

Gaining Extra Supervision via Multi-task learning for Multi-Modal Video Question Answering

Junyeong Kim*

School of Electrical Engineering
Korea Advanced Institute of Science and Technology
Daejeon, South Korea
junyeong.kim@kaist.ac.kr

Minuk Ma*

School of Electrical Engineering
Korea Advanced Institute of Science and Technology
Daejeon, South Korea
akalsdnr@kaist.ac.kr

Kyungsu Kim

Samsung Electronics
Seoul, South Korea
ks0326.kim@samsung.com

Sungjin Kim

Samsung Electronics
Seoul, South Korea
sj9373.kim@samsung.com

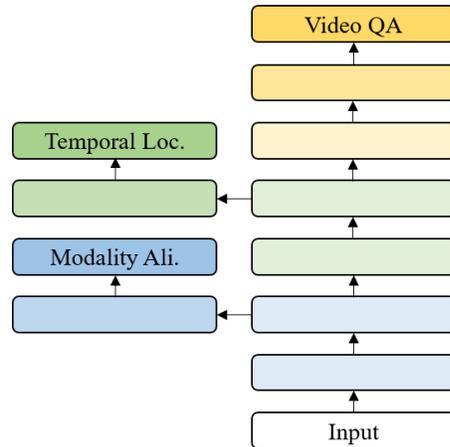
Chang D. Yoo

School of Electrical Engineering
Korea Advanced Institute of Science and Technology
Daejeon, South Korea
cd_yoo@kaist.ac.kr

Abstract—This paper proposes a method to gain extra supervision via multi-task learning for multi-modal video question answering. Multi-modal video question answering is an important task that aims at the joint understanding of vision and language. However, establishing large scale dataset for multi-modal video question answering is expensive and the existing benchmarks are relatively small to provide sufficient supervision. To overcome this challenge, this paper proposes a multi-task learning method which is composed of three main components: (1) multi-modal video question answering network that answers the question based on the both video and subtitle feature, (2) temporal retrieval network that predicts the time in the video clip where the question was generated from and (3) modality alignment network that solves metric learning problem to find correct association of video and subtitle modalities. By simultaneously solving related auxiliary tasks with hierarchically shared intermediate layers, the extra synergistic supervisions are provided. Motivated by curriculum learning, multi-task ratio scheduling is proposed to learn easier task earlier to set inductive bias at the beginning of the training. The experiments on publicly available dataset TVQA shows state-of-the-art results, and ablation studies are conducted to prove the statistical validity.

I. INTRODUCTION

Multi-modal video question answering is an important task that aims at the joint understanding of vision and language. Its application ranges from multimedia search engine to personal assistant. Deep neural network have been successfully applied to various QA tasks including textQA [1]–[4], imageQA [5]–[9], and videoQA [10]–[13] with significant performance improvement. Recently, the research on multi-modal videoQA [14]–[19] have also benefited from deep neural networks. One of the main challenges of multi-modal video question answering is that the size of existing benchmark datasets (e.g. MovieQA [14], PororoQA [20], and TVQA [15]) are relatively small to provide sufficient supervision to train QA models on the complex task. This paper proposes a method to gain extra supervision via multi-task learning for multi-modal video question answering. By solving related auxiliary tasks



multi-modal videoQA datasets have been introduced including MovieQA [14], PororoQA [20], and TVQA [15]. However, these benchmark datasets have a relatively small number of QA pairs considering the complexity of the task. MovieQA [14] consists of 6,462 QA pairs for the video+subtitles task, PororoQA consists of 8,913 QA pairs, and TVQA consists of 152.5K QA pairs. Comparing with imageQA benchmark datasets, VQA v1.0 [21] dataset consists of 614.2K QA pairs and VQA v2.0 dataset consists of 1.1M QA pairs.

Multi-Task Learning (MTL) is a learning paradigm in machine learning which jointly solves multiple tasks in a single model. MTL aims to leverage useful information contained in multiple related tasks to gain positive synergies across all the tasks. For example, tasks like temporal localization and visual-semantic alignment have been found useful for each other when trained jointly [22]. To solve multi-modal videoQA, this paper proceeds by analogy to human intelligence. Humans would first have to know the proper association between vision and language. On top of that, humans would attempt to localize the moment which is relevant to answering the question. Finally, humans could learn how to answer the questions. We formulated this as a multi-task learning problem and designed two auxiliary tasks.

This paper proposes a method to gain extra supervision via multi-task learning for multi-modal videoQA. Solving auxiliary tasks simultaneously with the QA task can provide synergistic learning signals. On top of the QA network based on Jie et al. [15], we introduce two auxiliary tasks that hierarchically share parameters with the QA network as depicted in Fig. 1. One auxiliary task is modality alignment which aims at correctly associating video and subtitle features. It shares parameters with the lower layers of the QA network. The other task is temporal localization which aims at finding the moment in the video clip that is most relevant to answering the current question. It shares parameters with the higher layers of the QA network. In order to control the timing and strength of the objective of each task, multi-task ratio scheduling is proposed. Motivated by curriculum learning [23], multi-task ratio scheduling attempts to learn easier task earlier to set inductive bias at the beginning of the training. The main contribution of this paper is summarized as follows. (1) Multi-task learning method for multi-modal videoQA is proposed which recorded state-of-the-art performance on TVQA benchmark. (2) Multi-task ratio scheduling is proposed to efficiently reflect the objectives of each task during training.

