

End-to-End Audio Visual Scene-Aware Dialog using Multimodal Attention-Based Video Features

Hori, C.; Alamri, H.; Wang, J.; Wichern, G.; Hori, T.; Cherian, A.; Marks, T.K.; Cartillier, V.;
Lopes, R.; Das, A.; Essa, I.; Batra, D.; Parikh, D.

TR2018-085 July 13, 2018

Abstract

Dialog systems need to understand dynamic visual scenes in order to have conversations with users about the objects and events around them. Scene-aware dialog systems for real-world applications could be developed by integrating state-of-the-art technologies from multiple research areas, including: end-to-end dialog technologies, which generate system responses using models trained from dialog data; visual question answering (VQA) technologies, which answer questions about images using learned image features; and video description technologies, in which descriptions/captions are generated from videos using multimodal information. We introduce a new dataset of dialogs about videos of human behaviors. Each dialog is a typed conversation that consists of a sequence of 10 question-and-answer (QA) pairs between two Amazon Mechanical Turk (AMT) workers. In total, we collected dialogs on 9,000 videos. Using this new dataset, we trained an end-to-end conversation model that generates responses in a dialog about a video. Our experiments demonstrate that using multimodal features that were developed for multimodal attention-based video description enhances the quality of generated dialog about dynamic scenes (videos). Our dataset, model code and pretrained models will be publicly available for a new Video Scene-Aware Dialog challenge.

arXiv

This work may not be copied or reproduced in whole or in part for any commercial purpose. Permission to copy in whole or in part without payment of fee is granted for nonprofit educational and research purposes provided that all such whole or partial copies include the following: a notice that such copying is by permission of Mitsubishi Electric Research Laboratories, Inc.; an acknowledgment of the authors and individual contributions to the work; and all applicable portions of the copyright notice. Copying, reproduction, or republishing for any other purpose shall require a license with payment of fee to Mitsubishi Electric Research Laboratories, Inc. All rights reserved.

End-to-End Audio Visual Scene-Aware Dialog using Multimodal Attention-Based Video Features

Chiori Hori[†], Huda Alamri^{*†}, Jue Wang[†], Gordon Wichern[†],
Takaaki Hori[†], Anoop Cherian[†], Tim K. Marks[†],
Vincent Cartillier^{*}, Raphael Gontijo Lopes^{*}, Abhishek Das^{*},
Irfan Essa^{*}, Dhruv Batra^{*} Devi Parikh^{*},

[†]Mitsubishi Electric Research Laboratories (MERL), Cambridge, MA, USA

^{*}School of Interactive Computing, Georgia Tech

Abstract

Dialog systems need to understand dynamic visual scenes in order to have conversations with users about the objects and events around them. Scene-aware dialog systems for real-world applications could be developed by integrating state-of-the-art technologies from multiple research areas, including: end-to-end dialog technologies, which generate system responses using models trained from dialog data; visual question answering (VQA) technologies, which answer questions about images using learned image features; and video description technologies, in which descriptions/captions are generated from videos using multimodal information. We introduce a new dataset of dialogs about videos of human behaviors. Each dialog is a typed conversation that consists of a sequence of 10 question-and-answer (QA) pairs between two Amazon Mechanical Turk (AMT) workers. In total, we collected dialogs on $\sim 9,000$ videos. Using this new dataset, we trained an end-to-end conversation model that generates responses in a dialog about a video. Our experiments demonstrate that using multimodal features that were developed for multimodal attention-based video description enhances the quality of generated dialog about dynamic scenes (videos). Our dataset, model code and pretrained models will be publicly available for a new Video Scene-Aware Dialog challenge.

1 Introduction

Spoken dialog technologies have been applied in real-world human-machine interfaces including smart phone digital assistants, car navigation systems, voice-controlled smart speakers, and human-facing robots [1, 2, 3]. Generally, a dialog system consists of a pipeline of data processing modules, including automatic speech recognition, spoken language understanding, dialog management, sentence generation, and speech synthesis. However, all of these modules require significant hand engineering and domain knowledge for training. Recently, end-to-end dialog systems have been gathering attention, and they obviate this need for expensive hand engineering to some extent. In end-to-end approaches, dialog models are trained using only paired input and output sentences, without relying on pre-designed data processing modules or intermediate internal data representations such as concept tags and slot-value pairs. End-to-end systems can be trained to directly map from a user’s utterance to a system response sentence and/or action. This significantly reduces the data preparation and system development cost. Several types of sequence-to-sequence models have been applied to end-to-end dialog systems, and it has been shown that they can be trained in a completely data-driven manner. End-to-end approaches have also been shown to better handle flexible conversations between the user and the system by training the model on large conversational datasets [4, 5].

