

From FiLM to Video: Multi-turn Question Answering with Multi-modal Context

Dat Tien Nguyen^{1*}, Shikhar Sharma^{2†}, Hannes Schulz^{2†}, Layla El Asri²

¹ University of Amsterdam, Amsterdam, Netherlands ² Microsoft Research, Montréal, Canada
t.d.nguyen@uva.nl, {shikhar.sharma, hannes.schulz, layla.elasri}@microsoft.com

Abstract

Understanding audio-visual content and the ability to have an informative conversation about it have both been challenging areas for intelligent systems. The Audio Visual Scene-aware Dialog (AVSD) challenge, organized as a track of the Dialog System Technology Challenge 7 (DSTC7), proposes a combined task, where a system has to answer questions pertaining to a video given a dialogue with previous question-answer pairs and the video itself. We propose for this task a hierarchical encoder-decoder model which computes a multi-modal embedding of the dialogue context. It first embeds the dialogue history using two LSTMs. We extract video and audio frames at regular intervals and compute semantic features using pre-trained I3D and VGGish models, respectively. Before summarizing both modalities into fixed-length vectors using LSTMs, we use FiLM blocks to condition them on the embeddings of the current question, which allows us to reduce the dimensionality considerably. Finally, we use an LSTM decoder that we train with scheduled sampling and evaluate using beam search. Compared to the modality-fusing baseline model released by the AVSD challenge organizers, our model achieves a relative improvements of more than 16%, scoring 0.36 BLEU-4 and more than 33%, scoring 0.997 CIDEr.

1 Introduction

Deep neural networks have been successfully applied to several computer vision tasks such as image classification (Krizhevsky, Sutskever, and Hinton 2012), object detection (Szegedy, Toshev, and Erhan 2013), video action classification (Karpathy et al. 2014), etc. They have also been successfully applied to natural language processing tasks such as machine translation (Cho et al. 2014), machine reading comprehension (Hermann et al. 2015), etc. There has also been an explosion of interest in tasks which combine multiple modalities such as audio, vision, and language together. Some popular multi-modal tasks combining these three modalities, and their differences are highlighted in Table 1.

Given an image and a question related to the image, the Visual Question Answering (VQA) challenge (Antol et al. 2015) tasked users with selecting an answer to the question.

*Work performed during internship at Microsoft Research

†These two authors contributed equally

Agrawal, Batra, and Parikh (2016) identified several sources of bias in the VQA dataset, which led to deep neural models answering several questions superficially. They found that in several instances, deep architectures exploited the statistics of the dataset to select answers ignoring the provided image. This prompted the release of VQA 2.0 (Goyal et al. 2017) which attempts to balance the original dataset. In it, each question is paired to two similar images which have different answers. Due to the complexity of VQA, understanding the failures of deep neural architectures for this task has been a challenge. It is not easy to interpret whether the system failed in understanding the question or in understanding the image or in reasoning over it. The CLEVR dataset (Johnson et al. 2017) was hence proposed as a useful benchmark to evaluate such systems on the task of visual reasoning. Extending question answering over images to videos, Tapaswi et al. (2016) have proposed MovieQA, where the task is to select the correct answer to a provided question given the movie clip on which it is based.

Task	Visual	Audio	Text Format
VQA	Image	No	QA
MovieQA	Video	Yes	QA
VisDial	Image	No	QA-Dialogue
AVSD	Video	Yes	QA-Dialogue

Table 1: Tasks with audio, visual and text modalities

Intelligent systems that can interact with human users for a useful purpose are highly valuable. To this end, there has been a recent push towards moving from single-turn Question Answering (QA) to multi-turn dialogue, which is a natural and intuitive setting for humans. Among multi-modal dialogue tasks, Visual Dialog (VisDial) (Das et al. 2017a) provides an image and dialogue where each turn is a QA pair. The task is to train a model to answer these questions within the dialogue. The AVSD challenge extends the VisDial task from images to the audio-visual domain.