II. RELATED WORK

In this section, we introduce the related works of our paper in four categories; multi-task learning, modality alignment, temporal localization and multi-modal video question answering.

A. Multi-Task Learning

Multi-task learning aims at jointly solving multiple related tasks with a single model. By sharing parameters across related tasks, the model can generalize better on the original task.

Most of the multi-task learning methods share the hidden layers across all tasks and have task-specific output layers for each task. Starting from the work of Kaufmann et al. [24], there have been rich research on multi-task learning across the majority of applications of machine learning from computer vision (CV) [25] to natural language processing (NLP) [26]. Kim et al. [27] proposed Deep Partial Person Re-identification (DPPR) that jointly learns person classification and person re-identification for partial person re-identification. Object detection architectures such as Fast R-CNN [25] and Faster R-CNN [28] used multi-task loss for bounding box regression and object classification. Tiao et al. [29] tackled the task of Person Search [29] by jointly learning pedestrian detection and person re-identification. Recently, Li et al. [30] proposed the invertible Question Answering Network (iQAN) to leverage the complementary relations between questions and answers in the image by jointly learning the Visual Question Answering (VQA) and Visual Question Generation (VQG) tasks.

B. Modality Alignment

As an auxiliary task of our proposed method, we jointly learn the modality alignment and the temporal localization along with the multi-modal video question answering. Both tasks have been extensively studied in the field of deep learning. Karpathy et al. [31] proposed a method that captures the inter-modal correspondences between vision and language to generate natural language description of the given image. The latent alignment between the segments of the sentence and the region of the image is learned with a structured max-margin loss. Castrejón et al. [32] proposed a method that learn cross-modal scene representations that transfer across modalities. By regularizing cross-modal CNNs to have shared representation, the resulting representation is agnostic of the modality. Yu et al. [33] proposed Joint Sequence Fusion (JSFusion) model that can measure semantic similarity between any pairs of multimodal sequence data. Hierarchical attention mechanism is utilized to learn matching representation patterns among modalities.

C. Temporal Localization

Temporal localization aims at localizing temporal parts from the given video. Hendricks et al. [34] proposed the Moment Context Network (MCN) for temporal localization with natural language query. The MCN effectively localizes temporal parts related to natural language query by integrating local and global video feature over time. Gao et al. [22] proposed a multi-task learning approach for temporal localization with natural language query. Location regression and visual-semantic alignment are jointly learned for temporal localization. Temporal Unit Regression Network (TURN) [35] jointly predicts action proposals and refines the temporal boundaries by temporal coordinate regression. Long untrimmed video is decomposed into video clips, which are reused as basic building blocks of temporal proposals for fast computation.

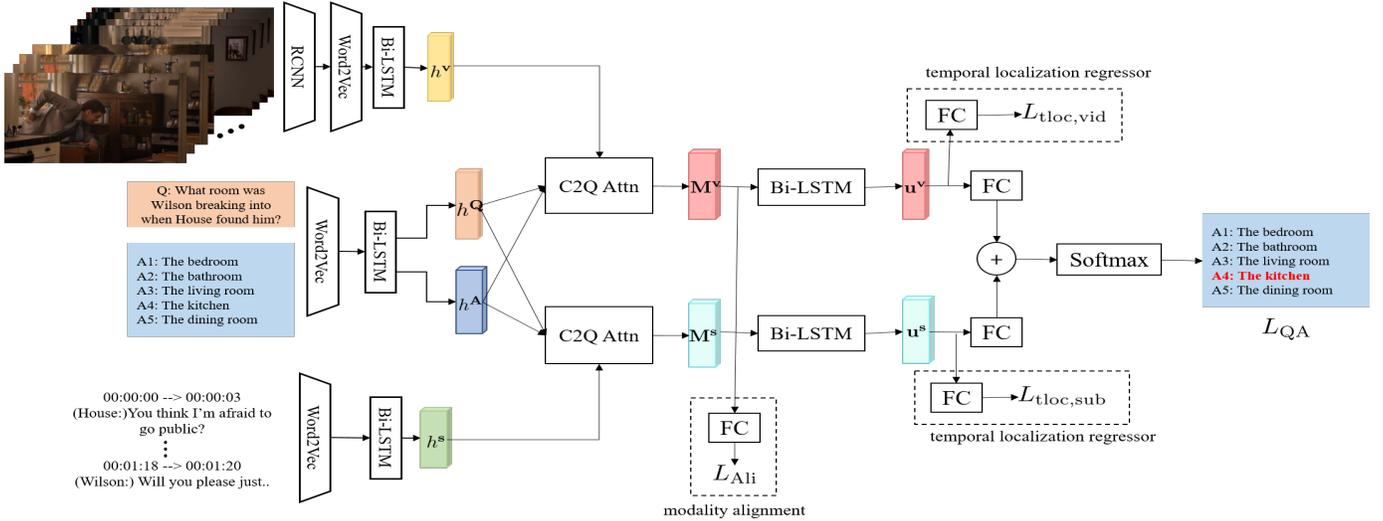


Fig. 2. The overall architecture of proposed method for multi-modal video question answering. The proposed network is composed of three networks which share intermediate layers; QA network, modality alignment network, and temporal localization network.

