

Context, Attention and Audio Feature Explorations for Audio Visual Scene-Aware Dialog

Shachi H Kumar Eda Okur Saurav Sahay

Juan Jose Alvarado Leanos Jonathan Huang Lama Nachman

Anticipatory Computing Lab
Intel Labs

{shachi.h.kumar, eda.okur, saurav.sahay, juan.jose.alvarado.leanos, jonathan.huang, lama.nachman}@intel.com

Abstract

With the recent advancements in AI, Intelligent Virtual Assistants (IVA) have become a ubiquitous part of every home. Going forward, we are witnessing a confluence of vision, speech and dialog system technologies that are enabling the IVAs to learn audio-visual groundings of utterances and have conversations with users about the objects, activities and events surrounding them. As a part of the 7th Dialog System Technology Challenges (DSTC7), for Audio Visual Scene-Aware Dialog (AVSD) track, We explore ‘topics’ of the dialog as an important contextual feature into the architecture along with explorations around multimodal Attention. We also incorporate an end-to-end audio classification ConvNet, AcNet, into our models. We present detailed analysis of the experiments and show that some of our model variations outperform the baseline system presented for this task.

Introduction

Training end-to-end dialog systems have obviated the need for building extensive engineered pipelines for various components of a Spoken Dialog System (SDS) (Collobert et al. 2011) (Serban et al. 2016). Recent progress in visual grounding techniques (Yu et al. 2016) (Antol et al. 2015) is enabling machines to understand semantic concepts that are shared across audio, language and vision modalities. Similar progress have been made in the domain of Audio Understanding (Hershey et al.) allowing machines to listen to the environment and understand various sensory events in the environment. With audio and visual grounding methods, end-to-end multimodal SDSes are now being trained to meaningfully communicate with us in natural language about the real dynamic audio-visual sensory world around us. Real world visual language interactions with machines are still in early stages of demonstrations and have been made possible due to Deep Learning based advancements in fields such as Machine Reading Comprehension, Neural Machine Translation (NMT), Visual Question Answering (VQA) and Visual Dialog (Das et al. 2017).

Topic models allow unsupervised learning of latent distributions of vector spaces for text documents. These models can help the SDS to consider additional information coming from the context of question topics and dialog history

topics while decoding the answers. In this work, we incorporate context information of the dialog - question, history of questions and answers, and captions in terms of adding underlying topical information expressed in these. We also build on top of the AVSD baseline model (Hori et al. 2018) by adding different configurations of attention layer to the dialog history LSTMs and audio/visual features to learn to attend to them while generating answers. The main difference of this approach from the baseline encoder-decoder is that it does not attempt to encode the dialog history purely as a recurrence layer. Instead, it encodes the history into a sequence of learnt vectors and chooses a subset of these vectors adaptively while decoding the answer. This allows the decoder to adaptively cope better with long interactions and not rely purely on the history sentence LSTM to try to remember long length interactions. We develop and test our approaches on the AVSD dataset (AlAmri et al. 2018) released as a part of the DSTC7 task.

Related Work

The encoder-decoder paradigm (Sutskever, Vinyals, and Le 2014) is the most commonly used paradigm for AVSD like applications and was initially developed for the Neural Machine Translation(NMT) domain. While the encoder maps the source encodings to fixed length vectors, the decoder translates the fixed length vectors to the desired target encodings. Many network variations have been proposed as part of this encoder decoder architecture (Cho et al. 2014). Attention-based models (Bahdanau, Cho, and Bengio 2014) can dynamically retrieve relevant pieces of the source via selective reading through a relatively simple matching operation. Recent architectures have proposed use of Transformers (Vaswani et al. 2017) that extend use of Attention mechanism in various interesting ways. Early Deep Learning based approaches in visual grounding (visual-to-text translation) used RNNs for image captioning (Vinyals et al. 2017). Early Video Captioning work simply extended the image captioning algorithms by average pooling the video frames (Venugopalan et al. 2015b). This simple extension fails to capture multiple events spanning longer duration videos. People developed more sophisticated ways to capture longer duration videos by using recurrent encoders for frames (Venugopalan et al. 2015a) or use of Attention (Yao et al. 2015) mechanism to selectively focus on relevant video features. Visual Ques-