We present our FiLM Attention Hierarchical Recurrent Encoder-Decoder (FA-HRED) model for the AVSD task. FA-HRED combines a Hierarchical Recurrent Encoder-Decoder (HRED) for encoding and generating QA-dialogue with a novel FiLM-based audio-visual feature extractor for videos

and an auxiliary multi-task learning-based decoder for decoding a summary of the video. It outperforms the baseline results for the AVSD dataset (Hori et al. 2018) and was ranked 2nd overall among the DSTC7 AVSD challenge participants.

In Section 2, we discuss existing literature on end-to-end dialogue systems with a special focus on multi-modal dialogue systems. Section 3 describes the AVSD dataset. In Section 4, we present the architecture of our FA-HRED model. We describe our evaluation and experimental setup in Section 5 and then conclude in Section 6.

2 Related Work

With the availability of large conversational corpora from sources like Reddit and Twitter, there has been a lot of recent work on end-to-end modelling of dialogue for open domains. Ritter, Cherry, and Dolan (2011) treated dialogue as a machine translation problem where they translate from the stimulus to the response. They observed this to be more challenging than machine translation tasks due the larger diversity of possible responses. Among approaches that just use the previous utterance to generate the current response, Shang, Lu, and Li (2015) proposed a response generation model based on the encoder decoder framework. Sordani et al. (2015) also proposed an encoder-decoder based neural network architecture that uses the previous two utterances to generate the current response. Among discriminative methods (i.e. methods that produce a score for utterances from a set and then rank them), Lowe et al. (2015) proposed a neural architecture to select the best next response from a list of responses by measuring their similarity to the dialogue context. Serban et al. (2016) extended prior work on encoder-decoder-based models to multi-turn conversations. They trained a hierarchical model called HRED for generating dialogue utterances where a recurrent neural network encoder encodes each utterance. A higher-level recurrent neural network maintains the dialogue state by further encoding the individual utterance encodings. This dialogue state is then decoded by another recurrent decoder to generate the response at that point in time. In followup work, Serban et al. (2017) used a latent stochastic variable to condition the generation process which aided their model in producing longer coherent outputs that better retain the context.

Datasets and tasks (Das et al. 2017a; Mostafazadeh et al. 2017; De Vries et al. 2017) have also been released recently to study visual-input based conversations. Das et al. (2017a) train several generative and discriminative deep neural models for the VisDial task. They observe that on this task, discriminative models outperform generative models and that models making better use of the dialogue history do better than models that do not use dialogue history at all. Unexpectedly, the performance between models that use the image features and models that do not use these features is not significantly different. As we discussed in Section 1, this is similar to the issues VQA models faced initially due to the imbalanced nature of the dataset, which leads us to believe that language is a strong prior on the VisDial dataset too. Das et al. (2017b) train two separate agents to play a cooperative game where one agent has to answer the other agent’s questions, which in turn has to predict the *fc7* features of

the Image obtained from VGGNet. Both agents are based on HRED models and they show that agents fine-tuned with Reinforcement Learning (RL) outperform agents trained solely with supervised learning. Mostafazadeh et al. (2017) train both generative and discriminative deep neural models on the Image-Grounded Conversations (IGC) dataset, where the task is to generate questions and answers to carry on a meaningful conversation. De Vries et al. (2017) train HRED-based models on GuessWhat?! dataset in which agents have to play a guessing game where one player has to find an object in the picture which the other player knows about and can answer questions about them.

Moving from image-based dialogue to video-based dialogue adds further complexity and challenges. Limited availability of such data is one of the challenges. Apart from the AVSD dataset, there does not exist a video dialogue dataset to the best of our knowledge and the AVSD data itself is fairly limited in size. Extracting relevant features from videos also contains the inherent complexity of extracting features from individual frames and additionally requires understanding their temporal interaction. The temporal nature of videos also makes it important to be able to focus on a varying-length subset of video frames as the action which is being asked about might be happening within them. There is also the need to encode the additional modality of audio which would be required for answering questions that rely on the audio track. With limited size of publicly available datasets based on the visual modality, learning useful features from high dimensional visual data has been a challenge even for the VisDial dataset, and we anticipate this to be an even more significant challenge on the AVSD dataset as it involves videos.