In these applications, however, all conversation is triggered by user speech input, and the contents of system responses are limited by the training data (a set of dialogs). Current dialog systems cannot understand dynamic scenes using multimodal sensor-based input

such as vision and non-speech audio, so machines using such dialog systems cannot have a conversation about what’s going on in their surroundings. To develop machines that can carry on a conversation about objects and events taking place around the machines or the users, dynamic scene-aware dialog technology is essential.

To interact with humans about visual information, systems need to understand both visual scenes and natural language inputs. One naive approach could be a pipeline system in which the output of a visual description system is used as an input to a dialog system. In this cascaded approach, semantic frames such as "who" is doing "what" and "where" must be extracted from the video description results. The prediction of frame type and the value of the frame must be trained using annotated data. In contrast, the recent revolution of neural network models allows us to combine different modules into a single end-to-end differentiable network. We can simultaneously input video features and user utterances into an encoder-decoder-based system whose outputs are natural-language responses.

Using this end-to-end framework, *visual question answering* (VQA) has been intensively researched in the field of computer vision [6, 7, 8]. The goal of VQA is to generate answers to questions about an imaged scene, using the information present in a single static image. As a further step towards conversational visual AI, the new task of *visual dialog* was introduced [9], in which an AI agent holds a meaningful dialog with humans about an image using natural, conversational language [10]. While VQA and visual dialog take significant steps towards human-machine interaction, they only consider a single static image. To capture the semantics of dynamic scenes, recent research has focused on *video description* (natural-language descriptions of videos). The state of the art in video description uses a multimodal attention mechanism that selectively attends to different input modalities (feature types), such as spatiotemporal motion features and audio features, in addition to temporal attention [11].

In this paper, we propose a new research target, a dialog system that can discuss dynamic scenes with humans, which lies at the intersection of multiple avenues of research in natural language processing, computer vision, and audio processing. To advance this goal, we introduce a new model that incorporates technologies for multimodal attention-based video description into an end-to-end dialog system. We also introduce a new dataset of human dialogues about videos. We are making our dataset, code, and model publicly available for a new Video Scene-Aware Dialog Challenge.

2 Audio Visual Scene-Aware Dialog Dataset

We collected text-based conversations data about short videos for Audio Visual Scene-Aware Dialog (AVSD) as described in [12] using from an existing video description dataset, Charades [13], for Dialog System Technology Challenge the 7th edition (DSTC7)¹. Charades is an untrimmed and multi-action dataset, containing 11,848 videos split into 7985 for training, 1863 for validation, and 2,000 for testing. It has 157 action categories, with several fine-grained actions. Further, this dataset also provides 27,847 textual descriptions for the videos, each video is associated with 1–3 sentences. As these textual descriptions are only available in the training and validation set, we report evaluation results on the validation set.

The data collection paradigm for dialogs was similar to the one described in [9], in which for each image, two different Mechanical Turk workers interacted via a text interface to yield a dialog. In [9], each dialog consisted of a sequence of questions and answers about an image. In the video scene-aware dialog case, two Amazon Mechanical Turk (AMT) workers had a discussion about events in a video. One of the workers played the role of an answerer who had already watched the video. The answerer answered questions asked by another AMT worker – the questioner. The questioner was not allowed to watch the whole video but only the first, middle and last frames of the video which were single static images. After having a conversation to ask about the events that happened between the frames through 10 rounds of QA, the questioner summarized the events in the video as a description.

In total, we collected dialogs for 7043 videos from the Charades training set and all of the validation set (1863 videos). Since we did not have scripts for the test set, we split the validation set into 732