D. Multi-Modal Video Question Answering

Recently, the research on multi-modal videoQA [14]–[19] leverages additional text modality such as subtitle along with video modality for the joint understanding of vision and language. There are various benchmark datasets for multi-modal videoQA including MovieQA [14], PororoQA [20] and TVQA [15]. Multi-modal videoQA is a challenging task for its relatively small size of benchmark datasets. The majority of the methods on multi-modal videoQA are motivated by memory-augmented architecture [2]. Tapaswi et al. [14] utilized memory network (MemN2N) [2] to store video clips into memory and retrieve required information for answering question. Read-Write Memory Network (RWMN) [16] replaces the fully-connected layers in memory network [2] into the convolutional layers to capture local information in each memory slot. After the video and subtitle features are fused using bilinear operation, convolutional write/read networks store/retrieve information, respectively. Focal Visual-Text Attention (FVTA) [17] applied hierarchical attention mechanism on three-dimensional tensor of question, video and text to dynamically determine which modality and what time to attend for question answering. Multimodal Dual Attention Memory (MDAM) [19] applied multi-head attention mechanism [36] to learn the latent representation of multi-modal inputs.

III. THE PROPOSED METHOD

In this section, we describe the proposed method and training procedure in details. Fig. 2 gives the overview of overall architecture which fully utilizes multi-modal inputs (video and subtitle) and QA pairs to answer the question. The proposed method is composed of three networks which hierarchically shares intermediate layers; question-answering (QA) network, modality alignment network and temporal localization network. Note that we utilized QA network proposed by Jie et al. [15].

A. Problem Formulation

The formal definition of Multi-modal Video Question Answering (QA) is as follows. The inputs to the model are (1) a video clip \mathbf{V} , (2) a subtitle corresponding to the video clip \mathbf{S} , (3) a question \mathbf{q} and (4) five candidate answers $\mathbf{a} = \{a_i\}_{i=0}^4$ where only one is correct. The task is to predict correct answer for the question and the objective of training is to learn model parameters to maximize the following log-likelihood:

$$\hat{\theta} = \operatorname{argmax}_{\theta} \sum_{\mathcal{D}} \log P(\mathbf{y} | \mathbf{v}, \mathbf{s}, \mathbf{q}, \mathbf{a}; \theta), \quad (1)$$

where \mathcal{D} denotes the dataset, θ represents the model parameters and \mathbf{y} denotes the correct answer out of five candidates.

B. Feature Extraction

Before describing the proposed method, we first introduce the feature extraction method. For a fair comparison, we followed the same video and text feature extraction procedure used in previous work [15] and fixed them during training.

1) *Video Features*: We extracted two types of video features; ImageNet feature and visual concept feature. First, the frames are extracted from each video clip \mathbf{v} with the rate of 3fps.

ImageNet Feature: For each frame, the feature vector of size 2048D is extracted from “Average Pooling” layer of ResNet-101 [37] trained on ImageNet Benchmark [38]. The frame feature corresponding to the same video clip are first L2-normalized and then stacked, forming ImageNet feature $\mathbf{V}^{img} \in \mathbb{R}^{n_{img} \times 2048}$, where n_{img} represents the number of frames extracted from the video clip.

Visual Concept Feature: Inspired by recent works [15], [39] that use detected object labels as visual inputs instead of using CNN features directly, we also extracted detected labels which are referred to as visual concept feature [15]. Faster R-CNN [28] trained on Visual Genome Benchmark [40] is

utilized to detect objects in each frame. After collecting every detected concept for each video clip over all of the frames and eliminating the overlapping concepts, we utilized GloVe [41] to embed each concept into feature representation. The resulting visual concept feature is represented as $\mathbf{V}^{cpt} \in \mathbb{R}^{n_{cpt} \times 300}$, where n_{cpt} denotes the number of concepts in video clip.

2) *Text Features*: We used GloVe [41] to embed words into feature representations. Every sentences in each subtitle \mathbf{S} are flattened and tokenized into a sequence of words and GloVe [41] embeds the sequence of words into subtitle feature denoted as $\mathbf{S}^{rep} \in \mathbb{R}^{n_s \times 300}$, where n_s represents the number of words in the subtitle. Question feature $\mathbf{q}^{rep} \in \mathbb{R}^{n_q \times 300}$ and candidate answer feature $a_i^{rep} \in \mathbb{R}^{n_{a_i} \times 300}$ are also embedded similarly, where n_q and n_{a_i} are the number of words in question \mathbf{q} and candidate answer a_i , respectively.