On the AVSD task, Hori et al. (2018) train an attention-based audio-visual scene-aware dialogue model which we use as the baseline model for this paper. They divide each video into multiple equal-duration segments and, from each of them, extract video features using an I3D (Carreira and Zisserman 2017) model, and audio features using a VGGish (Hershey et al. 2017) model. The I3D model was pre-trained on Kinetics (Kay et al. 2017) dataset and the VGGish model was pre-trained on Audio Set (Gemmeke et al. 2017). The baseline encodes the current utterance’s question with a Long Short-Term Memory (LSTM) (Hochreiter and Schmidhuber 1997) and uses the encoding to attend to the audio and video features from all the video segments and to fuse them together. The dialogue history is modelled with a hierarchical recurrent LSTM encoder where the input to the lower level encoder is a concatenation of question-answer pairs. The fused feature representation is concatenated with the question encoding and the dialogue history encoding and the resulting vector is used to decode the current answer using an LSTM decoder. Similar to the VisDial models, the performance difference between the best model that uses text and the best model that uses both text and video features is small. This indicates that the language is a stronger prior here and the baseline model is unable to make good use of the highly relevant video.

Automated evaluation of both task-oriented and non-task-oriented dialogue systems has been a challenge (Sharma et al. 2017; Liu et al. 2016) too. Most such dialogue systems are evaluated using per-turn evaluation metrics since there is

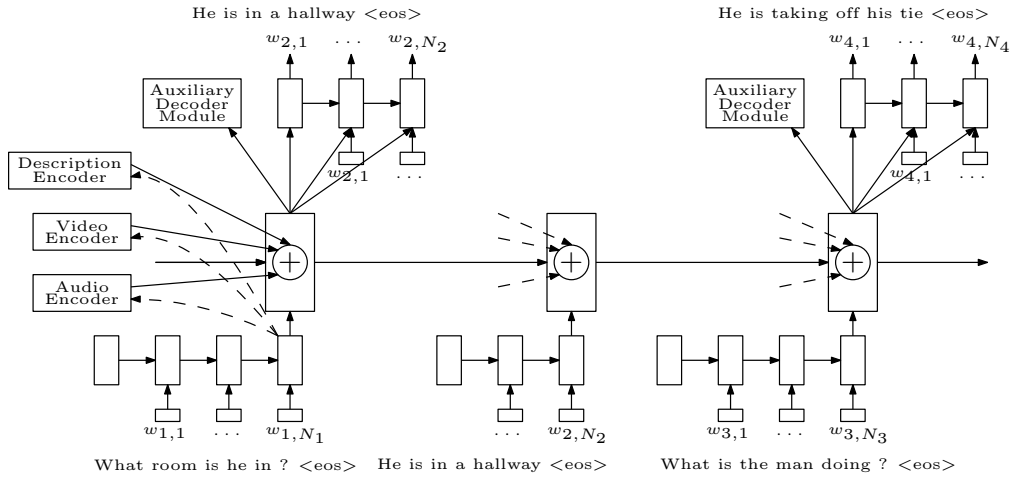


Figure 1: FA-HRED uses the last question’s encoding to attend to video description, audio, and video features. These features along with the dialogue state enable the model to generate the answer to the current question. The ground truth answer is encoded into the dialogue history for the next turn.

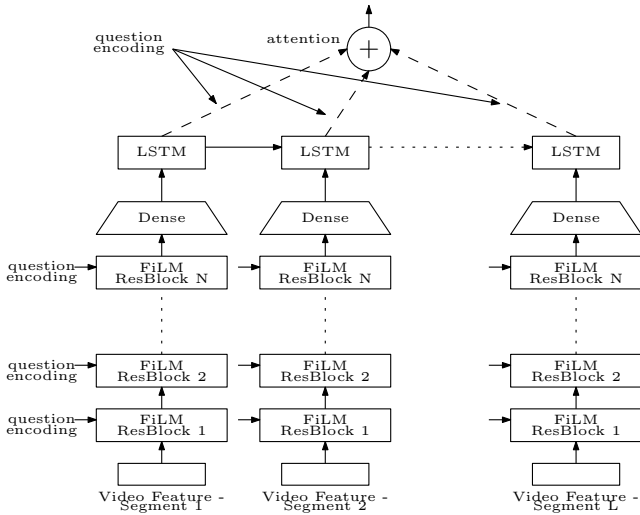


Figure 2: Video Encoder Module: FiLM for video features. Question encoding of the current question is used here.