¹<http://workshop.colips.org/dstc7/call.html>

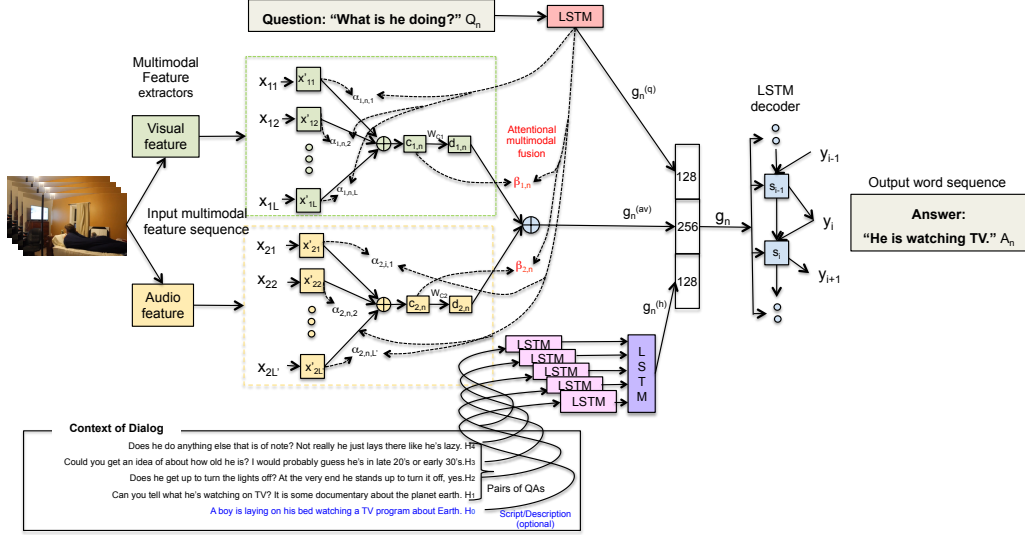


Figure 1: Our multimodal-attention based video scene-aware dialog system

and 733 videos and used them as our validation and test sets respectively. See Table 1 for statistics. The average numbers of words per question and answer are 8 and 10, respectively.

Table 1: Video Scene-aware Dialog Dataset on Charades

	training	validation	test
#dialogs	7043	732	733
#turns	123,480	14,680	14,660
#words	1,163,969	138,314	138,790

3 Video Scene-aware Dialog System

We built an end-to-end dialog system that can generate answers in response to user questions about events in a video sequence. Our architecture is similar to the Hierarchical Recurrent Encoder in Das *et al.* [9]. The question, visual features, and the dialog history are fed into corresponding LSTM-based encoders to build up a context embedding, and then the outputs of the encoders are fed into a LSTM-based decoder to generate an answer. The history consists of encodings of QA pairs. We feed multimodal attention-based video features into the LSTM encoder instead of single static image features. Figure 1 shows the architecture of our video scene-aware dialog system.

3.1 End-to-end Conversation Modeling

This section explains the neural conversation model of [4], which is designed as a sequence-to-sequence mapping process using recurrent neural networks (RNNs). Let X and Y be input and output sequences, respectively. The model is used to compute posterior probability distribution $P(Y|X)$. For conversation modeling, X corresponds to the sequence of previous sentences in a conversation, and Y is the system response sentence we want to generate. In our model, both X and Y are sequences of words. X contains all of the previous turns of the conversation, concatenated in sequence, separated by markers that indicate to the model not only that a new turn has started, but which speaker said that sentence. The most likely hypothesis of Y is obtained as

$$\hat{Y} = \arg \max_{Y \in \mathcal{V}^*} P(Y|X) \quad (1)$$

$$= \arg \max_{Y \in \mathcal{V}^*} \prod_{m=1}^{|Y|} P(y_m | y_1, \dots, y_{m-1}, X), \quad (2)$$

where \mathcal{V}^* denotes a set of sequences of zero or more words in system vocabulary \mathcal{V} .

Let X be word sequence x_1, \dots, x_T and Y be word sequence y_1, \dots, y_M . The encoder network is used to obtain hidden states h_t for $t = 1, \dots, T$ as:

$$h_t = \text{LSTM}(x_t, h_{t-1}; \theta_{enc}), \quad (3)$$

where h_0 is initialized with a zero vector. $\text{LSTM}(\cdot)$ is a LSTM function with parameter set θ_{enc} .

The decoder network is used to compute probabilities $P(y_m | y_1, \dots, y_{m-1}, X)$ for $m = 1, \dots, M$ as:

$$s_0 = h_T \quad (4)$$

$$s_m = \text{LSTM}(y_{m-1}, s_{m-1}; \theta_{dec}) \quad (5)$$

$$P(\mathbf{y} | y_1, \dots, y_{m-1}, X) = \text{softmax}(W_o s_m + b_o), \quad (6)$$

where y_0 is set to $\langle \text{eos} \rangle$, a special symbol representing the end of sequence. s_m is the m -th decoder state. θ_{dec} is a set of decoder parameters, and W_o and b_o are a matrix and a vector. In this model, the initial decoder state s_0 is given by the final encoder state h_T as in Eq. (4), and the probability is estimated from each state s_m . To efficiently find \hat{Y} in Eq. (1), we use a beam search technique since it is computationally intractable to consider all possible Y .