C. Question-Answering Network

Now, we describe the QA network proposed by Jie et al. [15]. First, bi-directional LSTM (bi-LSTM) is used to encode the both video and text features into embedding space. The bi-LSTM consists of two LSTM layers; forward LSTM \vec{f} and backward LSTM \overleftarrow{f} . For the input sequence of $\{x_i\}_{t=1}^T$, the forward LSTM \vec{f} encodes the input sequence in forward order (from x_1 to x_T) into hidden states (h_1, \dots, h_T) . The backward LSTM \overleftarrow{f} encodes the input sequence in backward order (from x_T to x_1) and generate hidden states $(\overleftarrow{h}_1, \dots, \overleftarrow{h}_T)$. The hidden states from both directions at each time step are stacked to obtain resulting hidden representation, i.e. $h_t = [\vec{h}_t; \overleftarrow{h}_t]$ where $[\cdot]$ represents concatenation. Subtitle \mathbf{S}^{rep} , question \mathbf{q}^{rep} , candidate answer a_i^{rep} and visual concept \mathbf{V}^{cpt} features are now encoded by bi-LSTM and denoted as $h^S \in \mathbb{R}^{n_s \times 2d}$, $h^q \in \mathbb{R}^{n_q \times 2d}$, $h^{a_i} \in \mathbb{R}^{n_{a_i} \times 2d}$ and $h^{cpt} \in \mathbb{R}^{n_{cpt} \times 2d}$, respectively. Here, d is the size of hidden state which is set to 150 in this experiment. Similarly for the ImageNet feature \mathbf{V}^{img} , it is first fed into fully-connected layer with \tanh activation function to project into word space, then encoded by Bi-LSTM producing $h^{img} \in \mathbb{R}^{n_{img} \times 2d}$.

The context-query attention layer [42], [43] is utilized to jointly model the encoded context (e.g. video, subtitle) and query (e.g. question, candidate answers). It feeds a set of context vectors $\mathbf{C} \in \mathbb{R}^{n \times 2d}$ and a set of query vectors $\mathbf{Q} \in \mathbb{R}^{m \times 2d}$ as inputs, and constructs context-to-query attention matrix \mathbf{A} . The context-to-query attention is generated as follows: First, the similarities between each pair of context vector and query vector are computed, producing a similarity matrix $\mathbf{S} \in \mathbb{R}^{n \times m}$, where S_{ij} represents the similarity between i -th context word \mathbf{C}_i and j -th query word \mathbf{Q}_j . Instead of the original trilinear function [42], dot product is utilized to calculate similarity, i.e. $S_{ij} = \mathbf{C}_i \cdot \mathbf{Q}_j$. Then, we normalize each row of similarity matrix by applying the softmax function, producing a matrix $\hat{\mathbf{S}}$. Finally, context-to-query attention which contains the attended query vectors for the entire context is computed as $\mathbf{A} = \hat{\mathbf{S}}\mathbf{Q} \in \mathbb{R}^{n \times 2d}$. The context-to-query attention signifies

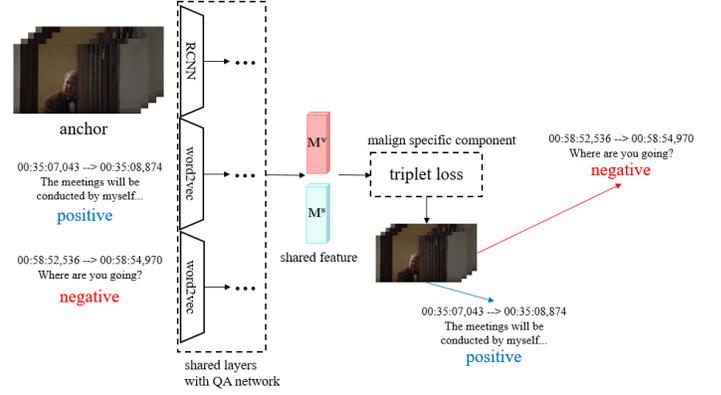


Fig. 3. Illustration of modality alignment network. Modality alignment solves multi-modal metric learning problem where it attempts to find correct association of video and subtitle.

which word in query is most relevant to the each word in context.

Consider the upper stream of the Fig.2 where the visual concept is used as the context for context-query attention layer. The question and candidate answer are considered as the query to generate the context-to-query attentions $\mathbf{A}^{cpt,q}, \mathbf{A}^{cpt,a_i} \in \mathbb{R}^{n_{cpt} \times 2d}$, respectively. The context-to-query attentions are then fused with context as follows:

$$\mathbf{M}^{cpt,a_i} = [H^{cpt}; \mathbf{A}^{cpt,q}; \mathbf{A}^{cpt,a_i}; H^{cpt} \odot \mathbf{A}^{cpt,q}; H^{cpt} \odot \mathbf{A}^{cpt,a_i}], \quad (2)$$

where \odot denotes element-wise multiplication.

Finally, the fused feature vector $\mathbf{M}^{cpt,a_i} \in \mathbb{R}^{n_{cpt} \times 10d}$ is again fed into Bi-LSTM and max-pooled along time to get final feature vector $\mathbf{u}^{cpt,a_i} \in \mathbb{R}^{10d}$ for each candidate answer a_i . The prediction score is obtained by applying linear fully-connected layer on a set of final feature vectors $\{\mathbf{u}^{cpt,a_i}\}_{i=1}^5$. Prediction score for bottom stream can be computed similarly by utilizing the subtitle as the context for context-query attention layer. The prediction score for each stream is summed to get final score and *softmax* function is applied to produce answer probability $\hat{\mathbf{y}}$. The cross-entropy loss function is used to train QA model:

$$L_{QA} = - \sum_{i=1}^5 \mathbf{y}_i \log \hat{\mathbf{y}}_i + (1 - \mathbf{y}_i) \log(1 - \hat{\mathbf{y}}_i). \quad (3)$$