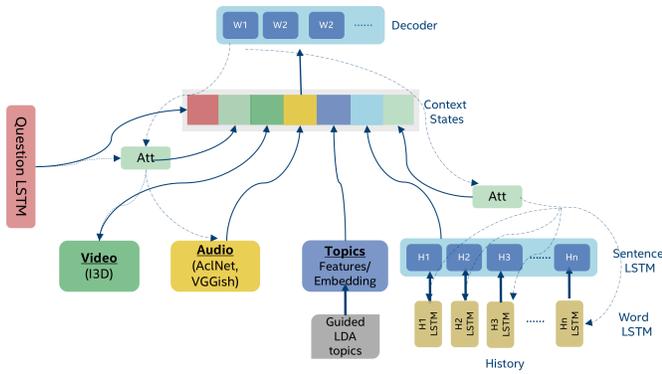


Figure 1: Architecture of Our System

tion Answering (VQA) (Antol et al. 2015) further extended the Video Captioning and Neural Machine Translation work and incorporated question-to-answer generation via end-to-end encoder-decoder framework. Visual Dialog task (Das et al. 2017) further extended VQA for multi-turn conversational dialog about static images addressing dialog handling challenges. AVSD task brings together all of these technologies for holding conversations about videos. To capture objects, events and other temporal information from audio and video streams, a lot of recent progress has been made in audio and video understanding area (Hori et al. 2017). AVSD system exploits these high level features learnt from various pre-trained networks such as VGGish and C3D (Yu et al. 2016).

Model Description

In this section, we describe the main architecture explorations of our work. Figure 1 shows our model with three main architectural extensions as will be discussed in the following sections.

Topic Model Explorations

Topics form a very important source of context in a dialog. A dialog is more likely to contain information about a few specific topics. Charades dataset (Sigurdsson et al. 2016) contains videos on common household activities such as watching TV, eating, cleaning, using a laptop, sleeping and so on. Since the dialog data involves the thread of conversation about a given video, there is a natural ‘topic’ of conversation based on the activity in the video. For example, if the video is about a person in the kitchen cooking, the dialog most likely revolves around the things or activities in the kitchen. Similarly, a video about a person watching TV would most likely involve a dialog around the activities in the living room. However, this segregation is not as straight-forward because only one person in the dialog has access to the entire video, and the other person is guessing and questioning about the activity. This demands a computational linguistics approach to understand the topics of the conversation. In this work, we use topic models to map parts of the dialog to topics in an unsupervised way.

Latent Dirichlet Allocation (LDA). Latent Dirichlet Allocation (Blei, Ng, and Jordan 2001) is one of the very commonly used topic models to determine the topic distribution of a document. We train topic models on questions, answers, QA pairs, captions and history. The topic distributions generated by these models are incorporated as features into our models or used to learn topic embeddings.

Guided LDA¹. Since we are interested in identifying specific activities based on domain knowledge, we use Guided LDA (Jagarlamudi, III, and Udupa 2012) that presents a way to guide topic models to learn topics of specific interest to a user. This technique uses a set of seed words provided by the user, representative of the topics in a corpus.

- **Seed words:** We derive the seed words for various topics based on the analysis presented in (Sigurdsson et al. 2016) for the Charades dataset. We consider topics such as entertainment, cooking, cleaning, resting, etc. A detailed list of seed words for the 9-topics configuration is presented in Table 1. We obtain the seed words by identifying a set of most common nouns (objects), verbs, scenes, and actions from the Charades dataset analysis.
- **Examples of text with topics:** Below are some of the examples with the predicted topics (having the highest probability within document-topic distribution vector) for given captions, questions and answers from the AVSD dataset:

Caption: *a person smiles as they cook food on the stove . the person takes out their laptop and consults a recipe .*

→ Predicted Topic: *Cooking / Kitchen*

Question: *does he take something out of the fridge ?*

Answer: *yes he takes something out of the refrigerator*

→ Predicted Topic: *Cooking / Kitchen*

Caption: *a person runs into the bedroom , undressing . the person checks their phone , turns off the light , then jumps in bed .*

→ Predicted Topic: *Rest / Bedroom*

Caption: *a person is playing with their camera while sitting on the stairs . the person goes to take a picture of themselves , but sneezes just as the flash goes off .*

→ Predicted Topic: *Entertainment / Living Room*

Caption: *a person walks into the bathroom and turns on the light . the person drinks from a cup of coffee . the person watches themselves in the mirror and smiles .*

→ Predicted Topic: *Cleaning / Bath*

Attention Explorations

In the AVSD baseline architecture, the context vector is the main bridge between the multimodal encoder and the answer decoder. The dialog history encoder encodes information via word level LSTMs followed by sentence level LSTMs. The sentence level LSTM may not be able to retain information from earlier sentences in the dialog history. Only the last

