

Hierarchical Object-oriented Spatio-Temporal Reasoning for Video Question Answering

Long Hoang Dang, Thao Minh Le, Vuong Le and Truyen Tran

Applied Artificial Intelligence Institute, Deakin University, Australia

{hldang, lethao, vuong.le, truyen.tran}@deakin.edu.au

Abstract

Video Question Answering (Video QA) is a powerful testbed to develop new AI capabilities. This task necessitates learning to reason about objects, relations, and events across visual and linguistic domains in space-time. High-level reasoning demands lifting from associative visual pattern recognition to symbol-like manipulation over objects, their behavior and interactions. Toward reaching this goal we propose an object-oriented reasoning approach in that video is abstracted as a dynamic stream of interacting objects. At each stage of the video event flow, these objects interact with each other, and their interactions are reasoned about with respect to the query and under the overall context of a video. This mechanism is materialized into a family of general-purpose neural units and their multi-level architecture called Hierarchical Object-oriented Spatio-Temporal Reasoning (HOSTR) networks. This neural model maintains the objects' consistent lifelines in the form of a hierarchically nested spatio-temporal graph. Within this graph, the dynamic interactive object-oriented representations are built up along the video sequence, hierarchically abstracted in a bottom-up manner, and converge toward the key information for the correct answer. The method is evaluated on multiple major Video QA datasets and establishes new state-of-the-arts in these tasks. Analysis into the model's behavior indicates that object-oriented reasoning is a reliable, interpretable and efficient approach to Video QA.

1 Introduction

Much of the recent impressive progress of AI can be attributed to the availability of suitable large-scale testbeds. A powerful testbed – largely under-explored – is Video Question Answering (Video QA). This task demands a wide range of cognitive capabilities including learning and reasoning about objects and dynamic relations in space-time, in both visual and linguistic domains. A major challenge of reasoning over video is extracting question-relevant high-level facts from low-level moving pixels over an extended period of time. These facts include objects, their motion profiles, actions, interactions,

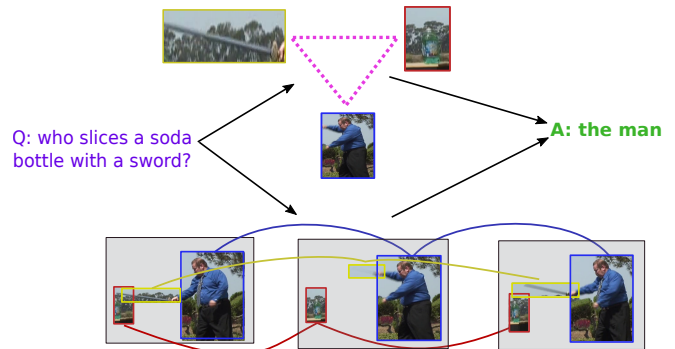


Figure 1: The key to Video QA is the effective relational reasoning on objects with their temporal lifelines (below sequences) interleaved with spatial interactions (upper graph) under the context set in the video and the perspective provided by the query. This object-oriented spatio-temporal reasoning is the main theme of this work.

events, and consequences distributed in space-time. Another challenge is to learn the long-term temporal relation of visual objects conditioning on the guidance clues from the question – effectively bridging the semantic gulf between the two domains. Finally, learning to reason from relational data is an open problem on its own, as it pushes the boundary of learning from simple one-step classification to dynamically construct question-specific computational graphs that realize the iterative reasoning process.

A highly plausible path to tackle these challenges is via *object-centric learning* since objects are fundamental to cognition [Spelke and Kinzler, 2007]. Objects pave the way towards more human-like reasoning capability and symbolic computing [Lake *et al.*, 2017; Spelke and Kinzler, 2007]. Unlike objects in static images, objects in video have unique evolving lives throughout space-time. As object lives throughout the video, it changes its appearance and position, and interacts with other objects at arbitrary time. When observed in the videos, all these behaviors play out on top of a background of rich context of the video scene. Furthermore, in the question answering setting, these object-oriented information must be considered from the specific view point set by the linguistic query. With these principles, we pinpoint the key to Video QA to be the effective high-level relational reasoning of spatio-temporal objects under the video context and the perspective

provided by the query (see Fig. 1). This is challenging for the complexity of the video spatio-temporal structure and the cross-domain compatibility gap between linguistic query and visual objects.

Toward such challenge we design a general-purpose neural unit called Object-oriented Spatio-Temporal Reasoning (OSTR) that operates on a set of video object sequences, a contextual video feature, and an external linguistic query. OSTR models object lifelong interactions and returns a summary representation in the form of a singular set of objects. The specialties of OSTR are in the partitioning of intra-object temporal aggregation and inter-object spatial interaction that leads to the efficiency of the reasoning process. Being flexible and generic, OSTR units are suitable building blocks for constructing powerful reasoning models.