D. Modality Alignment Network

Modality alignment network regards the pairing of video and subtitle as supervision and attempts to predict the correct alignment between two modalities. It shares parameters with the lower layers of the QA network. After the video and subtitles are embedded into the common space forming h^{cpt} and h^S , we denote $\{h_i^{cpt}, h_i^S\}$ as the positive pair and $\{h_i^{cpt}, h_{i^-}^S\}$ as the negative pair of encoded video-subtitle features where i represents the index of element in each mini-batch and i^- represents the index of elements in each mini-batch except i . Each training mini-batch, composed of total B video-subtitle pairs, includes single positive pair and $B-1$ negative samples.

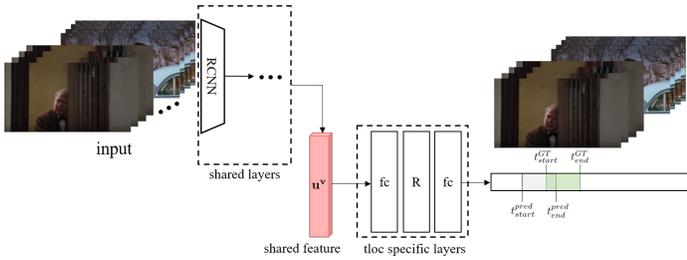


Fig. 4. Illustration of temporal localization network. For simplicity, only upper stream that process video input is drawn. Temporal localization network solves regression problem to predict the start and end time of ground truth moment where the question was generated from.

The objective of modality alignment network is to make the features of positive pair get closer in the embedding space, and the negative pairs get farther. Intuitively, the video-subtitle pair should have a high matching score if its words have a confident support in the video. We formulated this as a metric learning problem. Motivated by Hoffer et al. [44], we utilized max-margin loss to pull features of positive pairs and push features of negative pairs. The distance between positive pair is defined as $d^+ = \|h_i^{cpt} - h_i^S\|_2$ and the distance between negative pair is defined as $d^- = \|h_i^{cpt} - h_i^S\|_2$ where $\|\cdot\|_2$ denotes the l^2 -norm of vector. The modality alignment loss constrains the positive distance d^+ to be smaller than d^- by the margin τ :

$$L_{MA} = \max(0, \tau - d^- + d^+). \quad (4)$$

E. Temporal Localization Network

Temporal localization network localizes the temporal part relevant to the question. The moment in the clips where the question was generated from is regarded as supervision. It shares parameters with the higher layers of the QA network. We formulate the objective of temporal localization network as a regression problem. For each stream, the final feature vectors $\{\mathbf{u}^{a_i}\}_{i=1}^5$ are concatenated and used to regress the start point (t_{start}^{GT}) and end point (t_{end}^{GT}) of ground-truth moment of question generation. We normalized the $t_{start}^{GT}, t_{end}^{GT}$ by the length of the video clip to have the value between 0 and 1. The loss function for the temporal localization network contains two terms:

$$L_{TL} = L_{reg} - L_{overlap}. \quad (5)$$

The first term is straightforward regression loss which is the mean squared error between ground-truth and prediction. The second term is referred to as *overlap loss* which considers the overlap between the ground-truth and prediction. The two terms are formulated as follows:

$$L_{reg} = \|t^{GT} - t^{pred}\|_2, \quad (6)$$

$$L_{overlap} = \frac{O(t^{GT}, t^{pred})}{t_{end}^{GT} - t_{start}^{GT}}. \quad (7)$$

where $O(t^{GT}, t^{pred})$ represents the length of overlap between t^{GT} and t^{pred} which can be formulated as $O(t^{GT}, t^{pred}) = \max(0, \min(t_{end}^{GT}, t_{end}^{pred}) - \max(t_{start}^{GT}, t_{start}^{pred}))$.

F. Multi-Task Ratio Scheduling

The entire network is trained by simultaneously optimizing the aforementioned three loss functions; one for QA network, modality alignment network, and temporal localization network, respectively. The total loss to be minimized during training is the weighted sum of all three losses:

$$L = \alpha_{QA}L_{QA} + \alpha_{MA}L_{MA} + \alpha_{TL}L_{TL}$$

where α_{QA} , α_{MA} , and α_{TL} are the weights for each loss function. In order to control the timing and strength of the objective of each task, multi-task ratio scheduling is proposed to schedule the weights for each loss function. Motivated by curriculum learning [23], simple tasks are focused more at the early stage of training and complex tasks are focused later. Among the designed tasks, the task of modality alignment and temporal localization are easier than question-answering. Therefore initially, the weight for modality alignment α_{MA} is set higher than the other weights to facilitate solving modality alignment task. Then, the weight for temporal localization α_{TL} is set higher and finally, the weight for question-answering α_{QA} is set to the highest to solve the multi-modal video question-answering task.

IV. EXPERIMENTS

This section provides the experimental details and results of our proposed method. First, the benchmark dataset used to train and evaluate the proposed model is introduced. Then, we describe the experimental details. Finally, we provide quantitative results with ablation study.