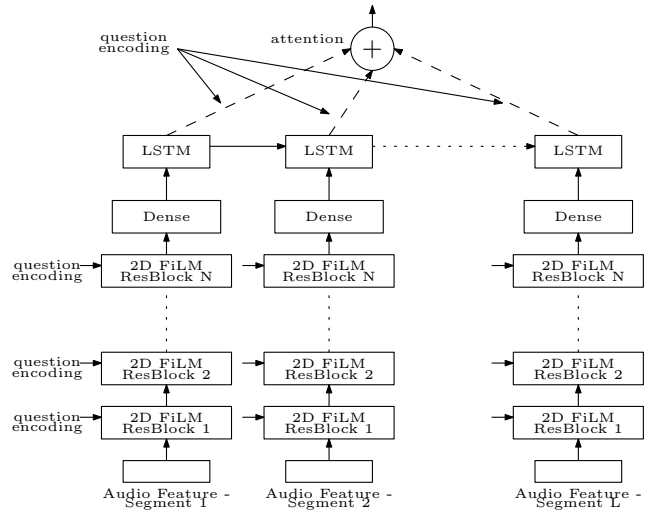


Figure 3: Audio Encoder Module: FiLM for audio features. Question encoding of the current question is used here.

no suitable per-dialogue metric as conversations do not need to happen in a deterministic ordering of turns. These per-turn evaluation metrics are mostly word-overlap-based metrics such as BLEU, METEOR, ROUGE, and CIDEr, borrowed from the machine translation literature. Due to the diverse nature of possible responses, world-overlap metrics are not highly suitable for evaluating these tasks. Human evaluation of generated responses is considered the most reliable metric for such tasks but it is cost prohibitive and hence the dialogue system literature continues to rely widely on word-overlap-based metrics.

3 The AVSD dataset and challenge

The AVSD dataset (Alamri et al. 2018) consists of dialogues collected via Amazon Mechanical Turk (AMT). Each dialogue is associated with a video from the Charades (Sigurd-

son et al. 2016) dataset and has conversations between two AMT workers related to the video. The Charades dataset has multi-action short videos and it provides text descriptions for these videos, which the AVSD challenge also distributes as the caption. The AVSD dataset has been collected using similar methodology as the VisDial dataset. In AVSD, each dialogue turn consists of a question and answer pair. One of the AMT workers assumes the role of *questioner* while the other AMT worker assumes the role of *answerer*. The *questioner* sees three static frames from the video and has to ask questions. The *answerer* sees the video and answers the questions asked by the *questioner*. After 10 such QA turns, the *questioner* wraps up by writing a summary of the video based on the conversation.

Dataset statistics such as the number of dialogues, turns, and words for the AVSD dataset are presented in Table 2.

	training	validation	test
# dialogs	7 659	1 787	1 710
# turns	153 180	35 740	13 490
# words	1 450 754	339 006	110 252
# dialogs	6 172	732	733
# turns	123 480	14 680	14 660
# words	1 163 969	138 314	138 790

Table 2: AVSD: Dataset Statistics. Top: official dataset. Bottom half: prototype dataset released earlier.

For the initially released prototype dataset, the training set of the AVSD dataset corresponds to videos taken from the training set of the Charades dataset while the validation and test sets of the AVSD dataset correspond to videos taken from the validation set of the Charades dataset. For the official dataset, training, validation and test sets are drawn from the corresponding Charades sets.

The Charades dataset also provides additional annotations for the videos such as action, scene, and object annotations, which are considered to be external data sources by the AVSD challenge, for which there is a special sub-task in the challenge. The action annotations also include the start and end time of the action in the video.