For Video QA problem, we use OSTR units to build up Hierarchical Object-oriented Spatio-Temporal Reasoning (HOSTR) model. The network consists of OSTR units arranged in layers corresponding to the levels of video temporal structure. At each level, HOSTR finds local object interactions and summarizes them toward a higher-level, longer-term representation with the guidance of the linguistic query.

HOSTR stands out with its authentic and explicit modeling of video objects leading to the effective and interpretable reasoning process. The hierarchical architecture also allows the model to efficiently scale to a wider range of video formats and lengths. These advantages are demonstrated in a comprehensive set of experiments on multiple major Video QA tasks and datasets.

In summary, this paper makes three major contributions: (1) A semantic-rich object-oriented representation of videos that paves the way for spatio-temporal reasoning (2) A general-purpose neural reasoning unit with dynamic object interactions per context and query; and (3) A hierarchical network that produces reliable and interpretable video question answering.

2 Related Work

Video QA has been developed on top of traditional video analysis schemes such as recurrent networks of frame features [Zhao *et al.*, 2019] or 3D convolutional operators [Tran *et al.*, 2018]. Video representations are then fused with or gated by the linguistic query through co-attention [Jang *et al.*, 2017; Ye *et al.*, 2017], hierarchical attention [Liang *et al.*, 2018; Zhao *et al.*, 2018], and memory networks [Kim *et al.*, 2017; Wang *et al.*, 2019]. More recent works advance the field by exploiting hierarchical video structure [Le *et al.*, 2020a] or separate reasoning out of representation learning [Le *et al.*, 2020b]. A share feature between these works is considering the whole video frames or segments as the unit component of reasoning. In contrast, our work make a step further by using detail objects from the video as primitive constructs for reasoning.

Object-centric Video Representation inherits the modern capability of object detection on images [Desta *et al.*, 2018] and continuous tracking through temporal consistency [Wojke *et al.*, 2017]. Tracked objects form tubelets [Kalogeiton *et al.*, 2017] whose representation contributes to breakthroughs in action detection [Xie *et al.*, 2018] and event segmentation

[Chao *et al.*, 2018]. For tasks that require abstract reasoning, the connection between objects beyond temporal object permanence can be established through relation networks [Baradel *et al.*, 2018]. The concurrence of objects’ 2D spatial- and 1D temporal- relations naturally forms a 3D spatio-temporal graphs [Wang and Gupta, 2018]. This graph can be represented as either a single flattened one where all parts connect together [Zeng *et al.*, 2019], or separated spatial- and temporal-graphs [Pan *et al.*, 2020]. They can also be approximated as a dynamic graph where objects live through the temporal axis of the video while their properties and connection evolve [Jain *et al.*, 2016].

Object-based Video QA is still in infancy. The works in [Yang *et al.*, 2020] and [Huang *et al.*, 2020] extract object features and feed them to generic relational engines without prior structure of reasoning through space-time. At the other extreme, detected objects are used to scaffold the symbolic reasoning computation graph [Yi *et al.*, 2020] which is explicit but limited in flexibility and cannot recover from object extraction errors. Our work is a major step toward the object-centric reasoning with the balance between explicitness and flexibility. Here video objects serve as active agents which build up and adjust their interactions dynamically in the spatio-temporal space as instructed by the linguistic query.

3 Method

3.1 Problem Definition

Given a video V and linguistic question q , our goal is to learn a mapping function $F_\theta(\cdot)$ that returns a correct answer \bar{a} from an answer set A as follows:

$$\bar{a} = \arg \max_{a \in A} F_\theta(a | q, V). \quad (1)$$

In this paper, a video V is abstracted as a collection of object sequences tracked in space and time. Function F_θ is designed to have the form of a hierarchical object-oriented network that takes the object sequences, modulates them with the overall video context, dynamically infers object interactions as instructed by the question q so that key information regarding a arises from the mix. This object-oriented representation is presented next in Sec. 3.2, followed by Sec. 3.3 describing the key computation unit and Sec. 3.4 the resulting model.

3.2 Data Representation

Video as a Set of Object Sequences

Different from most of the prominent VideoQA methods [Jang *et al.*, 2017; Gao *et al.*, 2018; Le *et al.*, 2020a] where videos are represented by frame features, we break down a video of length L into a list of N object sequences $O = \{o_{n,t}\}_{n=1,t=1}^{N,L}$ constructed by chaining corresponding objects of the same identity n across the video. These objects are represented by a combination of (1) appearance features (RoI pooling features) $o_{n,t}^a \in \mathbb{R}^{2048}$ (representing “what”); and (2) the positional features $o_{n,t}^p = [\frac{x_{min}}{W}, \frac{y_{min}}{H}, \frac{x_{max}}{W}, \frac{y_{max}}{H}, \frac{w}{W}, \frac{h}{H}, \frac{wh}{WH}]$ (“where”), with w and h are the sizes of the bounding box and W, H are those of the video frame, respectively.