¹<https://github.com/vi3k6i5/GuidedLDA>

Topic	Seed Words
Entertainment/LivingRoom	living, room, recreation, garage, basement, entryway, television, tv, phone, laptop, sofa, chair, couch, armchair, seat, picture, sit, play, smile, laugh, watch, listen, turn, ...
Cooking/Kitchen	kitchen, pantry, food, water, dish, sink, refrigerator, fridge, stove, microwave, toaster, kettle, oven, stewpot, saucepan, cook, wash, prepare, cut, chop, heat, bake, fry, ...
Eating/Dining	dining, room, table, chair, plate, fork, knife, spoon, bowl, glass, cup, mug, coffee, tea, sandwich, meal, breakfast, lunch, dinner, wine, bar, eat, drink, pour, grab, have, ...
Cleaning/Bath	bathroom, hallway, entryway, stairs, restroom, toilet, towel, broom, vacuum, floor, sink, water, mirror, cabinet, medicine, hairdryer, clean, wash, shower, brush, bath, ...
Dressing/Closet	walk-in, closet, clothes, wardrobe, shoes, shirt, pants, trousers, skirt, jacket, t-shirt, underwear, sweatshirt, coat, hanger, rack, dress, undress, wear, put, fit, hang, try, ...
Laundry	laundry, room, basement, clothes, clothing, cloth, basket, bag, box, towel, shelf, dryer, washer, washing, machine, do, wash, grasp, hold, get, throw, take, pick, dry, ...
Rest/Bedroom	bedroom, room, bed, pillow, blanket, mattress, bedstand, nightstand, commode, dresser, bedside, lamp, nightlight, night, light, lie, snuggle, sleep, rest, tidy, awaken, ...
Work/Study	home, office, den, workroom, room, garage, basement, laptop, desktop, computer, pc, monitor, mouse, keyboard, phone, desk, chair, light, work, study, write, type, ...
Sports/Exercise	recreation, room, garage, basement, hallway, stairs, gym, fitness, floor, bag, towel, ball, treadmill, bike, rope, mat, run, walk, workout, exercise, cycle, lift, jump, jog, ...

Table 1: Sample of Seed Words for 9 Topics

hidden state of sentence level LSTM is passed on to the context vector for the decoder to leverage information from dialog history. This single n-dimensional vector carries the burden of retaining useful information from the entire dialog history. We add attention layer to the dialog history representations and audio/video features to selectively focus on relevant parts of the dialog history and audio/video features at every step of the decoder. We calculate the attention weights corresponding to every dialog history turn, the multimodal features and the decoder representation and apply the weights to the history and multimodal features to compute the relevant representations. These help create a weighted combination of the dialog history and multimodal context that is richer than the unweighted single context vectors of the encoder hidden representations. We append the input encoding along with the AV multimodal feature encodings and pass it to the decoder LSTM for learning the output encodings.

Audio Feature Explorations

Following the successes of image classification, ConvNets have become popular for audio classification. (Hershey et al.) showed that large image classification ConvNets trained with huge amount of weakly labeled Youtube data leads to semantically meaningful representations, which has been effectively used for transfer learning in new tasks. For the sound recognition pipeline, we used an end-to-end audio classification ConvNet, called AclNet (Huang and Leanos 2018). AclNet takes raw, amplitude-normalized 44.1 kHz audio samples as input, and produces classification output without the need to compute spectral features. The network consists of two separate building blocks: the low-level features which uses two 1D convolutions that produces spectral-like outputs, followed by the high-level features which is inspired by the VGG architecture.

AclNet is trained using the ESC-50 (Piczak 2015) corpus, which is a dataset of 50 classes of environmental sounds organized in 5 semantic categories (animals, interior/domestic, exterior/urban, human, natural landscapes). Each sound file lasts for 5 seconds and there are 2000 audio recordings in total. ESC-50 has been widely adopted by the machine audition community as a testbed for novel architectures. At the time of this writing, AclNet shows the best one shot accuracy of 85.65% (behind an ensemble model) for the ESC-50.

The nature of training data and the classification architecture have several differences from the VGGish model, which is why we believe it provides a diversifying if not complementary effect for the model fusion. The files and labels of the ESC-50 are carefully curated, as opposed the weakly labeled Youtube data that VGGish is trained with. The spectral features learned by AclNet was completely data-driven, thus it may behave somewhat differently than the mel-spectral features from VGGish.

Dataset

We use the official-training, prototype-validation and prototype-test sets provided as a part of the DSTC7 Audio Visual Scene-Aware Dialog challenge. Table 2 shows the distribution of AVSD data across these different sets.

	Training	Validation	Test
# of Dialogs	7,659	732	733
# of Turns	153,180	14,680	14,660
# of Words	1,450,754	138,314	138,790

Table 2: Audio Visual Scene-Aware Dialog Dataset

AVSD Dataset Analysis

We analyze the official training set of AVSD dataset to gain a better insight on the dialogs, and we also use that knowledge to create certain subsets of the prototype test set to evaluate our results on these subsets in order to observe the desired effects of our models.