In the scene-aware-dialog scenario, a scene context vector including audio and visual features is also fed to the decoder. We modify the LSTM in Eqs. (4)–(6) as

$$s_{n,0} = \bar{\mathbf{0}} \quad (7)$$

$$s_{n,m} = \text{LSTM}([y_{n,m-1}^T, g_n^T]^T, s_{n,m-1}; \theta_{dec}), \quad (8)$$

$$P(\mathbf{y}_n | y_{n,1}, \dots, y_{n,m-1}, X) = \text{softmax}(W_o s_{n,m} + b_o), \quad (9)$$

where g_n is the concatenation of question encoding $g_n^{(q)}$, audio-visual encoding $g_n^{(av)}$ and history encoding $g_n^{(h)}$ for generating the n -th answer $A_n = y_{n,1}, \dots, y_{n,|Y_n|}$. Note that unlike Eq. (4), we feed all contextual information to the LSTM at every prediction step. This architecture is more flexible since the dimensions of encoder and decoder states can be different.

$g_n^{(q)}$ is encoded by another LSTM for the n -th question, and $g_n^{(h)}$ is encoded with hierarchical LSTMs, where one LSTM encodes each question-answer pair and then the other LSTM summarizes the question-answer encodings into $g_n^{(h)}$. The audio-visual encoding is obtained by multi-modal attention described in the next section.

3.2 Multimodal-attention based Video Features

To predict a word sequence in video description, prior work [14] extracted content vectors from image features of VGG-16 and spatiotemporal motion features of C3D, and combined them into one vector in the fusion layer as:

$$g_n^{(av)} = \tanh\left(\sum_{k=1}^K d_{k,n}\right), \quad (10)$$

where

$$d_{k,n} = W_{ck}^{(\lambda_D)} c_{k,n} + b_{ck}^{(\lambda_D)}, \quad (11)$$

and $c_{k,n}$ is a context vector obtained using the k -th input modality.

We call this approach Naïve Fusion, in which multimodal feature vectors are combined using projection matrices W_{ck} for K different modalities (input sequences x_{k1}, \dots, x_{kL} for $k = 1, \dots, K$).

To fuse multimodal information, prior work [11] proposed method extends the attention mechanism. We call this fusion approach *multimodal attention*. The approach can pay attention to specific modalities of input based on the current state of the decoder to predict the word sequence in video description. The number of modalities indicating the number of sequences of input feature vectors is denoted by K .

The following equation shows an approach to perform the attention-based feature fusion:

$$g_n^{(av)} = \tanh\left(\sum_{k=1}^K \beta_{k,n} d_{k,n}\right). \quad (12)$$

The similar mechanism for temporal attention is applied to obtain the multimodal attention weights $\beta_{k,n}$:

$$\beta_{k,n} = \frac{\exp(v_{k,n})}{\sum_{\kappa=1}^K \exp(v_{\kappa,n})}, \quad (13)$$

where

$$v_{k,n} = w_B^T \tanh(W_B g_n^{(q)} + V_{Bk} c_{k,n} + b_{Bk}). \quad (14)$$

Here the multimodal attention weights are determined by question encoding $g_n^{(q)}$ and the context vector of each modality $c_{k,n}$ as well as temporal attention weights in each modality. W_B and V_{Bk} are matrices, w_B and b_{Bk} are vectors, and $v_{k,n}$ is a scalar. The multimodal attention weights can change according to the question encoding and the feature vectors (shown in Figure 1).

This enables the decoder network to attend to a different set of features and/or modalities when predicting each subsequent word in the description. Naïve fusion can be considered a special case of Attentional fusion, in which all modality attention weights, $\beta_{k,n}$, are constantly 1.

4 Experiments for Multimodal attention-based Video Features

To select best video features for the video scene-aware dialog system, we firstly evaluate the performance of video description using multimodal attention-based video features in this paper.

4.1 Datasets

We evaluated our proposed feature fusion using the MSVD (YouTube2Text) [15], MSR-VTT [16], and Charades [13] video data sets.

- MSVD (YouTube2Text) covers a wide range of topics including sports, animals, and music. We applied the same condition defined by [15]: a training set of 1,200 video clips, a validation set of 100 clips, and a test set of the remaining 670 clips.
- MSR-VTT is split into training, validation, and testing sets of 6,513, 497, and 2,990 clips respectively. However, approximately 12 of the MSR-VTT videos on YouTube have been removed. We used the available data consists of 5,763, 419, and 2,616 clips for train, validation, and test respectively defined by [11].
- Charades [13] is split into 7985 clips for training and 1863 clips for validation. provides 27,847 textual descriptions for the videos, As these textual descriptions are only available in the training and validation set, we report the evaluation results on the validation set.