In practice, we use Faster R-CNN with ROI-pooling to extract positional and appearance features; and DeepSort for

multi-object tracking. We assume that the objects live from the beginning to the end of the video. Occluded or missed objects have their features marked with special null values which will be specially dealt with by the model.

Joint Encoding of “What” and “Where”

Since the appearance $o_{n,t}^a$ of an object may remain relatively stable over time while $o_{n,t}^p$ constantly changes, we must find *joint positional-appearance features* of objects to make them discriminative in both space and time. Specifically, we propose the following multiplicative gating mechanism to construct such features:

$$o_{n,t} = f_1(o_{n,t}^a) \odot f_2(o_{n,t}^p) \in \mathbb{R}^d, \quad (2)$$

where $f_2(o_{n,t}^p) \in (0, 1)$ serves as a position gate to (softly) turn on/off the localized appearance features $f_1(o_{n,t}^a)$. We choose $f_1(x) = \tanh(W_a x + b_a)$ and $f_2(x) = \text{sigmoid}(W_p x + b_p)$, where W_a and W_p are network weights with height of d .

Along with the object sequences, we also maintain global features of video frames which hold the information of the background scene and possible missed objects. Specifically, for each frame t , we form the the global features g_t as the combination of the frame’s appearance features (pretrained ResNet *pool5* vectors) and motion feature (pretrained ResNeXt-101) extracted from such frame.

With these ready, the video is represented as a tuple of object sequences O_n and frame-wise global features g_t : $\mathcal{V} = (\{O_n \mid O_n \in \mathbb{R}^{L \times d}\}_{n=1}^N, \{g_t\}_{t=1}^L)$.

Linguistic Representation

We utilize a BiLSTM running on top of GloVe embedding of the words in a query of length S to generate contextual embeddings $\{e_s\}_{s=1}^S$ for $e_s \in \mathbb{R}^d$, which share the dimension d with object features. We also maintain a global representation of the question by summarizing the two end LSTM states $q_g \in \mathbb{R}^d$. We further use q_g to drive the attention mechanism and combine contextual words into a unified query representation $q = \sum_{s=1}^S \alpha_s e_s$ where $\alpha_s = \text{softmax}_s(W_q(e_s \odot q_g))$.

3.3 Object-oriented Spatio-Temporal Reasoning (OSTR)

With the videos represented as object sequences, we need to design a scalable reasoning framework that can work natively on the structures. Such a framework must be modular so it is flexible to different input formats and sizes. Toward this goal, we design a generic reasoning unit called *Object-oriented Spatio-Temporal Reasoning (OSTR)* that operates on this object-oriented structure and supports layering and parallelism.

Algorithmically, OSTR takes as input a query representation q , a context representation c , a set of N object sequences $X = \{X_n \mid X_n \in \mathbb{R}^{T \times d}\}_{n=1}^N$ of equal length T , and individual identities $\{n\}$. In practice, X can be a subsegment of the whole object sequences O , and c is gathered from the frame features g_t constructed in Sec. 3.2. The output of the OSTR is a set of object instances of the same identity.

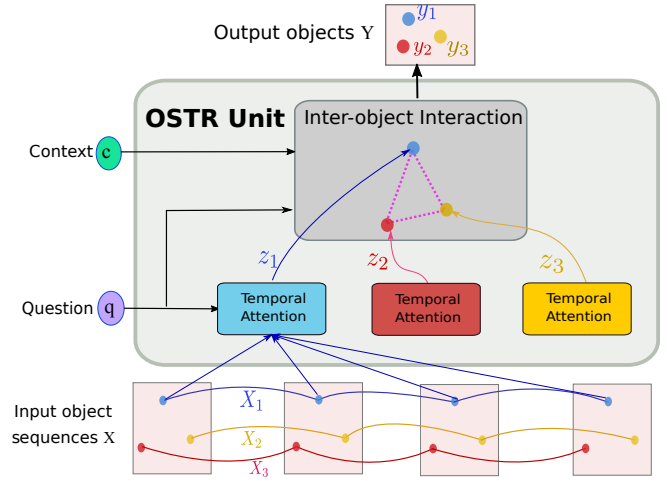


Figure 2: Object-oriented Spatio-Temporal Reasoning (OSTR) unit. Inputs include a set of object sequences X (with identity indicated by colors), a context c and a query q . Each object sequence is first summarized by a temporal attention module (matching color boxes). The inter-object S-T relations is modeled by a graph-based spatial interaction (gray box with pink graph). OSTR unit outputs a set of object instances $\{y_n\}$ with object IDs corresponding to those in the input sequences.

Across the space-time domains, real-world objects have distinctive properties (appearance, position, etc.) and behaviors (motion, deformation, etc.) throughout their lives. Meanwhile, different objects living in the same period can interact with each other. The OSTR closely reflects this nature by containing the two main components: (1) Intra-object temporal attention and (2) Inter-object interaction (see Fig. 2).