- **Q/A Lengths:** We analyzed the number of words in questions and answers in AVSD, and here are our findings:

$$\text{Questions: } \mu = 8.46, \sigma = 3.48, \text{range} = [2, 56]$$

$$\text{Answers: } \mu = 10.41, \sigma = 5.03, \text{range} = [5, 55]$$

Unlike similar previous datasets like VisDial (Das et al. 2017) and VQA (Antol et al. 2015), answers in AVSD are significantly longer and much more detailed. We observed a mean-length of 10.4 words in AVSD compared to 2.9 in Vis-Dial and 1.1 in VQA datasets. This indicates the existence of descriptive and conversational answers in AVSD. Similarly, questions are also lengthy as we see a mean-length of 8.5 words in AVSD compared to 5.1 in Vis-Dial and 6.2 in VQA datasets.

- **Binary vs. Non-Binary Q/A:** In VisDial, binary questions are defined as those starting with ‘do’, ‘did’, ‘have’, ‘has’, ‘is’, ‘are’, ‘was’, ‘were’, ‘can’, ‘could’. We expanded that keywords set with variations based on the most common 100 first words in AVSD questions, and we found that 61.19% of the questions are binary in AVSD. Similarly, answers that contain ‘yes’ or ‘no’ are assumed to be binary answers in VisDial, and we expanded the keywords with common variations which strongly indicate that the answer is binary. As a results, we observed that 43.07% of the answers in AVSD are binary. Among them, only 41.69% are ‘yes’ within all binary (yes/no) answers. We split the dataset into two for binary and non-binary questions, and we did the same for the answers.
- **Coreferences in Dialogs:** We analyzed the presence of coreferences in AVSD dialogs by checking whether they contain pronouns like ‘he’, ‘she’, ‘his’, ‘her’, ‘it’, ‘their’, ‘they’, ‘this’, ‘that’, ‘those’, etc. As the AVSD dialogs have 10 turns, the coreferences are quite common, especially towards the later turns as expected. We observed that 62.08% of the questions and 73.69% of the answers contain those pronouns, which indicates that the coreference resolution will be an issue and the dialog history should be taken into account more seriously.
- **Audio-related Q/A:** Similar to the binary vs. non-binary Q/A, we constructed a list of keywords which strongly indicate that the conversation involves audio-related questions and answers (such as ‘audio’, ‘sound’, ‘hear’, ‘noise’, ‘voice’, etc.). We found that 12.38% of the questions and 14.39% of the answers are audio-related in AVSD, and again we created the subsets accordingly.

Experiments and Results

We use the official-training and prototype-validation sets provided for the task for our experiments and report the results on the prototype-test set. Note that we cannot use the official-test set since the ground truth is not released yet; and we cannot use the official-validation set since it covers the dialogs in the prototype-test set, on which we are reporting our results in this paper. For all of the experiments, we use Adam optimizer and pick the best model based on the lowest validation perplexity for response generation.

In the subsections below, we present the results of our explorations on topics, attention and audio-based experiments. In general, the CIDEr metric (Consensus-based Image Description Evaluation) (Vedantam, Zitnick, and Parikh 2015) uses sentence similarity and inherently captures the notions of grammaticality, saliency, importance and accuracy. It is also shown to present higher agreement with the judgement of consensus assessed by humans compared to other metrics such as BLEU and ROUGE. Based on this, we focus more on the improvements in CIDEr scores for our experiments.

Topic Experiments

To include topics, we try various configurations. We use separate topic models trained on questions, answers, QA pairs (i.e., concatenated Q and A), captions, history (i.e., concatenated QA pairs until the current dialog turn), and his-

tory+captions (i.e., concatenated caption and QA pairs until the current turn) respectively to generate topics for samples from each category.

- **Standard LDA vs. Guided LDA:** We experimented with training topic models using both Standard LDA (without seed words) and Guided LDA (with seed words) methods separately for questions, answers, QA pairs, captions, and overall history.
- **Number of Topics:** To investigate the effects of variation in number of guided topics on the performance of our models, we experimented with training topic models using 5, 7, 9 and 11 topics. We observed that using 9 seeded topics yields the best scores for the responses. Thus, we continued with 9 topics for all of our topic experiments.

