

This WACV 2023 paper is the Open Access version, provided by the Computer Vision Foundation. Except for this watermark, it is identical to the accepted version; the final published version of the proceedings is available on IEEE Xplore.

Dense but Efficient VideoQA for Intricate Compositional Reasoning

Jihyeon Lee^{*} Kakao Brain gina.ai@kakaobrain.com Wooyoung Kang^{*} Kakao Brain edwin.kang@kakaobrain.com Eun-Sol Kim Department of Computer Science, Hanyang University eunsolkim@hanyang.ac.kr

Abstract

It is well known that most of the conventional video question answering (VideoQA) datasets consist of easy questions requiring simple reasoning processes. However, long videos inevitably contain complex and compositional semantic structures along with the spatio-temporal axis, which requires a model to understand the compositional structures inherent in the videos. In this paper, we suggest a new compositional VideoQA method based on transformer architecture with a deformable attention mechanism to address the complex VideoQA tasks. The deformable attentions are introduced to sample a subset of informative visual features from the dense visual feature map to cover a temporally long range of frames efficiently. Furthermore, the dependency structure within the complex question sentences is also combined with the language embeddings to readily understand the relations among question words. Extensive experiments and ablation studies show that the suggested dense but efficient model outperforms other baselines.

1. Introduction

Along with the immense success of deep learning methods to understand the contents of images and text, various applications requiring complex reasoning have been proposed. Especially, visual question answering (VQA) [2] is one of the most important tasks, which asks a diverse set of questions about the visual contents and requires understanding the semantic structures inherent in the contents. By virtue of the emergence of transformer architectures and their pre-training schemes, the performance of the VQA has shown successful performance [34, 39]; however, it is not straightforward to apply the architectures to the video domain. Compared to the image and text, video data involves more complex semantic structures along with not only spatial but also temporal-axis. As described in Figure 1, long videos inevitably contain multiple events, and the events can have multiple and complex correlations. Therefore, it is important to temporally ground the multiple events and their semantic structure.

Q: Was a laptop or a television the thing she went in front of after laughing? A: Laptop

Figure 1. Example of an intricate VideoQA problem. The semantic elements of the video, which consists of characters, their actions, and the relationship between the characters, is continually changing along with temporal axis. Therefore, it is hard to answer the questions requiring understanding the complex semantic structure.

Most of the previous datasets proposed for VideoQA consist of relatively short clips containing an event or a single action class, with relatively easy questions [22, 36, 31, 46]. For this reason, the understanding of the short clips can be sufficiently addressed with image-based architectures by selecting a few representative frames from the clips. However, in the case of long videos having various events and complex relationships between the events, conventional architectures struggle to learn over large timescales. For these cases, it is essential to address the temporal grounding of the various events by considering enough frames within videos.

In this paper, we suggest a novel video/text understanding method for intricate VideoQA tasks, which consist of complicated questions requiring multiple reasoning steps. The two main ideas of the suggested methods are 1) efficiently sampling as many as possible informative visual features from the videos to learn the inherent temporal semantic structures and 2) considering the hierarchical dependency model to understand complex questions requiring multiple reasoning steps.

First of all, we suggest a deformable sampling module, which allows dense but efficient visual token sampling. Obviously, the conventional sparse sampling method [29], selecting a few frames followed by temporal pooling to get a single feature vector to be applied to downstream tasks, causes incomplete understanding for the long and intricate videos. As can be seen in Figure 1, at least 3 intervals from the video should be considered to get a correct answer to the question. Unfortunately, there is a fundamental trade-

^{*} indicates equal contribution

off between computational cost and the number of frames to be learned with the model. To settle the problem, we introduce a deformable attention module that effectively selects a subset of meaningful visual features along spatiotemporal axis. Specifically, the suggested method considers the semantics of the given query sentence.

Secondly, we introduce a dependency attention module to learn dependency-aware feature vectors of question tokens. As the input videos contain a more complex semantic structure, getting complicated questions requiring multiple reasoning steps are inevitable. Therefore, it is necessary to take into account the semantic structure within questions to learn desirable embeddings of the question tokens. We suggest leveraging the semantic structure from the dependency parse tree of the questions. By combining the deformable sampling module and the dependency attention module, our method is able to deal with the intricate compositional reasoning problems.

In experiments, we evaluate our model on Action Genome Question Answering (AGQA,[18]) dataset. The AGQA dataset is one of the most challenging benchmarks for VideoQA because it requires complex reasoning steps on long videos. Extensive experiments not only show impressive quantitative results on QA accuracy but also verify the effectiveness of each module by a comprehensive ablation study.

In summary, our contributions are as follows:

- We empirically reveal that covering a long time span is advantageous for complex problems, which needs spatio-temporal reasoning.
- We introduce a deformable sampling-based VideoQA model, DSR, which aims to solve compositional reasoning problem.
- Our experiments on VideoQA benchmarks show that the proposed method has the ability to perform complex spatio-temporal reasoning.

2. Related Work

Visual Question Answering VQA is the task of understanding how two inputs, text-based questions and visual features, relate to one another, proposed by Antol *et al.* [2]. For image-based question answering tasks, a significant amount of works propose attention-based model architectures to fuse question and image representations [1, 35, 25, 17]. Kim *et al.* [25] show remarkable performance by utilizing a bilinear attention network that finds bilinear interactions between two modalities. Moreover, inspired by the recent success of pre-trained language models [13, 11], universal pre-training frameworks for a vision-language representation learning achieve state-of-the-art performances not only on VQA but also on general visual-language tasks [39, 34].