Intra-object Temporal Attention

The goal of the temporal attention module is to produce a *query-specific* summary of each object sequence $X_n = \{x_{n,t}\}_{t=1}^T$ into a single vector z_n . The attention weights are driven by the query q to reflect the fact that the relevance to the query varies across the sequence. In details, the summarized vector z_n is calculated by

$$z_n = \text{temporal_attention}(X_n) := \gamma * \sum_{t=1}^T \beta_t x_{n,t}, \text{ where } (3)$$

$$\beta_t = \text{softmax}_t(W_a((W_q q + b_q) \odot (W_x x_t + b_x))), \quad (4)$$

where \odot is the Hadamard product, $\{W_a, W_q, W_x\}$ are learnable weights, γ is a binary mask vector to handle the null values caused by missed detections as mentioned in Sec. 3.2.

When the sequential structure is particularly strong, we can optionally employ a BiLSTM to model the sequence. We can then either utilize the last state of the forward LSTM and the first state of the backward LSTM, or place attention across the hidden states instead of object feature embeddings.

Inter-object Interaction

Fundamentally, the lifelines of object sequences are not only described by their internal behavior through time but also by the interactions with their neighbor objects coexisting in the

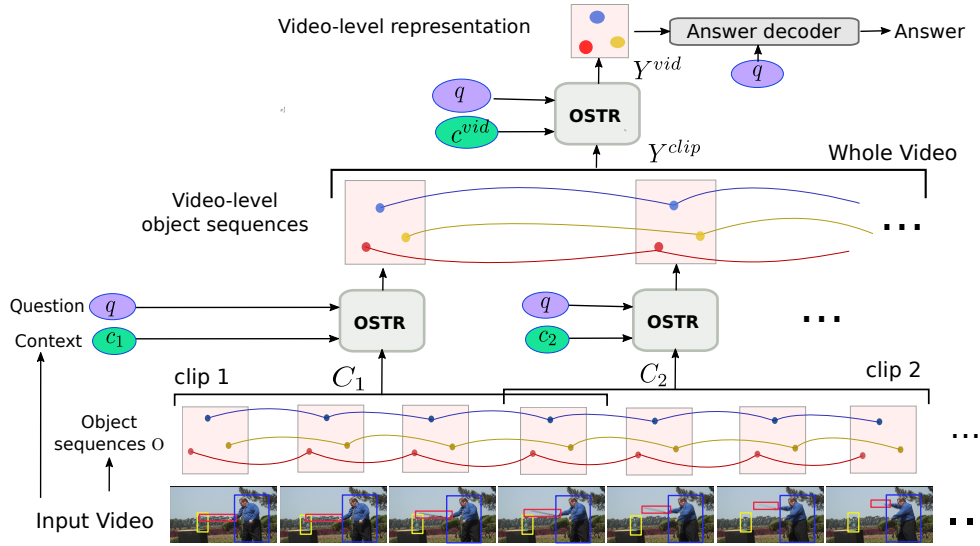


Figure 3: The architecture of Hierarchical Object-oriented Spatio-Temporal Reasoning (HOSTR) network for Video QA. HOSTR contains OSTR units operating at two levels: clip-level and video-level. Clip-level OSTR units model the interaction between object chunks within a particular clip under the modulation of the question and a clip-specific context representation. Their output objects are chained together and further sent to a video-level OSTR to capture the long-term dependencies between the objects existing in the whole video. Finally, a classifier taking as input the summarized output of video-level OSTR and the query is used for answer prediction.

same space. In order to represent such complex relationship, we build a spatio-temporal computation graph to facilitate the inter-object interactions modulated by the query q . This graph $\mathcal{G}(Z, E)$ contains vertices as the summarized objects $Z = \{z_n\}_{n=1}^N$ generated in Eq. 4, and the edges E represented by an adjacency matrix $A \in \mathbb{R}^{N \times N}$. A is calculated dynamically as the *query-induced correlation matrix* between the objects:

$$a_n = \text{norm}(W_a([z_n, z_n \odot q])), \quad (5)$$

$$A = a^\top a. \quad (6)$$

Here a_n is the relevance of object n w.r.t. the query q . The *norm* operator is implemented as a softmax function over objects in our implementation.

Given the graph \mathcal{G} , we use a Graph Convolutional Network (GCN) equipped with skip-connections to refine objects in relation with their neighboring nodes. Starting with the initialization $H^0 = (z_1, z_2, \dots, z_n) \in \mathbb{R}^{N \times d}$, the representations of nodes are updated through a number of refinement iterations. At iteration i , the new hidden states are calculated by:

$$\begin{aligned} \text{GCN}_i(H^{i-1}) &= W_2^{i-1} \sigma(AH^{i-1}W_1^{i-1} + b^{i-1}) \\ H^i &= \sigma(H^{i-1} + \text{GCN}_i(H^{i-1})), \end{aligned} \quad (7)$$

where $\sigma(\cdot)$ is a nonlinear activation (ELU in our implementation). After a fixed number of GCN iterations, the hidden states of the final layer are gathered as $H^{i_{max}} = \{h_n\}_{n=1}^N$.