In addition, we experimented with adding topics in various ways as described below:

- 1 **Topics as Features:** We explore incorporating the topic distribution vectors generated by Guided LDA as features for the questions and the dialog history. The question topics are added to the decoder state directly. In one variation, the dialog history topics are copied to all the decoder states directly. For this variation, we explored two configurations to construct dialog history topics: 1) We concatenate QA pair topics and caption topics at each turn, which we directly call our ‘GuidedLDA’ model; 2) We leverage all available topics (i.e., questions, answers, QA pairs, captions, history, history+captions), which we call ‘GuidedLDA-all’. In another variation, the dialog history topics are added as features to the history encoder LSTM to generate a richer representation, which we call HLSTM-with-topics. We also experiment adding GloVe vector representation (Pennington, Socher, and Manning 2014) for question and history. Specifically, we tried 200 dimension vectors with fine-tuning.
- 2 **Learning Topic Embeddings:** We learn topic embeddings from topics generated for the questions, QA pairs and captions. For each question, QA pair and caption, we pick the top-3 topics based on the topic proportions generated by the Guided LDA. We try to learn the embeddings for these topics similar to learning word-embeddings.

	Bleu1	Bleu2	Bleu3	Bleu4	Meteor	Rouge	CIDEr
Baseline	0.273	0.173	0.118	0.084	0.117	0.291	0.766
StandardLDA	0.255	0.164	0.113	0.082	0.114	0.285	0.772
GuidedLDA	0.265	0.170	0.117	0.084	0.118	0.293	0.812
GuidedLDA-all	0.272	0.173	0.118	0.085	0.119	0.293	0.793
GuidedLDA+GloVe	0.275	0.175	0.119	0.085	0.121	0.293	0.797
Topic-embeddings	0.257	0.165	0.114	0.083	0.115	0.287	0.772
HLSTM-with-topics	0.260	0.166	0.115	0.084	0.117	0.290	0.797

Table 3: Topic Experiments

Table 3 compares the baseline model with the addition of StandardLDA and GuidedLDA topic distributions as features. GuidedLDA performs better than the baseline for the CIDEr, ROUGE and METEOR metrics, while the baseline performs better for the BLEU scores (except BLEU4, which is the same for both). StandardLDA shows better CIDEr

	Bleu1	Bleu2	Bleu3	Bleu4	Meteor	Rouge	CIDEr
Binary							
Baseline	0.329	0.220	0.156	0.116	0.142	0.345	0.965
GuidedLDA	0.318	0.216	0.155	0.116	0.144	0.346	1.008
GuidedLDA+GloVe	0.328	0.220	0.156	0.115	0.146	0.347	0.995
Non-binary							
Baseline	0.233	0.139	0.090	0.060	0.099	0.250	0.541
GuidedLDA	0.226	0.136	0.089	0.061	0.100	0.253	0.590
GuidedLDA+GloVe	0.237	0.141	0.092	0.063	0.103	0.252	0.579

Table 4: Performance on Binary/Non-binary Answers

score compared to the baseline, but GuidedLDA performs better than StandardLDA for all metrics as expected. Note that when we incorporate all available topics at each turn using GuidedLDA-all model instead of GuidedLDA, we obtain much better BLEU scores; and we reach better or similar performances for all metrics compared to the baseline. GuidedLDA+GloVe is our best performing model and outperforms baseline in all metrics. It also outperforms all other topic-based models that we explored (except GuidedLDA in terms of CIDEr score). The HLSTM-with-topics model performs better and similar to the baseline in terms of CIDEr metric and BLEU4, METEOR and ROUGE metrics, respectively. The baseline however performs better than the Topic-embeddings model in all metrics except CIDEr.

We also compare the GuidedLDA models with the baseline by evaluating on subsets of the dataset consisting of binary and non-binary answers. From Table 4, for binary subset, GuidedLDA performs better than baseline for some metrics while GuidedLDA+GloVe shows similar/better performance as compared to baseline for all metrics. Interestingly, for the non-binary subset, our GuidedLDA+GloVe model outperforms baseline for all metrics. This shows that our model can generate better responses for the more complex, non-binary answer types.

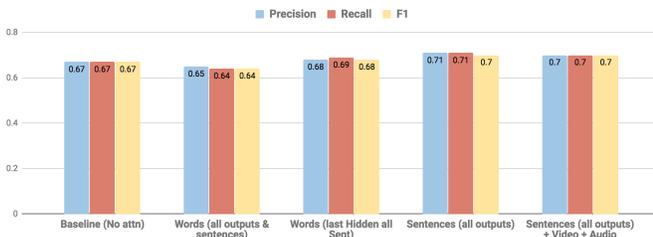


Figure 2: Precision, Recall, F1 Scores for Attention Experiments on Questions Containing Coreferences

Attention Experiments

In our model, we have modified the baseline architecture and added attention layer with 4 different configurations for the answer decoder to leverage information directly from the dialog history LSTMs and multimodal audio/video features. Since the attention layer pays attention to the dialog history in all our configurations while decoding, we wanted to evaluate the performance of attention with respect to the base-

line solely for questions that could benefit from dialog history. It is non-trivial to do such analysis quantitatively. One naive way to accomplish this is to isolate the questions containing coreferences and perform the analysis on these questions. Table 6 shows the performance of our models on the standard evaluation metrics on this coreference-subset of the dataset. We further wanted to check how semantically meaningful were the results that we generated using the attention variations. In order to compare the results at a more semantic level, we performed quantitative analysis on dialogs that contained binary answer in the ground truth, by creating the binary answers subset. We evaluate our models on their ability to predict these binary answers correctly (using precision, recall and F1 scores) and present this analysis in Figure 2.