A. Dataset

The TVQA Benchmark [15] is a multi-modal video question answering dataset. It is collected on 6 long-running TV shows from 3 genres: (1) sitcoms: *The Big Bang Theory*, *How I Met Your Mother*, *Friends*, (2) medical dramas: *Grey's Anatomy*, *House* and (3) crime drama: *Castle*. Total 21,793 short clips of 60/90 seconds are segmented for TVQA [15], accompanied with corresponding subtitles and character names. The questions in TVQA Benchmark is composed of following formal “[What/How/Where/Why/...] _____ [when/before/after] _____?” where second part localizes required point within the video clip for answering the question and first part provides the question on that point. The overall number of multiple-choice question-answer pairs are 152.5k, where train split contains 122,039 QA pairs, validation split contains 15,252 QA pairs and test split contains 7,623 QA pairs. Each QA pair has five candidate answers, but only one of them is correct. The performance of each model is measured by multiple-choice question answering accuracy.

B. Implementation Details

The proposed method was implemented using PyTorch framework. All of the experiments in this paper were performed under CUDA acceleration with single NVIDIA TITAN Xp (12GB of memory) GPU and trained using the Adam optimizer [45] with the learning rate of 0.0003 and mini-batch size of 32.

TABLE I

ACCURACY COMPARISON ON THE VALIDATION AND TEST SET OF TVQA BENCHMARK. THE SYMBOL MEANINGS ARE Q=QUESTION, S=SUBTITLE, V=VIDEO, IMG=IMAGENET FEATURES, REG=REGIONAL VISUAL FEATURES, CPT=VISUAL CONCEPT FEATURES. OUR METHOD ACHIEVES THE STATE-OF-THE-ART PERFORMANCE. THE TEST SET ACCURACY OF OUR PROPOSED METHOD IS OBTAINED FROM ONLINE EVALUATION SERVER. THE SYMBOL ‘-’ INDICATES THAT THE PERFORMANCE IS NOT PROVIDED.

Methods	Video Feature	valid Acc.	test Acc.
Random	-	-	20.00
Longest Answer	-	-	30.41
TVQA S+Q [15]	-	-	63.14
TVQA V+Q [15]	img	-	42.67
TVQA V+Q [15]	reg	-	42.75
TVQA V+Q [15]	cpt	-	43.38
TVQA S+V+Q [15]	img	-	63.57
TVQA S+V+Q [15]	reg	-	63.19
TVQA S+V+Q [15]	cpt	-	65.46
ours S+Q	-	64.36	64.63
ours V+Q	img	42.13	42.79
ours V+Q	cpt	43.45	44.42
ours S+V+Q	img	63.99	64.53
ours S+V+Q	cpt	66.22	67.05

C. Experimental Results

The experimental results are summarized in table I. We compared the performance of our proposed method with the results reported in the TVQA paper [15]. The random baseline shows 20.00% test accuracy for the task of multiple-choice question-answering with 5 candidate answers. The longest answer baseline selects the longest answer for each question. It achieves the performance of 30.41% which indicates that the correct answers tend to be longer than the wrong answers. Note that the validation accuracy of TVQA methods is not reported in the original paper [15].

Our subtitle-only (ours S+Q) method achieves test accuracy of 64.63% which is 34.22% higher than the longest answer baseline and 1.49% higher than the subtitle-only TVQA baseline (TVQA S+Q). Our video-only (ours V+Q) methods achieve the performance of 42.79% and 44.42% for ImageNet and visual concept feature, respectively. Compared to the TVQA baseline, our results on video-only methods obtains 0.12% and 1.04% performance boost. For our uni-model results (ours S+Q and ours V+Q), only temporal localization loss is utilized as an auxiliary loss. Our full model (ours S+V+Q) with ImageNet feature achieves the performance of 64.53% which is 0.96% higher than the TVQA S+V+Q with ImageNet feature. Our full (ours S+V+Q) method with visual concept feature achieves the state-of-the-art result on TVQA dataset with the performance of 67.05% which is 1.59% higher accuracy than the runner-up model, TVQA S+V+Q with visual concept feature. The experimental results verify that multi-task learning can bring additional performance boost especially when the task is complex. Especially, our S+Q model, which only uses subtitle and question but not video, achieved higher performance (64.63%) than TVQA S+V+Q model with ImageNet feature (63.57%). This shows competent

TABLE II

RESULT OF ABLATION STUDY BETWEEN MODEL VARIANTS. BOTH THE VIDEO AND SUBTITLES ARE UTILIZED IN THIS STUDY. THE USE OF TEMPORAL LOCALIZATION LOSS, MODALITY ALIGNMENT LOSS ARE DENOTED AS ‘+ TL’ AND ‘+ MA’. THE LAST COLUMN (DENOTED AS Δ) SHOWS THE PERFORMANCE DROP COMPARED TO THE FULL MODEL.

Methods	Video Feature	valid Acc.	Δ
QA	img	63.14	-0.85
QA + MA	img	63.67	-0.32
QA + TL	img	63.49	-0.50
QA + MA + TL	img	63.99	-
QA	cpt	65.03	-1.19
QA + MA	cpt	65.79	-0.43
QA + TL	cpt	65.64	-0.58
QA + MA + TL	cpt	66.22	-

performance improvement caused by extra supervision from multi-task learning.