To recover the underlying background scene information and compensate for possible undetected objects, we augment the object representations with the global context c :

$$y_n = \text{MLP}([h_n; c]). \quad (8)$$

These vectors line up to form the final output of the OSTR unit as a set of objects $Y = \{y_n\}_{n=1}^N$.

3.4 Hierarchical Object-oriented Spatio-Temporal Reasoning (HOSTR)

Even though the partitioning of temporal and spatial interaction in OSTR brings the benefits of efficiency and modularity, such separated treatment can cause the loss of spatio-temporal information, especially with long sequences. This limitation prevents us from using OSTR directly on the full video object sequences. To allow temporal and spatial reasonings to cooperate along the way, we break down a long video into multiple short (overlapping) clips and impose an hierarchical structure on top. With such division, the two types of interactions can be combined and interleaved across clips and allow full spatio-temporal reasoning.

Based on this motive, we design a novel hierarchical structure called Hierarchical Object-oriented Spatio-Temporal Reasoning (HOSTR) that follows the video multi-level structure and utilizes the OSTR units as building blocks. Our architecture shares the design philosophy of hierarchical reasoning structures with HCRN [Le *et al.*, 2020a] as well as the other general neural building blocks such as ResNet and InceptionNet. Thanks to the genericity of OSTR, we can build a hierarchy of arbitrary depth. For concreteness, we present here a two-layer HOSTR corresponding to the video structure: clip-level and video-level (see Fig. 3).

In particular, we first split all object sequences O constructed in Sec. 3.2 into K equal-sized chunks $C = \{C_1, C_2, \dots, C_K\}$ corresponding to the video clips, each of T frames. As the result, each chunk includes the object subsequences $C_k = \{o_{n,t}\}_{n=1, t=t_k}^{N, t_k+T}$, where $o_{n,t}$ are the object features extracted in Eq. 2, and t_k is the starting time of clip k .

Similarly, we divide the sequence of global frame features $\{g_t\}_{t=1}^L$ into K parts corresponding to the video clips. The

global context c_k for clip k is derived from each part by an identical operation with the temporal attention for objects in Eqs. 3,4: $c_k = \text{temporal_attention}(\{g_t\}_{t=t_k}^{t_k+T})$.

Clip-level OSTR units work on each of these subsequences C_k , context c_k^{clip} and query q , and generate the clip-level representation of the chunk $y_k^{\text{clip}} \in \mathbb{R}^{N \times d}$:

$$y_k^{\text{clip}} = \text{OSTR}(C_k, c_k^{\text{clip}}, q). \quad (9)$$

Outputs of the K clip-level OSTRs are K different sets of objects $y_k^{\text{clip}} = \{y_{n,k}\}_{n=1}^N$ whose identities n were maintained. Therefore, we can easily chain these objects of the same identity from different clips together to form the video-level sequence of objects $Y^{\text{clip}} = \{y_{n,k}^{\text{clip}}\}_{n=1,k=1}^{N,K}$.

At the video level, we have a single OSTR unit that takes in the object sequence Y^{clip} , query q , and video-level context c^{vid} . The context c^{vid} is again derived from the clip-level context c_k^{clip} by temporal attention: $c^{\text{vid}} = \text{temporal_attention}(\{c_k^{\text{clip}}\}_{k=1}^K)$.

The video-level OSTR models the long-term relationships between input object sequences in the whole video:

$$Y^{\text{vid}} = \text{OSTR}(Y^{\text{clip}}, c^{\text{vid}}, q).$$

The output of this unit is a set of N vectors $Y^{\text{vid}} = \{y_n^{\text{vid}} \mid y_n^{\text{vid}} \in \mathbb{R}^d\}_{n=1}^N$. The set is further summarized using an attention mechanism using the query q into the final representation vector r :

$$\delta_n = \text{softmax}_n(\text{MLP}[W_y y_n^{\text{vid}}; W_y y_n^{\text{vid}} \odot W_c q]), \quad (10)$$

$$r = \sum_{n=1}^N \delta_n y_n^{\text{vid}} \in \mathbb{R}^d. \quad (11)$$

3.5 Answer Decoders

We follow the common settings for answer decoders (e.g., see [Jang *et al.*, 2017]) which combine the final representation r with the query q using an MLP followed by a softmax to rank the possible answer choices. More details about the answer decoders per question types are available in the supplemental material. We use the cross-entropy as the loss function to training the model from end to end for all tasks except counting, where Mean Square Error is used.

4 Experiments

4.1 Datasets

We evaluate our proposed HOSTR on the three public video QA benchmarks, namely, TGIF-QA [Jang *et al.*, 2017], MSVD-QA [Xu *et al.*, 2017] and MSRVT-QA [Xu *et al.*, 2017]. More details are as follows.