Details of textual descriptions are summarized in Table 2.

Table 2: Sizes of textual descriptions in MSVD (YouTube2Text), MSR-VTT and Charades

Dataset	#Clips	#Description	#Descriptions per clip	#Word	Vocabulary size
MSVD	1,970	80,839	41.00	8.00	13,010
MSR-VTT	10,000	200,000	20.00	9.28	29,322
Charades	9,848	16,140	1.64	13.04	2,582

4.2 Video Processing

We used a sequence of 4096-dimensional feature vectors of the output from the fully-connected fc7 layer of a VGG-16 network pretrained on the ImageNet dataset for the image features.

The pretrained C3D [17] model is used to generate features for model motion and short-term spatiotemporal activity. The C3D network reads sequential frames in the video and outputs a fixed-length feature vector every 16 frames. 4096-dimensional features of activation vectors from fully-connected fc6-1 layer was applied to spatiotemporal features.

In addition to the VGG-16 and C3D features, we also adopted the state-of-the-art I3D features [18], spatiotemporal features that were developed for action recognition. The I3D model inflates the 2D filters and pooling kernels in the Inception V3 network along their temporal dimension, building 3D spatiotemporal ones. We used the output from the "Mixed_5c" layer of the I3D network to be used as video features in our framework. As a pre-processing step, we normalized all the video features to have zero mean and unit norm; the mean was computed over all the sequences in the training set for the respective feature.

In the experiments in this paper, we treated I3D-rgb (I3D features computed on a stack of 16 video frame images) and I3D-flow (I3D features computed on a stack of 16 frames of optical flow fields) as two separate modalities that are input to our multimodal attention model. To emphasize this, we refer to I3D in the results tables as I3D (rgb-flow).

4.3 Audio Processing

While the original MSVD (YouTube2Text) dataset does not contain audio features, we were able to collect audio data for 1,649 video clips (84% of the dataset) from the video URLs. In our previous work on multimodal attention for video description, we used two different types of audio features: concatenated mel-frequency cepstral coefficient (MFCC) features [19], and SoundNet [20] features [21]. In this paper, we also evaluate features extracted using a new state-of-the-art model, Audio Set VGGish [22].

Inspired by the VGG image classification architecture (Configuration A without the last group of convolutional/pooling layers), the Audio Set VGGish model operates on 0.96 s log Mel spectrogram patches extracted from 16 kHz audio, and outputs a 128-dimensional embedding vector. The model was trained to predict an ontology of labels from only the audio tracks of millions of YouTube videos. In this work, we overlap frames of input to the VGGish network by 50%, meaning an Audio Set VGGish feature vector is output every 0.48 s. For SoundNet [20], in which a fully convolutional architecture was trained to predict scenes and objects using a pretrained image model as a teacher, we take as input to the audio encoder the output of the second-to-last convolutional layer, which gives a 1024-dimensional feature vector every 0.67 s, and has a receptive field of approximately 4.16 s. For raw MFCC features, sequences of 13-dimensional MFCC features are extracted from 50 ms windows, every 25 ms, and then 20 consecutive frames are concatenated into a 260-dimensional vector and normalized to zero mean/unit variance (computed over the training set) and used as input to the BLSTM audio encoder.

4.4 Experimental Setup

The caption generation model, i.e., the decoder network, is trained to minimize the cross entropy criterion using the training set. Image features and deep audio features (SoundNet and VGGish) are fed to the decoder network through one projection layer of 512 units, while MFCC audio features are fed to a BLSTM encoder (one projection layer of 512 units and bidirectional LSTM layers of 512 cells) followed by the decoder network. The decoder network has one LSTM layer with 512 cells. Each word is embedded to a 256-dimensional vector when it is fed to the LSTM layer. In this video description task, we used L2 regularization for all experimental conditions and used RMSprop optimization.

4.5 Evaluation

The quality of the automatically generated sentences will be evaluated with objective measures to measure the similarity between the generated sentences and ground truth sentences. We will use the evaluation code for MS COCO caption generation² for objective evaluation of system outputs, which is a publicly available tool supporting various automated metrics for natural language generation such as BLEU, METEOR, ROUGE_L, and CIDEr.

²<https://github.com/tylin/coco-caption>

Table 3: Video description evaluation results on the MSVD (YouTube2Text) test set.