However, question-answering in the video domain is under-explored compared to those in the image domain. Contrary to the growing interest in measuring video reasoning capabilities [22, 18, 30, 31, 45], existing VideoQA models mostly deal with short clip videos or simple questions [28, 8, 38, 29]. Since a video is a sequence of images containing the temporal dimension, understanding richer spatio-temporal features and temporal localization of natural language is essential. To fuse the temporal feature, Fan et al. and Seo et al. attempt to utilize separate motion and appearance feature modules and integrate them with additional fusion network [28, 8, 38]. Le et al. [28] propose a hierarchical conditional relation network to embed the video input at different granularities. However, separated modules have limitations in effectively interacting with linguistic questions, and the performances fell behind as the transformer-based model rises. Kim et al. [26] propose a contrastive learning based training scheme that shows competitive performance, but only specializes in multiplechoice tasks. The current state-of-the-art model in VideoQA is ClipBERT [29], which is based on a cross-modal transformer. ClipBERT enables end-to-end learning by employing sparse sampling while it is unsuitable for intricate tasks that require advanced spatio-temporal reasoning since random sparse sampling loses several semantic structures. We propose a dense but efficient VideoQA model based on a transformer, which can maintain whole semantic structures.

Efficient Transformers Transformer architecture [40] has shown remarkable performance on various downstream tasks. However, computational cost and memory consumption of Transformer increase quadratically depending on the length of input sequences. There has been a surge of research interests recently in exploiting efficient transformer architectures to mitigate the problem. For example, various algorithms that approximate the quadratic cost attention matrix based on low-rank matrix factorization have been proposed in the field of natural language processing [41, 10, 44]. Also, in the vision domain, the scope of self-attention is restricted to local neighborhoods or specific axis based on the locality assumption of objects [21, 9, 4, 3]. However, the above algorithms are defined for the single modality. In contrast, we aim to address cross-modal sparsification based on a question conditional visual token sampling algorithm for the VideoQA task where non-local and fine-grained features are required.

3. DSR: Deformable Sampling-based VideoQA model for Compositional Reasoning

In this section, we introduce a detailed explanation of our model. We consider the intricate VideoQA problems, which



Figure 2. The overall architecture of DSR. Two details are missed for simplification; deformable sampling is conditioned on question context embedding, and global visual features are used as additional input to the cross-modal transformer.

require compositional spatio-temporal reasoning. Our goal is to learn a generalizable visual-reasoning representation with deformable sampling and dependency modeling.

3.1. Transformer-based dense sampling model

We propose Deformable Sampling-based VideoQA model for compositional Reasoning (DSR), dense but efficient one that utilizes deformable sampling for video features and dependency modeling for text questions. Figure 2 gives an overall architecture of DSR, which is based on a cross-modal transformer. Each visual feature and question token are independently encoded with a vision backbone model and a language encoder, respectively. Inputs of the cross-modal transformer are conditionally sampled video features and dependency guided question tokens. We denote visual and language inputs of transformer as V = $[v_1, v_2, ..., v_{L_v}] \in \mathbb{R}^{d \times L_v} \text{ and } L = [l_1, l_2, ..., l_{L_q}] \in \mathbb{R}^{d \times L_q},$ respectively, where L_v is the number of visual tokens sampled from conditional sampling module, L_q represents the number of question tokens, and d indicates representation dimension. These embeddings of two different modalities are concatenated as input to a 12-layer transformer for cross-modal fusion, with special tokens [CLS] and [SEP].

We first uniformly sample the frames from a video, which is sufficiently dense to cover the full length of a video. However, as the length of the video increases, usage of whole dense frames becomes impossible since it can not fit into a single transformer due to memory limitations. Thus, motivated by Zhu *et al.* [48], we introduce a deformable sampling module to only sample necessary visual features from full dense frames, conditionally to question embeddings. Consequently, relatively few visual features compared to initial dense features are sampled from the module. A detailed explanation of conditional sampling is stated in Section 3.2. Language inputs (*i.e.*, question tokens) also go through the pre-stage modeling step to enable compositional reasoning. The dependency attention module forces specific attention head of transformer to understand dependency parsing structure, representing the relationship between words in a question sequence. It will be explained in Section 3.3.

The output vector of [CLS] token, \mathbf{h}_{cls} , is an aggregated representation of the entire input sequence of the crossmodal transformer, used to predict the answer. We consider all the QAs as open-ended word tasks, which choose one correct word as the answer from a pre-defined answer set of size C. We calculate a classification score by applying a linear classifier and softmax function on the final output and train the model by minimizing cross-entropy loss,

$$L_{open} = -\sum_{c=1}^{C} \mathbb{1}\{y = c\} \log(p_c),$$
 (1)

in which $\mathbf{p} = softmax(FFN(\mathbf{h_{cls}})) \in \mathbb{R}^C$ and y is the ground truth answer label. During inference, conditionally sampled visual features and dependency modeled linguistic features are utilized to predict answers with proper reasoning, in the same manner with the training phase. In summary, our model achieves state-of-the-art performance on intricate VideoQA tasks by allowing end-to-end learning while covering temporally long and spatially fine-grained visual features, which are both important for advanced modeling. Different from the model that only observes a single or a few video clips, DSR can tackle the data that needs compositional reasoning.

3.2. Conditional Visual Feature Sampling

In this section, we describe how to effectively sample a subset of visual tokens from the long and dense feature map. Since video data has an additional temporal axis compared to image data, the feature map size of a video clip is much bigger than the feature map from an image. Thus, most of the VideoQA algorithms pool the feature map spatially [32] or temporally [29] and concatenate the sequence of visual features to question word vectors. Then, the concatenated feature is used as an input for a transformer-based OA model. However, the pooling-based approach would be sub-optimal for compositional VideoQA tasks that require long and fined-grained visual cues. Here, we assume that most visual features in spatio-temporal feature maps are redundant and uninformative for answering given questions. In the next section, we describe how to sample a few informative visual features from the dense feature map.