4 Models

Our FA-HRED model is based on the HRED framework for modelling dialogue systems. In our model, an utterance-level recurrent LSTM encoder encodes utterances and a dialogue-level recurrent LSTM encoder encodes the final hidden states of the utterance-level encoders, thus maintaining the dialogue state and dialogue coherence. We use the final hidden states of the utterance-level encoders in the attention mechanism that is applied to the outputs of the description, video, and audio encoders. The attended features from these encoders are fused with the dialogue-level encoder’s hidden states. An utterance-level decoder decodes the response for each such dialogue state following a question. We also add an auxiliary decoding module which is similar to the response decoder except that it tries to generate the caption and/or the summary of the video. We present our model in Figure 1 and describe the individual components in detail below.

Utterance-level Encoder

The utterance-level encoder is a recurrent neural network consisting of a single layer of LSTM cells. The input to the LSTMs are word embeddings for each word in the utterance. The utterance is concatenated with a special symbol $\langle \text{eos} \rangle$ marking the end of the sequence. We initialize our word embeddings using 300-dimensional GloVe (Pennington, Socher, and Manning 2014) and then fine-tune them during training. For words not present in the GloVe vocabulary, we initialize their word embeddings from a random uniform distribution.

Description Encoder

Similar to the utterance-level encoder, the description encoder is also a single-layer LSTM recurrent neural network. Its word embeddings are also initialized with GloVe and then fine-tuned during training. For the description, we use the caption and/or the summary for the video provided with the dataset. The description encoder also has access to the last hidden state of the utterance-level encoder, which it uses to generate an attention map over the hidden states of its LSTM. The final output of this module is the attention-weighted sum of the LSTM hidden states.

Video Encoder with Time-Extended FiLM

For the video encoder, we use an I3D model pre-trained on the Kinetics dataset (Kay et al. 2017) and extract the output of its *Mixed_7c* layer for L (30 for our models) equi-distant segments of the video. Over these features, we add N (2 for our models) FiLM (Perez et al. 2018) blocks which have been highly successful in visual reasoning problems. Each FiLM block applies a conditional (on the utterance encoding) feature-wise affine transformation on the features input to it, ultimately leading to the extraction of more relevant features. The FiLM blocks are followed by fully connected layers which are further encoded by a single layer recurrent LSTM network. The last hidden state of the utterance-level encoder then generates an attention map over the hidden states of its LSTM, which is multiplied by the hidden states to provide the output of this module. We also experimented with using convolutional *Mixed_5c* features to capture spatial information but on the limited AVSD dataset they did not yield any improvement. When not using the FiLM blocks, we use the final layer I3D features (provided by the AVSD organizers) and encode them with the LSTM directly, followed by the attention step. We present the video encoder in Figure 2.

Audio Encoder

The audio encoder is structurally similar to the video encoder. We use the VGGish features provided by the AVSD challenge organizers. Also similar to the video encoder, when not using the FiLM blocks, we use the VGGish features and encode them with the LSTM directly, followed by the attention step. The audio encoder is depicted in Figure 3.

Fusing Modalities for Dialogue Context

The outputs of the encoders for past utterances, descriptions, video, and audio together form the dialogue context C which is the input of the decoder. We first combine past utterances using a dialogue-level encoder which is a single-layer LSTM recurrent neural network. The input to this encoder are the final hidden states of the utterance-level LSTMs. To combine the hidden states of these diverse modalities, we found concatenation to perform better on the validation set than averaging or the Hadamard product.

