

Audio Visual Scene-Aware Dialog AVSD Challenge at DSTC7

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

TR2018-069 July 12, 2018

Abstract

Scene-aware dialog systems will be able to have conversations with users about the objects and events around them. Progress on such systems can be made by integrating state-of-the-art technologies from multiple research areas including end-to-end dialog systems visual dialog, and video description. We introduce the Audio Visual SceneAware Dialog (AVSD) challenge and dataset. In this challenge, which is one track of the 7th Dialog System Technology Challenges (DSTC7) workshop1 , the task is to build a system that generates responses in a dialog about an input video.

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.

Audio Visual Scene-Aware Dialog (AVSD) Challenge at DSTC7

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

^{*}School of Interactive Computing, Georgia Tech

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

Abstract

Scene-aware dialog systems will be able to have conversations with users about the objects and events around them. Progress on such systems can be made by integrating state-of-the-art technologies from multiple research areas including end-to-end dialog systems visual dialog, and video description. We introduce the Audio Visual Scene-Aware Dialog (AVSD) challenge and dataset. In this challenge, which is one track of the 7th Dialog System Technology Challenges (DSTC7) workshop¹, the task is to build a system that generates responses in a dialog about an input video.

1. Introduction

Spoken dialog technologies are becoming more common in real-world human-machine interfaces. Recently, end-to-end training of neural networks has been shown to be a promising approach for training dialogue systems from human-to-human dialogue corpora [2, 7]. A variety of neural conversation models were tested at DSTC6 [5]. However, current dialog systems are unable to have a conversation about what is going on in the user’s surroundings. The AVSD Challenge is motivated by the need for scene-aware dialog technology, so that machines can carry on a conversation with users about objects and events around them.

Entries to this challenge could capitalize on recent developments in related areas such as video description and Visual Dialog. Encoder-decoder networks developed for image captioning have recently been extended to the task of automatic *video description* [9], which is the generation of natural language descriptions of videos (e.g., a sentence that summarizes an input video). To enhance video description performance, [6] introduced an attention-based multimodal fusion approach that selectively attends to different input

modalities such as audio and video features. Visual Dialog [2, 3, 4] extends visual question answering (VQA) [1] from simple single-turn question answering to multi-turn dialog, in which the utterances in each turn may reference information from previous turns of the dialog.

In the AVSD challenge, we further extend Visual Dialog by extending the subject of the interaction from unimodal static images to multimodal videos, where input features could come from multiple domains including image features, motion features, non-speech audio, and speech audio.

2. Audio Visual Scene-Aware Dialog Challenge

2.1. Tasks

In this challenge, the system must generate responses to a user input in the context of a given dialog. This context consists of a dialog history (previous utterances by both user and system) in addition to video and audio information that comprise the scene. The quality of a system’s automatically generated sentences is evaluated using objective measures to determine whether or not the generated responses are natural and informative.

There are two tasks, each with two versions (a and b):

Task 1: Video and Text (a) Use the video and text training data provided but no external data sources, other than publicly available pre-trained feature extraction models. (b) External data may also be used for training.

Task 2: Text Only (a) Do not use the input videos for training or testing. Use only the text training data (dialogs and video descriptions) provided. (b) Any publicly available text data may also be used for training.

Challenge participants can select to submit entries in any or all of Task 1(a,b) and Task 2(a,b). Training data and a baseline system will be released to all participants of DSTC7.

The quality of the automatically generated sentences will be evaluated with objective measures to measure the similarity between the generated sentences and ground truth

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

sentences. We will use `nlg-eval`² for objective evaluation of system outputs.

2.2. Data collection

We are collecting text-based human dialog data for videos from human action recognition datasets such as CHARADES³ and Kinetics⁴. We have already collected text-based dialog data about short videos from CHARADES [8], which contains untrimmed and multi-action videos, along with video descriptions.

The data collection paradigm for dialogs was similar to that described in [2], in which for each image, two different Amazon Mechanical Turk (AMT) workers chatted via a text interface to yield a dialog. In [2], each dialog consisted of a sequence of questions and answers about an image. In our dataset, two 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 shown the video but was only shown three static images: the first, middle and last frames of the video. Having seen static frames from the video, the questioner already has good information about image- and appearance-based information in the video. Thus, rather than focusing on scene information that is available in the static images, the dialog instead revolves around the events and other temporal features in the video, which is the content of interest for our AVSD dataset. After 10 rounds of Q/A about the events that happened in the video, the questioner (who has not seen the video) is required to write a video description summarizing the events in the video.