Figure 3. Illustration of the proposed deformable sampling module. The figure represents a single head of a single CDA layer. For simplicity, we only visualize a deformable attention procedure of one reference point, which is solely colored as blue.

Conditional Deformable Attention Let $X \in \mathbb{R}^{d \times t \times h \times w}$ be a dense visual feature map extracted by a visual encoder such as ResNet [19]. The d, t, h, and w indicate dimension, temporal length, height, and width of the feature map, respectively. Based on the 2-d deformable attention module [48], we define our 3-d Conditional Deformable Attention (CDA) to sample question-conditional visual features from the spatio-temporal feature map X and a given question L as follows:

$$CDA(\mathbf{z}_q, p_q, X, L) = \sum_{m=1}^{M} W_m [\sum_{k=1}^{K} A_{mqk} \cdot W'_m X(p_q + \Delta p_{mqk})],$$

where, $A_{mqk} = W_m^A \bar{\mathbf{z}}_q, \quad \Delta p_{mqk} = W_m^{\Delta p} \bar{\mathbf{z}}_q, \quad \bar{\mathbf{z}}_q = \mathbf{z}_q \oplus pool(L),$
(2)

in which q is an element index for input query vector \mathbf{z}_q of a transformer layer and 3-d reference point p_q . In the first transformer layer, the input query $\mathbf{z}_q \in \mathbb{R}^{L_v \times d}$ is the learnable query, where L_v is the number of queries which is the same as the number of sampled visual features. Also, before feeding \mathbf{z}_q to the first transformer layer, we make a pooled question context $\hat{L} = pool(L) \in \mathbb{R}^{1 \times d}$ and add the question context to each learnable query by the broadcast vector addition, \oplus , to make CDA sample visual features conditioned on the given question context. For the rest of the layers, the $\bar{\mathbf{z}}_q$ is the output vector of the previous transformer layer. M and K denote the total number of attention heads and sampled key vectors, respectively. W, W', W_m^A , and $W_m^{\Delta p}$ are learnable linear projection layers. A_{mqk} denotes the attention weight of the k^{th} sampling point in the m^{th} attention head for a given query $\bar{\mathbf{z}}_q$, where $\sum_{k=1}^{K} A_{mqk} = 1$. $\Delta p_{mqk} \in \mathbb{R}^3$ is 3-d sampling offset. Since $p_q + \Delta p_{mqk}$ is a real-valued vector, we apply trilinear interpolation to compute $X(p_q + \Delta p_{mqk})$. With CDA, we can get L_v sampled visual tokens $V \in \mathbb{R}^{d \times L_v}$ where L_v is much smaller than the $t \times h \times w$, e.g., 25 vs. $30 \times 7 \times 7$. The overview of CDA is illustrated in Figure 3.

Regularization for Sampling Diversity For the question and answering task, sampled visual tokens from CDA are concatenated with question words, and a transformer-based model takes the concatenated features to predict an answer. Thus, it is important that the sampled visual tokens should be as diverse as possible to provide sufficient information for a given question. In Deformable DETR [48], offset predictions can be diverse without the collapse since each object query is trained to match a target object based on the Hungarian loss. However, in QA task, proper regularization is crucial to prevent the collapse because the model gets gradient feedback only from answering loss. To reinforce the diversity of sampled visual tokens, we explore three types of additional regularization terms. Here, we consider batched features where the V and X have shapes of $(N \times d \times L_v)$ and $(N \times d \times thw)$, respectively. The first regularization term is Soft Orthogonality (SO) [43], which is defined as follows:

(SO)
$$\lambda \sum_{i=1}^{N} \|V_i^{\top} V_i - I\|_F^2,$$
 (3)

where i indicates the index in a mini-batch. The SO aims the Gram matrix of the sampled tokens to be close to the identity matrix, I. Thus, each sampled visual token could be distinctive and independent.

The second regularization term is Maximal Coding Rate (MCR) [47], which formulated as follows:

(MCR)
$$-\lambda \sum_{i=1}^{N} \frac{1}{2} \log \det \left(I + \frac{d}{L_v \epsilon^2} V_i^{\top} V_i \right), \quad (4)$$

Maximizing the MCR results in the largest possible volume spanned by the vectors in the Gram matrix of the sampled tokens. Thus, the sampled tokens should be as independent as possible.

The last regularization term that we explore is the contrastive loss [6]. Here, we set the anchors, positive, and negative examples as sampled visual features V_i , the feature map X_i , and feature maps of others $X_{j\neq i}$ in the batch, respectively.

(Contrastive)
$$-\lambda \sum_{i=1}^{N} \log \frac{\exp(\sin(\hat{V}_i, \hat{X}_i)/\tau)}{\sum_{j=1}^{N} \mathbb{1}_{[j\neq i]} \exp(\sin(\hat{V}_i, \hat{X}_j)/\tau)},$$
 (5)



Figure 4. Example of dependency structure.



Figure 5. Adjacency matrix generated from dependency relations.

where \hat{V} and \hat{X} are the global averaged pooled of V and X, respectively. Also, we use the cosine similarity as the similarity function $\sin(\cdot, \cdot)$ and τ is set to 0.1 by default.