MSVD-QA consists of 50,505 QA pairs annotated from 1,970 short video clips. The dataset covers five question types: *What*, *Who*, *How*, *When*, and *Where*, of which 61% of the QA pairs for training, 13% for validation and 26% for testing.

MSRVTT-QA contains 10K real videos (65% for training, 5% for validation, and 30% for testing) with more than 243K

Model	Test Accuracy (%)	
	MSVD-QA	MSRVTT-QA
ST-VQA	31.3	30.9
Co-Mem	31.7	32.0
AMU	32.0	32.5
HME	33.7	33.0
HCRN	36.1	35.4
HOSTR	39.4	35.9

Table 1: Experimental results on MSVD-QA and MSRVT-QA.

Model	TGIF-QA			
	Action \uparrow	Trans. \uparrow	Frame \uparrow	Count \downarrow
ST-TP (R+C)	62.9	69.4	49.5	4.32
Co-Mem (R+F)	68.2	74.3	51.5	4.10
PSAC (R)	70.4	76.9	55.7	4.27
HME (R+C)	73.9	77.8	53.8	4.02
HCRN (R)	70.8	79.8	56.4	4.38
HCRN (R+F)	75.0	81.4	55.9	3.82
HOSTR (R)	75.6	82.1	58.2	4.13
HOSTR (R+F)	75.0	83.0	58.0	3.65

Table 2: Experimental results on TGIF-QA dataset. R: ResNet, F: Flow, and C: C3D, respectively. MSE is used as the evaluation metric for count while accuracy is used for others.

question-answer pairs. Similar to MSVD-QA, questions are of five types: *What*, *Who*, *How*, *When*, and *Where*.

TGIF-QA is one of the largest Video QA datasets with 72K animated GIFs and 120K question-answer pairs. Questions cover four tasks - *Action*, *Event Transition*, *FrameQA*, *Count*. We refer readers to the supplemental material for the dataset description and statistics.

4.2 Comparison Against SOTAs

Implementation: We use Faster R-CNN¹ for frame-wise object detection. The number of object sequences per video for MSVD-QA, MSRVT-QA is 40 and TGIF-QA is 50. We embed question words into 300-D vectors and initialize them with GloVe during training. Default settings are with 6 GCN layers for each OSTR unit. The feature dimension d is set to be 512 in all sub-networks.

We compare the performance of HOSTR against recent state-of-the-art (SOTA) methods on all three datasets. Prior results are taken from [Le *et al.*, 2020a].

MSVD-QA and MSRVT-QA: Table 1 shows detailed comparisons on MSVD-QA and MSRVT-QA datasets. It is clear that our proposed method consistently outperforms all SOTA models. Specifically, we significantly improve performance on the MSVD-QA by 3.3 absolute points while the improvement on the MSRVT-QA is more modest. As videos in the MSRVT-QA are much longer (3 times longer than those in MSVD-QA) and contain more complicated interaction, it might require a larger number of input object sequences than what in our experiments (40 object sequences).

TGIF-QA: Table 2 presents the results on TGIF-QA dataset. As pointed out in [Le *et al.*, 2020a], short-term motion features are helpful for action task while long-term motion

¹<https://github.com/airsplay/py-bottom-up-attention>

features are crucial for event transition and count tasks. Hence, we provide two variants of our model: *HOSTR (R)* makes use of ResNet features as the context information for OSTR units; and *HOSTR (R+F)* makes use of the combination of ResNet features and motion features (extracted by ResNeXt, the same as in HCRN) as the context representation. *HOSTR (R+F)* shows exceptional performance on tasks related to motion. Note that we only use the context modulation at the video level to concentrate on the long-term motion. Even without the use of motion features, *HOSTR (R)* consistently shows more favorable performance than existing works.

The quantitative results prove the effectiveness of the object-oriented reasoning compared to the prior approaches of totally relying on frame-level features. Incorporating the motion features as context information also shows the flexibility of the design, suggesting that *HOSTR* can leverage a variety of input features and has the potentials to apply to other problems.

4.3 Ablation Studies

To provide more insight about our model, we examine the contributions of different design components to the model’s performance on the MSVD-QA dataset. We detail the results in Table 3. As shown, intra-object temporal attention seems to be more effective in summarizing the input object sequences to work with relational reasoning than BiLSTM. We hypothesize that it is due to the selective nature of the attention mechanism – it keeps only information relevant to the query.

As for the inter-object interaction, increasing the number of GCN layers up to 6 generally improves the performance. It gradually degrades when we stack more layers due to the gradient vanishing. The ablation study also points out the significance of the contextual representation as the performance steeply drops from 39.4 to 37.8 without them.

Last but not least, we conduct two experiments to demonstrate the significance of video hierarchical modeling. “*1-level hierarchy*” refers to when we replace all clip-level OSTRs with the global average pooling operation to summarize each object sequence into a vector while keeping the OSTR unit at the video level. “*1.5-level hierarchy*”, on the other hand, refers to when we use an average pooling operation at the video level while keeping the clip-level the same as in our *HOSTR*. Empirically, it shows that going deeper in hierarchy consistently improves performance on this dataset. The hierarchy may have greater effects in handling longer videos such as those in the MSRVTT-QA and TGIF-QA datasets.