In total, we have collected dialogs for 7043 videos from the CHARADES training set plus 1465 videos from the validation set. See Table 1 for statistics.

Table 1. Audio Visual Scene-Aware Dialog Dataset on CHARADES. Since we did not have scripts for the test set, we split the validation set into 732 and 733 videos and use them as our validation and test sets, respectively.

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

3. Summary

We introduce a new challenge task and dataset—Audio Visual Scene-Aware Dialog (AVSD)—that form the basis of one track of the 7th Dialog System Technology Challenges (DSTC7) workshop. We collected human dialog data for

²<https://github.com/Maluuba/nlg-eval>

³<http://allenai.org/plato/charades/>

⁴<https://deepmind.com/research/open-source/open-source-datasets/kinetics/>

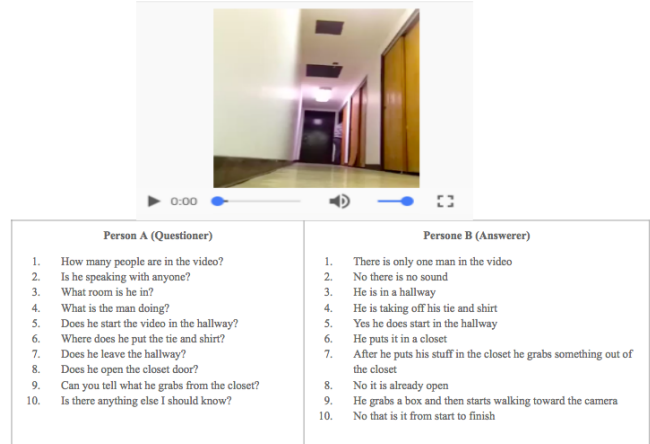


Figure 1. A sample from our Audio Visual Scene-Aware Dialog (AVSD) dataset. The task of Scene-aware Dialog requires an agent to generate a meaningful response about a video in the context of the dialog.

videos from the CHARADES dataset and plan to collect more for videos from the Kinetics dataset. The information provided to participants will include a detailed description of the baseline system, instructions for submitting results for evaluation, and details of the evaluation scheme.

References

- [1] S. Antol, A. Agrawal, J. Lu, M. Mitchell, D. Batra, C. L. Zitnick, and D. Parikh. VQA: Visual Question Answering. In *ICCV*, 2015.
- [2] A. Das, S. Kottur, K. Gupta, A. Singh, D. Yadav, J. M. F. Moura, D. Parikh, and D. Batra. Visual dialog. *CoRR*, abs/1611.08669, 2016.
- [3] A. Das, S. Kottur, J. M. Moura, S. Lee, and D. Batra. Learning cooperative visual dialog agents with deep reinforcement learning. In *ICCV*, 2017.
- [4] H. De Vries, F. Strub, S. Chandar, O. Pietquin, H. Larochelle, and A. Courville. Guesswhat?! visual object discovery through multi-modal dialogue. In *Proc. of CVPR*, 2017.
- [5] C. Hori and T. Hori. End-to-end conversation modeling track in DSTC6. volume abs/1706.07440, 2017.
- [6] C. Hori, T. Hori, T.-Y. Lee, Z. Zhang, B. Harsham, J. R. Hershey, T. K. Marks, and K. Sumi. Attention-based multimodal fusion for video description. In *ICCV*, Oct 2017.
- [7] R. Lowe, N. Pow, I. Serban, and J. Pineau. The ubuntu dialogue corpus: A large dataset for research in unstructured multi-turn dialogue systems. *arXiv preprint arXiv:1506.08909*, 2015.
- [8] G. A. Sigurdsson, G. Varol, X. Wang, I. Laptev, A. Farhadi, and A. Gupta. Hollywood in homes: Crowdsourcing data collection for activity understanding. 2016.
- [9] S. Venugopalan, H. Xu, J. Donahue, M. Rohrbach, R. J. Mooney, and K. Saenko. Translating videos to natural language using deep recurrent neural networks. In *NAACL HLT*, 2015.