Global Context Features The sampled visual features from our CDA represent fine-grained local information that is required to answer the given question. However, the interaction between local and global information is also crucial to solving the spatio-temporally complex QA task more accurately. Thus, we introduce additional global information $X_g \in \mathbb{R}^{N \times d \times t}$ that is extracted by applying the spatial pooling to feature maps $X \in \mathbb{R}^{N \times d \times t \times h \times w}$. As a result, we get the global-local visual feature $X_{gl} \in \mathbb{R}^{N \times d \times (L_v+t)}$ by concatenating the global and sampled local visual features.

3.3. Dependency Attention Module

In this section, we explain the details of the dependency attention module that extracts dependency-aware vector from question tokens. Motivated from Deguchi *et al.* [12], we introduce a self-attention module that incorporates dependency relations. Previous studies show that the performance of neural machine translation has been improved by incorporating sentence structures [5, 15, 42]. While most visual-language learning tasks only rely on a pre-trained language model to encode question embeddings, we believe that comprehension of sentence structure is crucial for non-conventional questions, and the dependency-based attention module would also work for the VideoQA task.

Language features are first learned through the dependency attention module before feeding into the cross-modal transformer. The module is consist of an *L*-layer transformer, where one attention head of the *i*-th multi-head selfattention layer is trained with constraints based on dependency parsing value. Let $O^{i-1} \in \mathbb{R}^{L_q \times d}$ is the output of previous layer. The dependency attention module first maps ${\cal O}^{i-1}$ to $d_{head}\mbox{-dimensional subspaces of multi-head attention as}$

$$Q_{dep} = O^{i-1} W^{Q_{dep}}, (6)$$

$$K_{dep} = O^{i-1} W^{\kappa_{dep}},\tag{7}$$

$$V_{dep} = O^{i-1} W^{V_{dep}},\tag{8}$$

where $W^{Q_{dep}}$, $W^{K_{dep}}$, and $W^{V_{dep}}$ are $d \times d_{head}$ dimensional parameter matrices. Dependency attention weight matrix, A_{dep} , is calculated by the bi-affine operation [14] as follows,

$$A_{dep} = softmax(Q_{dep}UK_{dep}^{T}), \tag{9}$$

where $U \in \mathbb{R}^{d_{head} \times d_{head}}$. Each value of A_{dep} represents the dependency relationship between two words, and the probability of token q_1 being the governor of token q_2 is modeled as $A_{dep}[q_2, q_1]$. Then, likewise the original selfattention module, attention output is obtained by multiplying A_{dep} and V_{dep} . Finally, attention outputs of all the heads (*i.e.*, one dependency outputs and $n_{head} - 1$ conventional outputs) are concatenated and the rest are calculated like the conventional multi-head attention.

Although A_{dep} can be learned by additional dependency loss function as in Deguchi *et al.* [12], we explicitly force the correct dependency value not only in the training phase but also in inference. The gold parse provides an upper bound for using dependency relations and enables accurate structural modeling. Figure 4 shows the example of dependency relationships and Figure 5 represents the relations in an adjacency matrix that utilized as a gold value of A_{dep} . The gold value forces each token to only give attention to its governor.

To apply dependency relations to the transformer module, we reorganize the adjacency matrix for subword sequences. When a word is separated into multiple subwords by BPE [37], the governor (*i.e.*, the head) of the rightmost subword is set to the governor of the original word and the governor of each subword other than the rightmost one is set to the right adjacent subword.

4. Experiment

In this section, we evaluate our proposed model on compositional spatio-temporal reasoning datasets. We first introduce the details of benchmark datasets in Section 4.1. Section 4.2 describes experimental setup including implementation details. We also provide extensive quantitative experiments and ablation studies in Section 4.3, to show how each of the proposed modules works. Lastly, we qualitatively confirm that our model samples reasonable visual frames conditioned on given questions, in supplementary materials.

4.1. Dataset

Action Genome Question Answering We validate DSR using AGQA dataset proposed by Grunde-McLaughlin *et al.* [18], the most challenging benchmark for VideoQA. While most of the existing benchmarks only utilize short video clips, use simple and biased questions, and focus on questions that require commonsense or external knowledge, AGQA consists of long video clips with an average length of 30 seconds. Each question generated by a hand-crafted program necessarily requires spatio-temporal reasoning steps.

We adopt a balanced, novel composition, and more composition version of AGQA [18]. A balanced dataset, 3.9M of QA pairs associated with 9.6K videos, minimizes the bias by balancing the answer distributions and types of question structures. Novel composition is constructed to test whether models can disentangle distinct concepts and combine them well. For example, compositions like "before standing up" are removed from the training set, while each word "before" and "standing" appear. It tests how well the model performs on questions with those novel compositions in inference. More composition tests whether models generalize to more compositional steps. The training set only contains simpler questions with $\leq S$ compositional steps, while the test set contains only questions with > S reasoning steps. The model generalized to novel compositions and more compositional steps can be regarded as a successful VideoQA model that understands compositional semantic structures.

Open answer questions have many possible answers, while binary questions have answers that are yes/no or before/after. Except for the first table, all the tables adopt a 10% version of the balanced dataset for the training and inference phase. We provide details of the dataset we used in supplementary materials.

TVQA [30] is an intricate multiple-choice VideoQA dataset composed of 60-90 second long video clips. Although the video clips, questions, answers, subtitles, timestamps, and objects are given from the dataset, we only utilize video clips and QA pairs to verify intricate compositional reasoning ability. Most baselines use subtitles and propose the model that maximizes the subtitle knowledge since the performance gained from subtitles is much larger than that from video, however, we claim only to use videos, questions, and answers for *video question answering* tasks to show the video understanding ability.