4.4 Qualitative Analysis

To provide more analysis on the behavior of *HOSTR* in practice, we visualize the spatio-temporal graph formed during *HOSTR* operation on a sample in MSVD-QA dataset. In Fig. 4, the spatio-temporal graph of the two most important clips (judged by the temporal attention scores) are visualized in order of their appearance in the video’s timeline. Blue boxes indicate the six objects with highest edge weights (row summation of the adjacency matrix A calculated in Eq.6). The red lines indicates the most prominent edges of the graph with intensity scaled to the edge strength.

In this example, *HOSTR* attended mostly on the objects related to the concepts relevant to answer the question (the

Model	Test Acc. (%)
Default config. (*)	39.4
Temporal attention (TA)	
Attention at both levels	39.4
BiLSTM at clip, TA at video level	39.4
BiLSTM at both levels	38.8
Inter-object Interaction	
SR with 1 GCN layer	38.3
SR with 4 GCN layers	38.7
SR with 8 GCN layers	39.0
Contextual representation	
w/o contextual representation	37.8
Hierarchy	
1-level hierarchy	38.0
1.5-level hierarchy	38.7

Table 3: Ablation results on MSVD-QA dataset. Default config. (*): 6 GCN layers, Attention at clip & BiLSTM at video level



Q: Who played on the grass with a dog and a ball? A: Girl

Figure 4: A visualization of the spatio-temporal graph formed in *HOSTR*. The six most attended objects in each clip are drawn in blue bounding boxes. Red links indicates the importances of the edges.

girl, ball and dog). Furthermore, the relationships between the girl and her surrounding objects are the most important among the edges, and this intuitively agrees with how human might visually examine the scene given the question.

5 Conclusion

We presented a new object-oriented approach to Video QA where objects living in the video are treated as the primitive constructs. This brings us closer to symbolic reasoning, which is arguably more human-like. To realize this high-level idea, we introduced a general-purpose neural unit dubbed Object-oriented Spatio-Temporal Reasoning (OSTR). The unit reasons about its contextualized input – which is a set of object sequences – as instructed by the linguistic query. It first selectively transforms each sequence to an object node, then dynamically induces links between the nodes to build a graph. The graph enables iterative relational reasoning through collective refinement of object representation, gearing toward reaching an answer to the given query. The units are then stacked in a hierarchy that reflects the temporal structure of a typical video, allowing higher-order reasoning across space and time. Our architecture establishes new state-of-the-arts on major Video QA datasets designed for complex compositional questions, relational, temporal reasoning. Our analysis shows that object-oriented reasoning is a reliable, interpretable and effective approach to Video QA.