1 *Attention on Dialog History Word LSTMs, all output states*: In this configuration, we remove the sentence level dialog history LSTM and the decoder computes the attention scores directly between the decoder state and the word level output states for all dialog history. We first padded the Word LSTM outputs from Dialog History LSTMs (see Word LSTM in History in Figure 1) to the maximum sentence length of all the sentences. We summed up all the attention scores from each of the sentence context vectors with the query decoder state. Directly attending to the output states of the word LSTMs in the dialog history encoder didn't perform well on any of the evaluation metrics compared to the baseline. This attention mechanism possibly attended to way more information than needed. Using this kind of attention, we had hoped that the system could remember answers that were already given (directly or indirectly) in the earlier turns of the dialog. Unfortunately we couldn't perform such quantitative analysis on the dataset in this work. The binary answer evaluation in this setup also performed poorer than the baseline as shown in Figure 2.

2 *Attention on Dialog History Word LSTMs, last hidden states*: This configuration is similar to the previous configuration with the difference that we only use the last hidden state output representations of the word LSTMs corresponding to the different turns in the dialog. Simpler than the previous setup, we stack up the hidden states from the history sentences for attention computation. From table 6, this model performs better than the previous setup, slightly worse than the baseline on standard evaluation metrics and slightly better than the baseline on binary answer evaluation from figure 2.

3 *Attention on Sentence LSTM output states*: The baseline architecture only leverages the last hidden state information from the sentence LSTM in the dialog history encoder. Instead, we extract the output states from all timesteps of the LSTM corresponding to n turns of the dialog history. This variation helps the decoder consider all the dialog turn compressed sentence representations via the attention mechanism. As shown in Table 6, this model performed better than the baseline model on BLEU3, BLEU4, METEOR, CIDEr and ROUGE. Figure 2 shows that this model also performed better than the baseline model on binary answer evaluation.

Question:	<i>is there sound to the video ?</i>	<i>Did he make a sound when blowing his nose ?</i>	<i>does she say anything ?</i>
Ground Truth	yes there is sound . nothing important .	he didnt blow his nose	she does not say anything .
Baseline + VGGish	no there is no sound in the video	no he did not smile in the video	no there is no sound in the video .
Baseline + AclNet	yes there is sound in the video	no he didnt sneeze in the video	no she does not say anything .

Table 5: Audio Examples (VGGish vs. AclNet)

	Bleu1	Bleu2	Bleu3	Bleu4	Meteor	Rouge	CIDEr
Baseline	0.248	0.151	0.101	0.071	0.110	0.256	0.664
Word LSTM (all output states)	0.223	0.138	0.092	0.065	0.103	0.262	0.591
Word LSTM (last hidden states)	0.229	0.139	0.093	0.065	0.105	0.250	0.661
Sentence LSTM (all output states)	0.242	0.151	0.103	0.073	0.110	0.261	0.707
Sentence LSTM (all outputs) +video+audio	0.234	0.146	0.099	0.070	0.109	0.254	0.690

Table 6: Decoder Attention over Dialog History and Multimodal Features

	Bleu1	Bleu2	Bleu3	Bleu4	Meteor	Rouge	CIDEr
Overall							
Baseline(B)	0.273	0.173	0.118	0.084	0.117	0.291	0.766
B+VGGish	0.271	0.172	0.118	0.085	0.116	0.292	0.791
B+AclNet	0.274	0.175	0.121	0.087	0.117	0.294	0.789
Audio-related							
Baseline(B)	0.267	0.179	0.128	0.096	0.120	0.285	0.919
B+VGGish	0.266	0.181	0.131	0.099	0.118	0.285	0.907
B+AclNet	0.266	0.183	0.132	0.100	0.120	0.287	0.944

Table 7: Audio Experiments & Performance on Overall vs. Audio-related Questions

4 *Attention on Sentence LSTM output states and Multimodal Audio/Video Features*: This configuration is similar to the last one with the difference that we add multimodal audio/video features as additional state to the attention module. This mechanism allows the decoder to selectively focus on the multimodal features along with the dialog history sentences. This configuration didn't really help improve the evaluation metrics compared to the baseline (see Table 6) except the CIDEr metric score. The results on binary answer evaluation (Figure 2) improved compared to the baseline but slightly degraded compared to the last configuration.