Benchmarks with short videos MSRVTT-QA [46] is created based on videos in MSRVTT and questions are automatically generated from video descriptions. It consists of 10k videos and 243k QA pairs, with an average video length of 15 seconds. TGIF-QA [23] is web GIF VQA, containing 165K QA pairs on 72K GIF videos with an average length of 3 seconds long. MSRVTT and TGIF are not only short but also easy videos. These videos only require simple spatial reasoning while AGQA requires intricate spatio-temporal reasoning. According to the original paper, MSRVTT is a set of simple clips that each can be described with a single sentence, thus confined to a single domain. Lei *et al.* [29] support it by showing that adding more clips does not improve performance for both datasets; even somewhat has a negative effect. Our model does not stand out in simple tasks since it aims to solve intricate reasoning problems by modeling dense features, but even shows competitive results on the datasets.

4.2. Experimental Setup

Baselines We compare our approach against four recent VideoQA methods [33, 8, 28, 29]. PSAC [33] utilizes a co-attention block after unimodal self-attention blocks to simultaneously attend to both modalities. HME [8] models question, appearance, and motion features with different LSTM encoders. Additional visual and question memories help the multimodal fusion. HCRN [28] designs conditional relation networks and stacks them to accommodate diverse input modalities and conditioning features. ClipBERT [29] inputs a few short clips independently to a cross-modality transformer and aggregates prediction scores from multiple clips as the final score. For PSAC, HME, and HCRN, the performances reported in Grunde-McLaughlin *et al.* [18] are utilized.

Implementation Details 2D ResNet-50 [20] and word embedding layers of BERT-base model [13] are adopted as visual and language backbones. Specifically, 5 Conv blocks of ResNet-50 and an extra convolution layer are used for spatial down-sampling. We initialize visual/text encoder and cross-modal transformer with image-text pretrained weights proposed from ClipBERT [29], which leverages large-scale image-text datasets [7, 27].

We use a 4-layer transformer to construct our CDA. In each transformer layer, we set the number of attention heads and sampling points to 4 and 8, respectively. Also, we use a 2-layer transformer for the dependency attention encoder. The first attention head of the first layer corresponds to a dependency-guided self-attention module. 3D and 1D positional embeddings are applied for visual and language embeddings, respectively. We also add different type embeddings to both video and text inputs of the cross-modal transformer to indicate their source type. We report more details such as hyperparameters in supplementary materials.

4.3. Quantitative Results

Comparision with baselines on varied benchmarks We compare DSR with state-of-the-art models on the aforementioned datasets. As shown in Table 1, DSR consistently outperforms all the baselines on AGQA dataset. Compared to the best baseline, ClipBERT, our model achieves better

	Types	PSAC [33]	HME [8]	HCRN [28]	ClipBERT [29]	DSR(Ours)
Full Balanced	Binary Open All	54.19 27.20 40.40	59.77 36.23 47.74	58.11 37.18 47.42	63.83 48.54 53.03	65.92 (+2.09) 49.54 (+1.00) 54.36 (+1.33)
Novel Composition	Binary Open All	43.00 14.80 32.49	52.39 19.46 40.11	43.40 23.72 36.06	53.87 36.45 40.82	59.57 (+5.70) 38.73 (+2.28) 43.96 (+3.14)
More Composition	Binary Open All	35.39 28.00 31.13	48.09 33.47 39.70	42.46 34.81 38.00	42.93 45.93 45.32	47.79 (-0.30) 48.08 (+2.15) 48.02 (+2.70)

Table 1. Quantitative comparison with the baselines on AGQA datset. Full balanced, novel composition, and more composition represent different subset of AGQA as described in Section 4.1. The bold represents the best score.

Methods	MSRVTT-QA	Action	Transition	FrameQA	TVQA
Co-Memory [16]	32.0	68.2	74.3	51.5	-
PSAC [33]	-	70.4	76.9	55.7	-
HME [8]	33.0	73.9	77.8	53.8	-
HCRN [28]	35.6	75.0	81.4	55.9	-
QueST [24]	34.6	79.5	81.0	59.7	-
multi-stream [30]	-	-	-	-	43.8
ClipBERT [29]	37.4	82.4	87.3	58.8	44.4
DSR(Ours)	37.2	81.7	87.6	58.3	48.8

Table 2. Experimental results on benchmark datasets.

points of 1.33, 3.14, and 2.70 on full balanced, novel composition, and more composition datasets, respectively. We run three independent trials on Table 1,2,5 and confirmed the statistical significance of DSR via t-test. The noticeable thing is that DSR, the model concentrates on complex spatio-reasoning, especially records high scores on novel composition and more composition subsets. Since these two subsets are intentionally curated to test the generalizability and the reasoning ability of the model, the results give proof of the quality of DSR. The experimental results according to the number of compositional steps are in supplementary materials.

Table 2 shows the results on MSRVTT-QA, TGIF-QA, and TVQA dataset. We experiment over three tasks (*i.e., Action, Transition, FrameQA*) in TGIF-QA benchmark. Even though the MSRVTT-QA and TGIF-QA mostly only require an understanding of the spatial features rather than temporal reasoning of given questions, our method achieves a score comparable with ClipBERT. Moreover, DSR achieves the state-of-the-art result in the V+Q setting of TVQA, where subtitles and timestamps are not used for training.