For ablation study, we only report the validation accuracy since test accuracy can only be measured through an online evaluation server for finite times. Table II summarizes the results of ablation study. Overall, the visual concept feature gives higher performance than ImageNet feature as reported in [15]. The first block of Table II shows the ablation study with the ImageNet feature and the second block of Table II shows the ablation study with the visual concept feature. Multi-task learning with temporal localization and modality alignment showed a meaningful increase in performance for both cases when used ImageNet and visual concept features. The ablation study suggests that solving auxiliary tasks together with the main task of video question answering can improve performance. Modality alignment brought higher performance gain than temporal localization. The amount of performance gain may differ by implementation detail though, the performance gain can differ by the choice and the scheduling of auxiliary tasks. The more relevant and helpful the auxiliary task is to the main task, the higher the performance gain is.

V. CONCLUSION

In this paper, we proposed a method to gain extra supervision via multi-task learning for multi-modal video question answering. We argue that the existing benchmark datasets on multi-modal video question answering are relatively small to provide sufficient supervision. To overcome this challenge, this paper proposes a multi-task learning method which is composed of three main components: (1) multi-modal video question answering network that answers the question based on the both video and subtitle feature, (2) temporal retrieval network that predicts the time in the video clip where the question was generated from and (3) modality alignment network that solves metric learning problem to find correct association of video and subtitle modalities. Motivated by curriculum learning, multi-task ratio scheduling is proposed to learn easier task earlier to set inductive bias at the beginning of the training. The experiments on publicly available dataset TVQA shows state-of-the-art results, and ablation studies are conducted to prove the statistical validity.