MSVD (YouTube2Text) Full Dataset					
Modalities (feature types)			Evaluation metric		
Image	Spatiotemporal	Audio	BLEU4	METEOR	CIDEr
VGG-16	C3D		0.524	0.320	0.688
VGG-16	C3D	MFCC	0.539	0.322	0.674
	I3D (rgb-flow)		0.525	0.330	0.742
		MFCC	0.527	0.325	0.702
	I3D (rgb-flow)	SoundNet	0.529	0.319	0.719
		VGGish	0.554	0.332	0.743

Table 4: Video description evaluation results on MSR-VTT Subset. Approximately 12% of the MSR-VTT videos have been removed from YouTube, so we train and test on the remaining Subset of MSR-VTT videos that we were able to download. The normalization for the visual features was not applied to MSR-VTT in this experiments.

MSR-VTT Subset					
Modalities (feature types)			Evaluation metric		
Image	Spatiotemporal	Audio	BLEU4	METEOR	CIDEr
VGG-16	C3D	MFCC	0.397	0.255	0.400
	I3D (rgb-flow)		0.347	0.241	0.349
		MFCC	0.364	0.253	0.393
	I3D (rgb-flow)	SoundNet	0.366	0.246	0.387
		VGGish	0.390	0.263	0.417

4.6 Results and Discussion

Tables 3, 4, and 5 show the evaluation results on the MSVD (YouTube2Text), MSR-VTT Subset, and Charades datasets. The I3D spatiotemporal features outperformed the combination of VGG-16 image features and C3D spatiotemporal features. We also tried a combination of VGG-16 image features plus I3D spatiotemporal features, but we do not report those results because they did not improve performance over I3D features alone. We believe this is because I3D features already include enough image information for the video description task. In comparison to C3D, which uses the VGG-16 base architecture and was trained on the Sports-1M dataset [23], I3D uses a more powerful Inception-V3 network architecture and was trained on the larger (and cleaner) Kinetics [24] dataset. As a result, I3D has demonstrated state-of-the-art performance for the task of human action recognition in video sequences [18]. Further, the Inception-V3 architecture has significantly fewer network parameters than the VGG-16 network, making it more efficient.

In terms of audio features, the Audio Set VGGish model provided the best performance. While we expected the deep features (SoundNet and VGGish) to provide improved performance compared to MFCC, there are several possibilities as to why VGGish performed better than SoundNet. First, the VGGish model was trained on more data, and had audio specific labels, whereas SoundNet used pre-trained image classification networks to provide labels for training the audio network. Second, the large Audio Set ontology used to train VGGish likely provides the ability to learn features more relevant to text descriptions than the broad scene/object labels used by SoundNet.

Table 5: Video description evaluation results on Charades.

Charades Dataset					
Modalities (feature types)			Evaluation metric		
Image	Spatiotemporal	Audio	BLEU4	METEOR	CIDEr
	I3D (rgb-flow)		0.094	0.149	0.236
		MFCC	0.098	0.156	0.268
	I3D (rgb-flow)	SoundNet	-	-	-
		VGGish	0.100	0.157	0.270

Since it is intractable to enumerate all possible word sequences in vocabulary \mathcal{V} , we usually limit them to the n -best hypotheses generated by the system. Although in theory the distribution $P(Y'|X)$ should be the true distribution, we instead estimate it using the encoder-decoder model.

Table 6: System response generation evaluation results with objective measures.

Input features	Attentional fusion	BLEU1	BLEU2	BLEU3	BLEU4	METEOR	ROUGE_L	CIDEr
QA	-	0.236	0.142	0.094	0.065	0.101	0.257	0.595
QA + Captions	-	0.245	0.152	0.103	0.073	0.109	0.271	0.705
QA + VGG16	-	0.231	0.141	0.095	0.067	0.102	0.259	0.618
QA + I3D	no	0.246	0.153	0.104	0.073	0.109	0.269	0.680
QA + I3D	yes	0.250	0.157	0.108	0.077	0.110	0.274	0.724
QA + I3D + VGGish	no	0.249	0.155	0.106	0.075	0.110	0.275	0.701
QA + I3D + VGGish	yes	0.256	0.161	0.109	0.078	0.113	0.277	0.727

5 Experiments for Video-scene-aware Dialog

In this paper, we extended an end-to-end dialog system to scene-aware dialog with multimodal fusion. As shown in Fig. 1, we embed the video and audio features selected in Section 2.

5.1 Conditions

We evaluated our proposed system with the dialog data for Charades we collected. Table 1 shows the size of each data set. We compared the performance between models trained from various combinations of the QA text, visual and audio features. In addition, we tested an efficacy of multimodal-attention mechanism for dialg response generation. We employed an ADAM optimizer [25] with the cross-entropy criterion and iterated the training process up to 20 epochs. For each of the encoder-decoder model types, we selected the model with the lowest perplexity on the expanded development set.