QA Performance based on Sequence Length of Visual Features Here, we analyze the efficiency and effectiveness of DSR when addressing long sequence visual features. We first explore how the QA accuracy varies as we increase the number of frames so that the visual feature covers a

N_{Frames}	N_{V_Tokens}	Binary	Open	All
2	$2 \times 7 \times 7$	60.22	46.05	50.21
4	$4 \times 7 \times 7$	60.32	47.40	51.20
8	$8 \times 7 \times 7$	61.29	46.37	50.75
32 w/ DSR	32 + 25	64.47	48.58	53.24

Table 3. Accuracy based on various sequence length of visual features.

longer temporal range. In this experiment, we set fps as 1 by default. From Table 3, we observe that the QA accuracy increases as we show more frames to the model. However, since the computational cost of self-attention operations in the transformer-based QA module increases quadratically based on the input sequence length, there is a limitation to consider a more long-ranged sequence without any sparsification of visual features. On the contrary, DSR can sample a subset of informative visual features from the dense feature map, the number of visual features can be controllable as a hyperparameter. As a result, we achieve higher accuracy even with a much less number of visual tokens (57 vs. 392). More detailed analyses of memory efficiency according to sequence lengths of our DSR are discussed in supplementary materials.

Sparse Sampling vs. Dense Sampling In this experiment, we compare the effectiveness of randomly sampled sparse features and densely sampled features for the intricate compositional reasoning task. In the ClipBERT [29], they propose a sparse sampling-based training strategy due to the high computation cost and memory consumption. The Sparse Random in Table 4 follows the training convention of ClipBERT. They randomly sample multiple clips across the whole video, and each clip consists of 2 consecutive frames with fps 2. Then, a shared transformer-based QA model independently predicts answers based on the multiple clips. Finally, the answer logits from each clip are averaged as a final decision. In contrast to sparse sampling, dense sampling aims to see longer sequences temporally with just one clip to address the intricate spatio-temporal

Sampling Method	N_{Frames}	N_{Clips}	Acc.
	2	1	50.57
Sparse Random	2	2	52.17
	2	4	52.80
	2	16	52.93
Dense w/ DSR	32	1	53.24

Table 4. Comparison of sparse sampling and dense sampling strategies.

reasoning task, AGQA. We observe that dense sampling with DSR shows higher accuracy than sparse sampling. Since DSR can sample a few diverse informative visual features from the spatio-temporally dense feature map, the model can effectively associate question words and sampled visual features, which leads to the highest accuracy in Table 4.

Ablation study In this section, we conduct extensive ablation experiments about the hyperparameters of DSR. The first row in Table 5 is our best configuration among all controllable variables. Firstly, we observe that there is an improvement with the dependency attention module. The dependency encoder helps structurization of question sequences by forcing dependency relations. Then, we explore the best number of visual features to be sampled. When we increase the number of object queries from 5 to 25, there are consistent improvements in the QA accuracy. However, if we set the number of object queries to 50, the QA accuracy drops slightly. We analyze that noisy and redundant visual features could be sampled if we consider too many sampling points. Thus, we set the number of object queries to 25 by default for the experiments in prior sections.

The next ablation is about the effectiveness of the global context features. From the sixth and seventh rows in Table 5, we observe that the global context features are notably helpful to boost the QA accuracy. While the "only local" model shows lower accuracy than the "only global" model, we achieve the best performance with the combination of global features and local features. This indicates that a proper association of global and local features is crucial to address the complex spatio-temporal reasoning task.

Subsequently, we explore the 3 types of sampling regularization terms. We find that the Soft Orthogonality (Eq 3) regularization achieves the best performance. The MCR regularization shows a high variance in the norm of gradients, which causes an unstable training process. We conjecture that the high variance comes from the logdet operator. Also, the contrastive loss shows the lowest accuracy. This could be due to the small batch size caused by the 12-layer transformer and ResNet-50 taking video data as input.

Finally, we compare DSR to a random sampling strategy of visual features. For the random sampling strategy, we randomly sample 57 visual features from the dense feature map during the training phase. Then, we uniformly sample

Dep.	# of Obj. q	Glob.	Reg.	Sampl.	Acc.
✓	32 + 25	both	SO	Deform	53.24
X	32 + 25	both	SO	Deform	52.26
~	32 + 5	both	SO	Deform	47.55
\checkmark	32 + 10	both	SO	Deform	51.29
\checkmark	32 + 50	both	SO	Deform	52.04
~	32	only global	SO	Deform	50.64
\checkmark	25	only local	SO	Deform	45.72
1	32 + 25	both	-	Deform	50.97
\checkmark	32 + 25	both	Cont.	Deform	49.84
\checkmark	32 + 25	both	MCR	Deform	51.06
1	57	local	-	Rand	51.81

Table 5. Results on ablation experiments.

the visual features from the flattened dense feature map at the inference phase. As expected, DSR with the diversity regularization and the global-local fusion achieves higher accuracy by a large margin than the random sampling under the same number of visual tokens. This indicates that our two strategies, avoiding collapsed sampling and globallocal information interaction, are essential in this samplingbased VideoQA task.

In addition, we 1) analyze the diversity and suitability of sampled tokens by visualizing the output of the deformable sampler, 2) visualize the effectiveness of the dependency attention module, and 3) validate that the performance gain of DSR comes from novel modules we proposed, but not from the increased parameter, in supplementary materials.

5. Conclusion

This paper presents the state-of-the-art compositional reasoning model for video question answering tasks, DSR, which utilizes deformable sampling module and dependency attention module for efficient video-text representation learning. Based on our finding that the dense model performs better than the sparse model on the compositional reasoning dataset, which is a different point of view from previous work, we conditionally sample question-related visual features from a dense feature map. This process remarkably reduces the number of visual tokens needed for cross-modal transformer while rather improving the efficiency; maximum allowed batch size and performance increase. The dependency-based attention module eases the model to conduct multi-step reasoning by guiding a particular attention head with structuralized dependency relations. Extensive experiments verify our model especially stands out against others on intricate benchmarks. Comprehensive ablation studies demonstrate each factor fairly contributes to our model.