Audio Experiments

Features extracted using a new state-of-the-art model, Audio Set VGGish were provided as a part of the DSTC7 AVSD challenge. We explore the use of the AclNet features based on the softmax output of the 50-classes from the model described in the previous section. Table 7 shows the comparison of the baseline model without audio features, and baseline with the addition of the VGGish and the AclNet features. From the table, Baseline+AclNet (B+AclNet) shows improved scores on all metrics as compared to the baseline and Baseline+VGGish (although B+VGGish gives the best score for this metric).

We further try to analyse the effects of the audio features specifically on audio-related questions. We split the test dataset into audio-related and non-audio-related dialogs

and evaluate the performance of our models on these subsets. From Table 7, we observe that B+AclNet performs the best on audio-related questions as well.

Table 5 presents a qualitative analysis of the addition of the VGGish and AclNet features to the baseline model. The question, "*is there sound in the video?*" is incorrectly answered by the B+VGGish model and correctly answered by the AclNet addition. Another interesting observation is made for the question "*Did he make a sound when blowing his nose?*". B+AclNet refers to "*sneeze*" in the response which is close to the ground truth answer "*he didnt blow his nose*", while the VGGish model does not generate a relevant response.

Conclusion and Future Work

In this paper, we present some explorations and techniques to improve contextual and multimodal end-to-end audio-visual scene aware dialog system. We incorporate context of the dialog in the form of topics, we use various attention mechanisms to enable the decoder to focus on relevant parts of the dialog history and audio/video features, and we incorporate audio features from an end-to-end audio classification architecture, AclNet. As part of the 7th Dialog System Technology Challenges (DSTC7), we validate our approaches on the AVSD dataset and show that some of our models perform better than the baseline on various metrics. Our topic-based contextual models also show that guiding the topic models with seed-words help in improving the overall performance as compared to the baseline. We also present a quantitative evaluation of our models on their ability to predict binary answers correctly. We show that our attention-based mechanisms perform well on dialog data involving coreference. We also show qualitatively and quantitatively that incorporating our audio pipeline improves the performance as compared to the VGGish features. AVSD is a new and exiting research area and the current work involved addition of contextual features and Attention mechanisms. As future work, besides fine-tuning the current techniques, we plan to explore multimodal fusion techniques and incorporate better representations for other modalities, such as object, pose and activity recognition algorithms into the pipeline.