REFERENCES

- [1] J. Weston, S. Chopra, and A. Bordes, "Memory networks," in *International Conference on Learning Representations (ICLR)*, 2015.
- [2] S. Sukhbaatar, J. Weston, R. Fergus *et al.*, "End-to-end memory networks," in *Advances in Neural Information Processing Systems (NIPS)*, 2015.
- [3] S. Min, V. Zhong, R. Socher, and C. Xiong, "Efficient and robust question answering from minimal context over documents," in *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (ACL)*, 2018.
- [4] C. Xiong, V. Zhong, and R. Socher, "DCN+: Mixed objective and deep residual coattention for question answering," in *International Conference on Learning Representations (ICLR)*, 2018.
- [5] M. Malinowski, M. Rohrbach, and M. Fritz, "Ask your neurons: A neural-based approach to answering questions about images," in *IEEE International Conference on Computer Vision (ICCV)*, 2015.
- [6] Z. Yang, X. He, J. Gao, and A. Smola, "Stacked attention networks for image question answering," in *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2016.
- [7] H. Nam, J.-W. Ha, and J. Kim, "Dual attention networks for multimodal reasoning and matching," in *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2017.
- [8] H. Ben-Younes, R. Cadène, N. Thome, and M. Cord, "Mutan: Multimodal tucker fusion for visual question answering," in *IEEE International Conference on Computer Vision (ICCV)*, 2017.
- [9] P. Anderson, X. He, C. Buehler, D. Teney, M. Johnson, S. Gould, and L. Zhang, "Bottom-up and top-down attention for image captioning and visual question answering," in *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2018.
- [10] Y. Jang, Y. Song, Y. Yu, Y. Kim, and G. Kim, "Tgif-qa: Toward spatio-temporal reasoning in visual question answering," in *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2017.
- [11] D. Xu, Z. Zhao, J. Xiao, F. Wu, H. Zhang, X. He, and Y. Zhuang, "Video question answering via gradually refined attention over appearance and motion," in *ACM Multimedia*.
- [12] J. Gao, R. Ge, K. Chen, and R. Nevatia, "Motion-appearance co-memory networks for video question answering," in *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2018.
- [13] Y. Yu, J. Kim, and G. Kim, "A joint sequence fusion model for video question answering and retrieval," in *European Conference on Computer Vision (ECCV)*, 2018.
- [14] M. Tapaswi, Y. Zhu, R. Stiefelwagen, A. Torralba, R. Urtasun, and S. Fidler, "Movieqa: Understanding stories in movies through question-answering," in *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2016.
- [15] J. Lei, L. Yu, M. Bansal, and T. L. Berg, "Tvqa: Localized, compositional video question answering," in *Conference on Empirical Methods in Natural Language Processing (EMNLP)*, 2018.
- [16] S. Na, S. Lee, J. Kim, and G. Kim, "A read-write memory network for movie story understanding," in *IEEE International Conference on Computer Vision (ICCV)*, 2017.
- [17] J. Liang, L. Jiang, L. Cao, L.-J. Li, and A. Hauptmann, "Focal visual-text attention for visual question answering," in *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2018.
- [18] B. Wang, Y. Xu, Y. Han, and R. Hong, "Movie question answering: Remembering the textual cues for layered visual contents," in *AAAI Conference on Artificial Intelligence*, 2018.
- [19] K.-M. Kim, S.-H. Choi, and B.-T. Zhang, "Multimodal dual attention memory for video story question answering," in *European Conference on Computer Vision (ECCV)*, 2018.
- [20] K. Kim, M. Heo, S. Choi, and B. Zhang, "Deepstory: Video story QA by deep embedded memory networks," in *International Joint Conference on Artificial Intelligence, (IJCAI)*, 2017.
- [21] A. Agrawal, J. Lu, S. Antol, M. Mitchell, C. L. Zitnick, D. Parikh, and D. Batra, "Vqa: Visual question answering," *Int. J. Comput. Vision*, vol. 123, no. 1, pp. 4–31, May 2017.
- [22] J. Gao, C. Sun, Z. Yang, and R. Nevatia, "Tall: Temporal activity localization via language query," in *IEEE International Conference on Computer Vision (ICCV)*, 2017.
- [23] Y. Bengio, J. Louradour, R. Collobert, and J. Weston, "Curriculum learning," in *Proceedings of the 26th Annual International Conference on Machine Learning (ICML)*, 2009.
- [24] R. Caruana, "Multitask learning: A knowledge-based source of inductive bias," in *Proceedings of the Tenth International Conference on Machine Learning (ICML)*, 1993.
- [25] R. Girshick, "Fast r-cnn," in *IEEE International Conference on Computer Vision (ICCV)*.
- [26] R. Collobert and J. Weston, "A unified architecture for natural language processing: Deep neural networks with multitask learning," in *Proceedings of the 25th International Conference on Machine Learning (ICML)*, 2008.
- [27] J. Kim and C. D. Yoo, "Deep partial person re-identification via attention model," in *International Conference on Image Processing (ICIP)*, 2017.
- [28] S. Ren, K. He, R. Girshick, and J. Sun, "Faster r-cnn: Towards real-time object detection with region proposal networks," in *Advances in Neural Information Processing Systems (NIPS)*, 2015.
- [29] T. Xiao, S. Li, B. Wang, L. Lin, and X. Wang, "Joint detection and identification feature learning for person search," in *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2017.
- [30] Y. Li, N. Duan, B. Zhou, X. Chu, W. Ouyang, X. Wang, and M. Zhou, "Visual question generation as dual task of visual question answering," in *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2018.
- [31] A. Karpathy and L. Fei-Fei, "Deep visual-semantic alignments for generating image descriptions," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 39, no. 4, pp. 664–676, 2017.
- [32] L. Castrejón, Y. Aytar, C. Vondrick, H. Pirsiavash, and A. Torralba, "Learning aligned cross-modal representations from weakly aligned data," in *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2016.
- [33] Y. Yu, J. Kim, and G. Kim, "A joint sequence fusion model for video question answering and retrieval," in *European Conference on Computer Vision (ECCV)*, 2018.
- [34] L. A. Hendricks, O. Wang, E. Shechtman, J. Sivic, T. Darrell, and B. C. Russell, "Localizing moments in video with natural language," in *IEEE International Conference on Computer Vision (ICCV)*, 2017.
- [35] J. Gao, Z. Yang, C. Sun, K. Chen, and R. Nevatia, "TURN TAP: temporal unit regression network for temporal action proposals," in *IEEE International Conference on Computer Vision (ICCV)*, 2017.
- [36] A. Vaswani, N. Shazeer, N. Parmar, J. Uszkoreit, L. Jones, A. N. Gomez, L. u. Kaiser, and I. Polosukhin, "Attention is all you need," in *Advances in Neural Information Processing Systems (NIPS)*, 2017.
- [37] K. He, X. Zhang, S. Ren, and J. Sun, "Deep residual learning for image recognition," *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2016.
- [38] J. Deng, W. Dong, R. Socher, L. Li, K. Li, and L. Fei-Fei, "Imagenet: A large-scale hierarchical image database," in *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2009.
- [39] X. Yin and V. Ordonez, "Obj2text: Generating visually descriptive language from object layouts," in *Conference on Empirical Methods in Natural Language Processing (EMNLP)*.
- [40] R. Krishna, Y. Zhu, O. Groth, J. Johnson, K. Hata, J. Kravitz, S. Chen, Y. Kalantidis, L.-J. Li, D. A. Shamma, M. S. Bernstein, and L. Fei-Fei, "Visual genome: Connecting language and vision using crowdsourced dense image annotations," *International Journal of Computer Vision*, vol. 123, no. 1, pp. 32–73, May 2017.
- [41] J. Pennington, R. Socher, and C. D. Manning, "Glove: Global vectors for word representation," in *Conference on Empirical Methods in Natural Language Processing (EMNLP)*, 2014.
- [42] A. F. H. H. Minjoon Seo, Aniruddha Kembhavi, "Bidirectional attention flow for machine comprehension," in *International Conference on Learning Representations (ICLR)*, 2017.
- [43] A. W. Yu, D. Dohan, Q. Le, T. Luong, R. Zhao, and K. Chen, "Fast and accurate reading comprehension by combining self-attention and convolution," in *International Conference on Learning Representations (ICLR)*, 2018.
- [44] E. Hoffer and N. Ailon, "Deep metric learning using triplet network," in *International Conference on Learning Representations Workshop Track (ICLR)*, 2015.
- [45] D. P. Kingma and J. Ba, "Adam: A method for stochastic optimization," *CoRR*, vol. abs/1412.6980, 2014. [Online]. Available: <http://arxiv.org/abs/1412.6980>