Acknowledgements This work was supported by Korea research grant from IITP (2022-0-00264/50%, 2022-0-00951/20%, 2022-0-00612/20%, 2020-0-01373/10%).

References

- [1] Peter Anderson, Xiaodong He, Chris Buehler, Damien Teney, Mark Johnson, Stephen Gould, and Lei Zhang. Bilinear Attention Networks. In *CVPR*, 2018.
- [2] Stanislaw Antol, Aishwarya Agrawal, Jiasen Lu, Margaret Mitchell, Dhruv Batra, C. Lawrence Zitnick, and Devi Parikh. Vqa: Visual question answering. In *ICCV*, 2015.
- [3] Anurag Arnab, Mostafa Dehghani, Georg Heigold, Chen Sun, Mario Lučić, and Cordelia Schmid. Vivit: A video vision transformer. In *ICCV*, 2021.
- [4] Gedas Bertasius, Heng Wang, and Lorenzo Torresani. Is space-time attention all you need for video understanding? In Proc. the International Conference on Machine Learning (ICML), July 2021.
- [5] Kehai Chen, Rui Wang, Masao Utiyama, Lemao Liu, Akihiro Tamura, Eiichiro Sumita, and Tiejun Zhao. Neural machine translation with source dependency representation. In *Proc. of the Conference on Empirical Methods in Natural Language Processing (EMNLP)*, 2017.
- [6] Ting Chen, Simon Kornblith, Mohammad Norouzi, and Geoffrey Hinton. A simple framework for contrastive learning of visual representations. In *Proc. the International Conference on Machine Learning (ICML)*, 2020.
- [7] Xinlei Chen, Hao Fang, Tsung-Yi Lin, Ramakrishna Vedantam, Saurabh Gupta, Piotr Dollár, and C. Lawrence Zitnick. Microsoft coco captions: Data collection and evaluation server. *Computing Research Repository (CoRR)*, 2015.
- [8] Fan Chenyou, Zhang Xiaofan, Zhang Shu, Wang Wensheng, Zhang Chi, and Huang Heng. Heterogeneous memory enhanced multimodal attention model for video question answering. In *CVPR*, 2019.
- [9] Rewon Child, Scott Gray, Alec Radford, and Ilya Sutskever. Generating long sequences with sparse transformers. arXiv preprint arXiv:1904.10509, 2019.
- [10] Krzysztof Marcin Choromanski, Valerii Likhosherstov, David Dohan, Xingyou Song, Andreea Gane, Tamas Sarlos, Peter Hawkins, Jared Quincy Davis, Afroz Mohiuddin, Lukasz Kaiser, David Benjamin Belanger, Lucy J Colwell, and Adrian Weller. Rethinking attention with performers. In *ICLR*, 2021.
- [11] Kevin Clark, Minh-Thang Luong, Quoc V. Le, and Christopher D. Manning. ELECTRA: Pre-training text encoders as discriminators rather than generators. In *ICLR*, 2020.
- [12] Hiroyuki Deguchi, Akihiro Tamura, and Takashi Ninomiya. Dependency-based self-attention for transformer NMT. In Proceedings of the International Conference on Recent Advances in Natural Language Processing (RANLP), 2019.
- [13] Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. BERT: Pre-training of deep bidirectional transformers for language understanding. In Proc. of the Annual Conference of the North American Chapter of the Association for Computational Linguistics (NAACL), 2019.
- [14] Timothy Dozat and Christopher D. Manning. Deep biaffine attention for neural dependency parsing. In *ICLR*, 2017.
- [15] Akiko Eriguchi, Yoshimasa Tsuruoka, and Kyunghyun Cho. Learning to parse and translate improves neural machine

translation. In Proc. the Annual Meeting of the Association for Computational Linguistics (ACL), 2017.