References

- AlAmri, H.; Cartillier, V.; Lopes, R. G.; Das, A.; Wang, J.; Essa, I.; Batra, D.; Parikh, D.; Cherian, A.; Marks, T. K.; and Hori, C. 2018. Audio visual scene-aware dialog (AVSD) challenge at DSTC7. *CoRR* abs/1806.00525.
- Antol, S.; Agrawal, A.; Lu, J.; Mitchell, M.; Batra, D.; Zitnick, C. L.; and Parikh, D. 2015. Vqa: Visual question answering. In *2015 IEEE International Conference on Computer Vision (ICCV)*, 2425–2433.
- Bahdanau, D.; Cho, K.; and Bengio, Y. 2014. Neural machine translation by jointly learning to align and translate. *CoRR* abs/1409.0473.
- Blei, D. M.; Ng, A. Y.; and Jordan, M. I. 2001. Latent dirichlet allocation. In Dietterich, T. G.; Becker, S.; and Ghahramani, Z., eds., *Advances in Neural Information Processing Systems 14 [Neural Information Processing Systems: Natural and Synthetic, NIPS 2001]*. MIT Press.
- Cho, K.; van Merriënboer, B.; Bahdanau, D.; and Bengio, Y. 2014. On the properties of neural machine translation: Encoder-decoder approaches. In Wu, D.; Carpuat, M.; Carreras, X.; and Vecchi, E. M., eds., *Proceedings of SSST@EMNLP 2014, Eighth Workshop on Syntax, Semantics and Structure in Statistical Translation*. Association for Computational Linguistics.
- Collobert, R.; Weston, J.; Bottou, L.; Karlen, M.; Kavukcuoglu, K.; and Kuksa, P. P. 2011. Natural language processing (almost) from scratch. *Journal of Machine Learning Research* 12:2493–2537.
- Das, A.; Kottur, S.; Gupta, K.; Singh, A.; Yadav, D.; Moura, J. M. F.; Parikh, D.; and Batra, D. 2017. Visual dialog. *2017 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*.
- Hershey, S.; Chaudhuri, S.; Ellis, D. P.; Gemmeke, J. F.; Jansen, A.; Moore, R. C.; Plakal, M.; Platt, D.; Saurous, R. A.; Seybold, B.; et al. Cnn architectures for large-scale audio classification. In *Acoustics, Speech and Signal Processing (ICASSP), 2017 IEEE International Conference*. IEEE.
- Hori, C.; Hori, T.; Lee, T.-Y.; Zhang, Z.; Harsham, B.; Hershey, J. R.; Marks, T. K.; and Sumi, K. 2017. Attention-based multimodal fusion for video description. *2017 IEEE International Conference on Computer Vision (ICCV)*.
- Hori, C.; AlAmri, H.; Wang, J.; Wichern, G.; Hori, T.; Cherian, A.; Marks, T. K.; Cartillier, V.; Lopes, R. G.; Das, A.; Essa, I.; Batra, D.; and Parikh, D. 2018. End-to-end audio visual scene-aware dialog using multimodal attention-based video features. *CoRR* abs/1806.08409.
- Huang, J. J., and Leanos, J. J. A. 2018. Aclnet: Efficient end-to-end audio classification cnn. *arXiv preprint arXiv:1811.06669*.
- Jagarlamudi, J.; III, H. D.; and Udupa, R. 2012. Incorporating lexical priors into topic models. In Daelemans, W.; Lapata, M.; and Màrquez, L., eds., *EACL 2012, 13th Conference of the European Chapter of the Association for Computational Linguistics*. The Association for Computer Linguistics.
- Pennington, J.; Socher, R.; and Manning, C. D. 2014. Glove: Global vectors for word representation. In *Empirical Methods in Natural Language Processing (EMNLP)*, 1532–1543.
- Piczak, K. J. 2015. ESC: Dataset for Environmental Sound Classification. In *Proceedings of the 23rd Annual ACM Conference on Multimedia*, 1015–1018. ACM Press.
- Serban, I. V.; Sordoni, A.; Bengio, Y.; Courville, A. C.; and Pineau, J. 2016. Building end-to-end dialogue systems using generative hierarchical neural network models. In *AAAI*, volume 16, 3776–3784.
- Sigurdsson, G. A.; Varol, G.; Wang, X.; Farhadi, A.; Laptev, I.; and Gupta, A. 2016. Hollywood in homes: Crowdsourcing data collection for activity understanding. In Leibe, B.; Matas, J.; Sebe, N.; and Welling, M., eds., *Computer Vision - ECCV 2016 - 14th European Conference, Proceedings*.
- Sutskever, I.; Vinyals, O.; and Le, Q. V. 2014. Sequence to sequence learning with neural networks. In *Advances in neural information processing systems*, 3104–3112.
- Vaswani, A.; Shazeer, N.; Parmar, N.; Uszkoreit, J.; Jones, L.; Gomez, A. N.; Kaiser, L.; and Polosukhin, I. 2017. Attention is all you need. In Guyon, I.; von Luxburg, U.; Bengio, S.; Wallach, H. M.; Fergus, R.; Vishwanathan, S. V. N.; and Garnett, R., eds., *Advances in Neural Information Processing Systems 30: Annual Conference on Neural Information Processing Systems 2017*.
- Vedantam, R.; Zitnick, C. L.; and Parikh, D. 2015. Cider: Consensus-based image description evaluation. In *IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2015, Boston, MA, USA, June 7-12, 2015*, 4566–4575. IEEE Computer Society.
- Venugopalan, S.; Rohrbach, M.; Donahue, J.; Mooney, R.; Darrell, T.; and Saenko, K. 2015a. Sequence to sequence – video to text. *2015 IEEE International Conference on Computer Vision (ICCV)*.
- Venugopalan, S.; Xu, H.; Donahue, J.; Rohrbach, M.; Mooney, R. J.; and Saenko, K. 2015b. Translating videos to natural language using deep recurrent neural networks. In Mihalcea, R.; Chai, J. Y.; and Sarkar, A., eds., *NAACL HLT 2015, The 2015 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*.
- Vinyals, O.; Toshev, A.; Bengio, S.; and Erhan, D. 2017. Show and tell: Lessons learned from the 2015 mscoco image captioning challenge. *IEEE Transactions on Pattern Analysis and Machine Intelligence* 39(4):652663.
- Yao, L.; Torabi, A.; Cho, K.; Ballas, N.; Pal, C.; Larochelle, H.; and Courville, A. 2015. Describing videos by exploiting temporal structure. In *2015 IEEE International Conference on Computer Vision (ICCV)*, 4507–4515.
- Yu, H.; Wang, J.; Huang, Z.; Yang, Y.; and Xu, W. 2016. Video paragraph captioning using hierarchical recurrent neural networks. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, 4584–4593.