We used the parameters of the LSTMs with #layer=2 and #cells=128 for encoding history and question sentences. Video features were projected to 256 dimensional feature space before modality fusion. The decoder LSTM had a structure of #layer=2 and #cells=128 as well.

5.2 Evaluation Results

Table 6 shows the response sentence generation performance of our models, training and decoding methods using objective measures, BLEU1-4, METEOR, ROUGE_L, and CIDEr, which were computed with the evaluation code for MS COCO caption generation as done for video description. We investigated different input features including question-answering dialog history plus last question (QA), human-annotated captions (Captions), video features of VGG16 or I3D rgb and flow features (I3D), and audio features (VGGish).

First we evaluated response generation quality with only QA features as a baseline without any video scene features. Then, we added the caption features to QA, and the performance improved significantly. This is because each caption provided the scene information in natural language and helped the system answer the question correctly. However, such human annotations are not available for real systems.

Next we added VGG16 features to QA, but they did not increase the evaluation scores from those of QA-only features. This result indicates that QA+VGG16 is not enough to let the system generate better responses than those of QA+Captions. After that, we replaced VGG16 with I3D, and obtained a certain improvement from the QA-only case. As in the video description, it has been shown that the I3D features are also useful for scene-aware dialog. Furthermore, we applied the multi-modal attention mechanism (attentional fusion) for I3D rgb and flow features, and obtained further improvement in all the metrics.

Finally, we examined the efficacy of audio features. The table shows that VGGish obviously contributed to increasing the response quality especially when using the attentional fusion. The following example of system response was obtained with or without VGGish features, which worked better for the questions regarding audios:

Question:	was there audio ?
Ground truth:	there is audio , i can hear music and background noise .
I3D:	no , there is no sound in the video .
I3D+VGGish:	yes there is sound in the video .

6 Conclusion

In this paper, we propose a new research target, a dialog system that can discuss dynamic scenes with humans, which lies at the intersection of multiple avenues of research in natural language processing, computer vision, and audio processing. To advance this goal, we introduce a new model that incorporates technologies for multimodal attention-based video description into an end-to-end dialog system. We also introduce a new dataset of human dialogues about videos. Using this new dataset, we trained an end-to-end conversation model that generates system responses in a dialog about an input video. Our experiments demonstrate that using multimodal features that were developed for multimodal attention-based video description enhances the quality of generated dialog about dynamic scenes. We are making our data set and model publicly available for a new Video Scene-Aware Dialog challenge.

References

- [1] Michael F McTear, “Spoken dialogue technology: enabling the conversational user interface,” *ACM Computing Surveys (CSUR)*, vol. 34, no. 1, pp. 90–169, 2002.
- [2] Steve J Young, “Probabilistic methods in spoken–dialogue systems,” *Philosophical Transactions of the Royal Society of London A: Mathematical, Physical and Engineering Sciences*, vol. 358, no. 1769, pp. 1389–1402, 2000.
- [3] Victor Zue, Stephanie Seneff, James R Glass, Joseph Polifroni, Christine Pao, Timothy J Hazen, and Lee Hetherington, “Juplter: a telephone-based conversational interface for weather information,” *IEEE Transactions on speech and audio processing*, vol. 8, no. 1, pp. 85–96, 2000.
- [4] Oriol Vinyals and Quoc Le, “A neural conversational model,” *arXiv preprint arXiv:1506.05869*, 2015.
- [5] Ryan Lowe, Nissan Pow, Iulian Serban, and Joelle Pineau, “The ubuntu dialogue corpus: A large dataset for research in unstructured multi-turn dialogue systems,” *arXiv preprint arXiv:1506.08909*, 2015.
- [6] Stanislaw Antol, Aishwarya Agrawal, Jiasen Lu, Margaret Mitchell, Dhruv Batra, C. Lawrence Zitnick, and Devi Parikh, “VQA: Visual Question Answering,” in *International Conference on Computer Vision (ICCV)*, 2015.
- [7] Peng Zhang, Yash Goyal, Douglas Summers-Stay, Dhruv Batra, and Devi Parikh, “Yin and Yang: Balancing and answering binary visual questions,” in *Conference on Computer Vision and Pattern Recognition (CVPR)*, 2016.
- [8] Yash Goyal, Tejas Khot, Douglas Summers-Stay, Dhruv Batra, and Devi Parikh, “Making the V in VQA matter: Elevating the role of image understanding in Visual Question Answering,” in *Conference on Computer Vision and Pattern Recognition (CVPR)*, 2017.
- [9] Abhishek Das, Satwik Kottur, Khushi Gupta, Avi Singh, Deshraj Yadav, José M. F. Moura, Devi Parikh, and Dhruv Batra, “Visual dialog,” *CoRR*, vol. abs/1611.08669, 2016.
- [10] Abhishek Das, Satwik Kottur, José M.F. Moura, Stefan Lee, and Dhruv Batra, “Learning cooperative visual dialog agents with deep reinforcement learning,” in *International Conference on Computer Vision (ICCV)*, 2017.
- [11] Chiori Hori, Takaaki Hori, Teng-Yok Lee, Ziming Zhang, Bret Harsham, John R. Hershey, Tim K. Marks, and Kazuhiko Sumi, “Attention-based multimodal fusion for video description,” in *The IEEE International Conference on Computer Vision (ICCV)*, Oct 2017.
- [12] Huda Alamri, Vincent Cartillier, Raphael Gontijo Lopes, Abhishek Das, Jue Wang, Irfan Essa, Dhruv Batra, Devi Parikh, Anoop Cherian, Tim K Marks, and Chiori Hori, “Audio visual scene-aware dialog (avsd) challenge at dstc7,” *arXiv preprint arXiv:1806.00525*, 2018.