- [16] Jiyang Gao, Runzhou Ge, Kan Chen, and Ram Nevatia. Motion-appearance co-memory networks for video question answering. In *CVPR*, 2018.
- [17] Peng Gao, Zhengkai Jiang, Haoxuan You, Pan Lu, Steven CH Hoi, Xiaogang Wang, and Hongsheng Li. Dynamic fusion with intra-and inter-modality attention flow for visual question answering. In CVPR, 2019.
- [18] Madeleine Grunde-McLaughlin, Ranjay Krishna, and Maneesh Agrawala. Agqa: A benchmark for compositional spatio-temporal reasoning. In *CVPR*, 2021.
- [19] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In CVPR, 2016.
- [20] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In CVPR, 2016.
- [21] Jonathan Ho, Nal Kalchbrenner, Dirk Weissenborn, and Tim Salimans. Axial attention in multidimensional transformers. arXiv preprint arXiv:1912.12180, 2019.
- [22] Yunseok Jang, Yale Song, Chris Dongjoo Kim, Youngjae Yu, Youngjin Kim, and Gunhee Kim. Video Question Answering with Spatio-Temporal Reasoning. *IJCV*, 2019.
- [23] Yunseok Jang, Yale Song, Youngjae Yu, Youngjin Kim, and Gunhee Kim. Tgif-qa: Toward spatio-temporal reasoning in visual question answering. In CVPR, 2017.
- [24] Jianwen Jiang, Ziqiang Chen, Haojie Lin, Xibin Zhao, and Yue Gao. Divide and conquer: Question-guided spatiotemporal contextual attention for video question answering. In AAAI, 2020.
- [25] Jin-Hwa Kim, Jaehyun Jun, and Byoung-Tak Zhang. Bilinear Attention Networks. In *NeurIPS*, 2018.
- [26] Seonhoon Kim, Seohyeong Jeong, Eunbyul Kim, Inho Kang, and Nojun Kwak. Self-supervised pre-training and contrastive representation learning for multiple-choice video QA. *CoRR*, 2020.
- [27] Ranjay Krishna, Yuke Zhu, Oliver Groth, Justin Johnson, Kenji Hata, Joshua Kravitz, Stephanie Chen, Yannis Kalantidis, Li-Jia Li, David A. Shamma, Michael S. Bernstein, and Li Fei-Fei. Visual genome: Connecting language and vision using crowdsourced dense image annotations. *IJCV*, 2016.
- [28] Thao Minh Le, Vuong Le, Svetha Venkatesh, and Truyen Tran. Hierarchical conditional relation networks for video question answering. arXiv preprint arXiv:2002.10698, 2020.
- [29] Jie Lei, Linjie Li, Luowei Zhou, Zhe Gan, Tamara L. Berg, Mohit Bansal, and Jingjing Liu. Less is more: Clipbert for video-and-language learningvia sparse sampling. In CVPR, 2021.
- [30] Jie Lei, Licheng Yu, Mohit Bansal, and Tamara L Berg. Tvqa: Localized, compositional video question answering. In Proc. of the Conference on Empirical Methods in Natural Language Processing (EMNLP), 2018.
- [31] Jie Lei, Licheng Yu, Tamara L Berg, and Mohit Bansal. Tvqa+: Spatio-temporal grounding for video question answering. In *Tech Report, arXiv*, 2019.

- [32] Linjie Li, Yen-Chun Chen, Yu Cheng, Zhe Gan, Licheng Yu, and Jingjing Liu. Hero: Hierarchical encoder for video+ language omni-representation pre-training. In Proc. of the Conference on Empirical Methods in Natural Language Processing (EMNLP), 2020.
- [33] Xiangpeng Li, Jingkuan Song, Lianli Gao, Xianglong Liu, Wenbing Huang, Xiangnan He, and Chuang Gan. Beyond rnns: Positional self-attention with co-attention for video question answering. In AAAI, 2019.
- [34] Jiasen Lu, Dhruv Batra, Devi Parikh, and Stefan Lee. Vilbert: Pretraining task-agnostic visiolinguistic representations for vision-and-language tasks. In *NeurIPS*, 2019.
- [35] Zhou Luowei, Palangi Hamid, Zhang Lei, Hu Houdong, J. Corso Jason, and Gao Jianfeng. Unified vision-language pre-training for image captioning and vqa. arXiv preprint arXiv:1909.11059, 2019.
- [36] Jonghwan Mun, Paul Hongsuck Seo, Ilchae Jung, and Bohyung Han. Marioqa: Answering questions by watching gameplay videos. In *ICCV*, 2017.
- [37] Rico Sennrich, Barry Haddow, and Alexandra Birch. Neural machine translation of rare words with subword units. In *Proc. the Annual Meeting of the Association for Computational Linguistics (ACL)*, 2016.
- [38] Ahjeong Seo, Gi-Cheon Kang, Joonhan Park, and Byoung-Tak Zhang. Attend what you need: Motion-appearance synergistic networks for video question answering. In Proc. the Annual Meeting of the Association for Computational Linguistics (ACL), 2021.
- [39] Hao Tan and Mohit Bansal. Lxmert: Learning crossmodality encoder representations from transformers. In *Proc. of the Conference on Empirical Methods in Natural Language Processing (EMNLP)*, 2019.
- [40] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Ł ukasz Kaiser, and Illia Polosukhin. Attention is all you need. In I. Guyon, U. V. Luxburg, S. Bengio, H. Wallach, R. Fergus, S. Vishwanathan, and R. Garnett, editors, *NeurIPS*, 2017.
- [41] Sinong Wang, Belinda Z Li, Madian Khabsa, Han Fang, and Hao Ma. Linformer: Self-attention with linear complexity. *arXiv preprint arXiv:2006.04768*, 2020.
- [42] Shuangzhi Wu, Dongdong Zhang, Zhirui Zhang, Nan Yang, Mu Li, and Ming Zhou. Dependency-to-dependency neural machine translation. *IEEE/ACM Transactions on Audio*, *Speech, and Language Processing*, 2018.
- [43] Di Xie, Jiang Xiong, and Shiliang Pu. All you need is beyond a good init: Exploring better solution for training extremely deep convolutional neural networks with orthonormality and modulation. In *CVPR*, 2017.
- [44] Yunyang Xiong, Zhanpeng Zeng, Rudrasis Chakraborty, Mingxing Tan, Glenn Fung, Yin Li, and Vikas Singh. Nyströmformer: A nyström-based algorithm for approximating self-attention. 2021.
- [45] 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.

- [46] Jun Xu, Tao Mei, Ting Yao, and Yong Rui. Msr-vtt: A large video description dataset for bridging video and language. In *CVPR*, 2016.
- [47] Yaodong Yu, Kwan Ho Ryan Chan, Chong You, Chaobing Song, and Yi Ma. Learning diverse and discriminative representations via the principle of maximal coding rate reduction. *Advances in Neural Information Processing Systems*, 2020.
- [48] Xizhou Zhu, Weijie Su, Lewei Lu, Bin Li, Xiaogang Wang, and Jifeng Dai. Deformable detr: Deformable transformers for end-to-end object detection. arXiv preprint arXiv:2010.04159, 2020.