference on Artificial Intelligence (AAAI)*, 4334–4340.