- [13] Gunnar A. Sigurdsson, Gül Varol, Xiaolong Wang, Ivan Laptev, Ali Farhadi, and Abhinav Gupta, “Hollywood in homes: Crowdsourcing data collection for activity understanding,” *ArXiv*, 2016.
- [14] Haonan Yu, Jiang Wang, Zhiheng Huang, Yi Yang, and Wei Xu, “Video paragraph captioning using hierarchical recurrent neural networks,” *CoRR*, vol. abs/1510.07712, 2015.
- [15] Sergio Guadarrama, Niveda Krishnamoorthy, Girish Malkarnenkar, Subhashini Venugopalan, Raymond Mooney, Trevor Darrell, and Kate Saenko, “Youtube2text: Recognizing and describing arbitrary activities using semantic hierarchies and zero-shot recognition,” in *Proceedings of the IEEE International Conference on Computer Vision*, 2013, pp. 2712–2719.
- [16] Jun Xu, Tao Mei, Ting Yao, and Yong Rui, “Msr-vtt: A large video description dataset for bridging video and language,” in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2016.
- [17] Du Tran, Lubomir D. Bourdev, Rob Fergus, Lorenzo Torresani, and Manohar Paluri, “Learning spatiotemporal features with 3d convolutional networks,” in *2015 IEEE International Conference on Computer Vision, ICCV 2015, Santiago, Chile, December 7-13, 2015*, 2015, pp. 4489–4497.
- [18] Joao Carreira and Andrew Zisserman, “Quo vadis, action recognition? a new model and the kinetics dataset,” in *CVPR*, 2017.
- [19] Chiori Hori, Takaaki Hori, Teng-Yok Lee, Ziming Zhang, Bret Harsham, John R Hershey, Tim K Marks, and Kazuhiko Sumi, “Attention-based multimodal fusion for video description,” in *ICCV*, 2017.
- [20] Yusuf Aytar, Carl Vondrick, and Antonio Torralba, “Soundnet: Learning sound representations from unlabeled video,” in *NIPS*, 2016.
- [21] Chiori Hori, Takaaki Hori, Tim K Marks, and John R Hershey, “Early and late integration of audio features for automatic video description,” in *ASRU*, 2017.
- [22] S. Hershey, S. Chaudhuri, D. P. W. Ellis, J. F. Gemmeke, A. Jansen, R. C. Moore, M. Plakal, D. Platt, R. A. Saurous, B. Seybold, M. Slaney, R. J. Weiss, and K. Wilson, “CNN architectures for large-scale audio classification,” in *ICASSP*, 2017.
- [23] Andrej Karpathy, George Toderici, Sanketh Shetty, Thomas Leung, Rahul Sukthankar, and Li Fei-Fei, “Large-scale video classification with convolutional neural networks,” in *Proceedings of the IEEE conference on Computer Vision and Pattern Recognition*, 2014, pp. 1725–1732.
- [24] Will Kay, Joao Carreira, Karen Simonyan, Brian Zhang, Chloe Hillier, Sudheendra Vijayanarasimhan, Fabio Viola, Tim Green, Trevor Back, Paul Natsev, et al., “The kinetics human action video dataset,” *arXiv*, 2017.
- [25] Diederik Kingma and Jimmy Ba, “Adam: A method for stochastic optimization,” *arXiv preprint arXiv:1412.6980*, 2014.