References

- [Baradel *et al.*, 2018] Fabien Baradel, Natalia Neverova, Christian Wolf, Julien Mille, and Greg Mori. Object level visual reasoning in videos. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 105–121, 2018.
- [Chao *et al.*, 2018] Yu-Wei Chao, Sudheendra Vijayanarasimhan, Bryan Seybold, David A Ross, Jia Deng, and Rahul Sukthankar. Rethinking the faster r-cnn architecture for temporal action localization. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 1130–1139, 2018.
- [Desta *et al.*, 2018] Mikyas T Desta, Larry Chen, and Tomasz Kornuta. Object-based reasoning in VQA. In *2018 IEEE Winter Conference on Applications of Computer Vision (WACV)*, pages 1814–1823. IEEE, 2018.
- [Gao *et al.*, 2018] Jiyang Gao, Runzhou Ge, Kan Chen, and Ram Nevatia. Motion-appearance co-memory networks for video question answering. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 6576–6585, 2018.
- [Huang *et al.*, 2020] Deng Huang, Peihao Chen, Runhao Zeng, Qing Du, Mingkui Tan, and Chuang Gan. Location-aware graph convolutional networks for video question answering. In *AAAI*, pages 11021–11028, 2020.
- [Jain *et al.*, 2016] Ashesh Jain, Amir R Zamir, Silvio Savarese, and Ashutosh Saxena. Structural-rnn: Deep learning on spatio-temporal graphs. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 5308–5317, 2016.
- [Jang *et al.*, 2017] Yunseok Jang, Yale Song, Youngjae Yu, Youngjin Kim, and Gunhee Kim. 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, 2017.
- [Kalogeiton *et al.*, 2017] Vicky Kalogeiton, Philippe Weinzaepfel, Vittorio Ferrari, and Cordelia Schmid. Action tubelet detector for spatio-temporal action localization. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 4405–4413, 2017.
- [Kim *et al.*, 2017] Kyung-Min Kim, Min-Oh Heo, Seong-Ho Choi, and Byoung-Tak Zhang. DeepStory: video story QA by deep embedded memory networks. In *IJCAI*, pages 2016–2022. AAAI Press, 2017.
- [Lake *et al.*, 2017] Brenden M Lake, Tomer D Ullman, Joshua B Tenenbaum, and Samuel J Gershman. Building machines that learn and think like people. *Behavioral and Brain Sciences*, 40, 2017.
- [Le *et al.*, 2020a] Thao Minh Le, Vuong Le, Svetha Venkatesh, and Truyen Tran. Hierarchical conditional relation networks for video question answering. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 9972–9981, 2020.
- [Le *et al.*, 2020b] Thao Minh Le, Vuong Le, Svetha Venkatesh, and Truyen Tran. Neural reasoning, fast and slow, for video question answering. In *IJCNN*, 2020.
- [Liang *et al.*, 2018] Junwei Liang, Lu Jiang, Liangliang Cao, Li-Jia Li, and Alexander G Hauptmann. Focal visual-text attention for visual question answering. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 6135–6143, 2018.
- [Pan *et al.*, 2020] Boxiao Pan, Haoye Cai, De-An Huang, Kuan-Hui Lee, Adrien Gaidon, Ehsan Adeli, and Juan Carlos Niebles. Spatio-temporal graph for video captioning with knowledge distillation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 10870–10879, 2020.
- [Spelke and Kinzler, 2007] Elizabeth S Spelke and Katherine D Kinzler. Core knowledge. *Developmental science*, 10(1):89–96, 2007.
- [Tran *et al.*, 2018] Du Tran, Heng Wang, Lorenzo Torresani, Jamie Ray, Yann LeCun, and Manohar Paluri. A closer look at spatiotemporal convolutions for action recognition. In *Proceedings of the IEEE conference on Computer Vision and Pattern Recognition*, pages 6450–6459, 2018.
- [Wang and Gupta, 2018] Xiaolong Wang and Abhinav Gupta. Videos as space-time region graphs. In *Proceedings of the European conference on computer vision (ECCV)*, pages 399–417, 2018.
- [Wang *et al.*, 2019] Anran Wang, Anh Tuan Luu, Chuan-Sheng Foo, Hongyuan Zhu, Yi Tay, and Vijay Chandrasekhar. Holistic multi-modal memory network for movie question answering. *IEEE Transactions on Image Processing*, 29:489–499, 2019.
- [Wojke *et al.*, 2017] Nicolai Wojke, Alex Bewley, and Dietrich Paulus. Simple online and realtime tracking with a deep association metric. In *ICIP*. IEEE, 2017.
- [Xie *et al.*, 2018] Saining Xie, Chen Sun, Jonathan Huang, Zhuowen Tu, and Kevin Murphy. Rethinking spatiotemporal feature learning: Speed-accuracy trade-offs in video classification. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 305–321, 2018.
- [Xu *et al.*, 2017] Dejing Xu, Zhou Zhao, Jun Xiao, Fei Wu, Hanwang Zhang, Xiangnan He, and Yueting Zhuang. Video question answering via gradually refined attention over appearance and motion. In *ACM Multimedia*, 2017.
- [Yang *et al.*, 2020] Zekun Yang, Noa Garcia, Chenhui Chu, Mayu Otani, Yuta Nakashima, and Haruo Takemura. Bert representations for video question answering. In *The IEEE Winter Conference on Applications of Computer Vision*, pages 1556–1565, 2020.
- [Ye *et al.*, 2017] Yunan Ye, Zhou Zhao, Yimeng Li, Long Chen, Jun Xiao, and Yueting Zhuang. Video question answering via attribute-augmented attention network learning. In *Proceedings of the 40th International ACM SIGIR Conference on Research and Development in Information Retrieval*, pages 829–832. ACM, 2017.
- [Yi *et al.*, 2020] Kexin Yi, Chuang Gan, Yunzhu Li, Pushmeet Kohli, Jiajun Wu, Antonio Torralba, and Joshua B. Tenenbaum. CLEVRER: collision events for video representation and reasoning. In *ICLR*, 2020.
- [Zeng *et al.*, 2019] Runhao Zeng, Wenbing Huang, Mingkui Tan, Yu Rong, Peilin Zhao, Junzhou Huang, and Chuang Gan. Graph convolutional networks for temporal action localization. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 7094–7103, 2019.
- [Zhao *et al.*, 2018] Zhou Zhao, Xinghua Jiang, Deng Cai, Jun Xiao, Xiaofei He, and Shiliang Pu. Multi-turn video question answering via multi-stream hierarchical attention context network. In *IJCAI*, pages 3690–3696, 2018.
- [Zhao *et al.*, 2019] Zhou Zhao, Zhu Zhang, Shuwen Xiao, Zhenxin Xiao, Xiaohui Yan, Jun Yu, Deng Cai, and Fei Wu. Long-form video question answering via dynamic hierarchical reinforced networks. *IEEE Transactions on Image Processing*, 28(12):5939–5952, 2019.