Decoders

The *answer decoder* consists of a single-layer recurrent LSTM network and generates the answer to the last question utterance. At each time-step, it is provided with the

	BLEU-1	BLEU-2	BLEU-3	BLEU-4	METEOR	ROUGE	CIDEr
Challenge Baseline	0.626	0.485	0.383	0.309	0.215	0.487	0.746
(Task 1.a.i) FA-HRED w/o summary	0.648	0.505	0.399	0.323	0.231	0.510	0.843
(Task 1.a.ii) FA-HRED w/ summary	0.695	0.553	0.444	0.360	0.249	0.544	0.997
(Task 2.a.i) FA-HRED w/o summary	0.662	0.520	0.416	0.340	0.228	0.518	0.851
(Task 2.a.ii) FA-HRED w/ summary	0.686	0.541	0.429	0.343	0.243	0.536	0.920
Winning Team (Task 1.a.ii)	0.718	0.584	0.478	0.394	0.267	0.563	1.094
(Task 1.a.i) FA-HRED w/o summary	0.650	0.506	0.397	0.316	0.224	0.505	0.795
(Task 1.a.ii) FA-HRED w/ summary	0.685	0.542	0.433	0.349	0.242	0.536	0.947
(Task 2.a.i) FA-HRED w/o summary	0.635	0.500	0.398	0.323	0.220	0.501	0.799
(Task 2.a.ii) FA-HRED w/ summary	0.656	0.507	0.398	0.319	0.228	0.513	0.836
Previous state-of-the-art (Hori et al. 2018)	0.256	0.161	0.109	0.078	0.113	0.277	0.727
(Task 1.a.i) FA-HRED w/o summary	0.300	0.189	0.131	0.095	0.136	0.362	0.968
(Task 1.a.ii) FA-HRED w/ summary	0.312	0.197	0.136	0.098	0.140	0.362	0.993
(Task 2.a.i) FA-HRED w/o summary	0.299	0.187	0.128	0.093	0.135	0.360	0.961
(Task 2.a.ii) FA-HRED w/ summary	0.297	0.187	0.130	0.095	0.136	0.358	0.987

Table 3: Scores achieved by our model on different tasks of the AVSD challenge test set. Task 1 model configurations use both video and text features while Task 2 model configurations only use text features. **First section:** train on official, test on official. **Second section:** train on prototype, test on official. **Third section:** train on prototype, test on prototype.

dialogue-level state and produces a softmax over a vector corresponding to vocabulary words and stops when 30 words were produced or an end of sentence token is encountered.

The *auxiliary decoder* is functionally similar to the *answer decoder*. The decoded sentence is the caption and/or description of the video. We use the Video Encoder state instead of the Dialogue-level Encoder state as input since with this module we want to learn a better video representation capable of decoding the description.

Loss Function

For a given context embedding C_t at dialogue turn t , we minimize the negative log-likelihood of the ground truth response words $0 \leq r_{t,m} \leq V$ (vocabulary size), normalized by the number of words M in the ground truth response \mathbf{r} ,

$$\mathcal{L}(C_t, \mathbf{r}) = -\frac{1}{M} \sum_{m=1}^M \sum_i^V ([r_{t,m}=i] \log p(r_{t,m}=i | C_t, r_{t,m-1}^*)),$$

where the probabilities $p(\cdot | \cdot)$ are given by the decoder LSTM output,

$$r_{t,m-1}^* = \begin{cases} r_{t,m-1} & ; s > 0.2, s \sim U(0, 1) \\ v \sim p(r_{t,m-1} | C_t, r_{t,m-2}^*) & ; \text{else} \end{cases}$$

is given by scheduled sampling (Bengio et al. 2015), and $r_{t,0}^*$ is a symbol denoting the start of a sequence. We optimize the model using the AMSGrad algorithm (Reddi, Kale, and Kumar 2018) and use a per-condition random search to determine hyperparameters. We train the model using the BLEU-4 score on the validation set as our stopping criterion.

5 Experiments

The AVSD challenge tasks we address here are:

1. Video and Text

- 1.a. Allowed to use publicly available pre-trained models for feature extraction + QA + captions
 - 1.a.i. Not allowed to use summary
 - 1.a.ii. Also allowed to use summary
2. Text only
 - 2.a. Allowed to use QA + captions
 - 2.a.i. Not allowed to use summary
 - 2.a.ii. Also allowed to use summary

We train our FA-HRED model for Task 1.a and Task 2.a of the challenge and we present the results in Table 3. Our model outperforms the baseline model released by Hori et al. (2018) on all of these tasks. The scores for the winning team have been released to challenge participants and are also included. Their approach, however, is not public as of yet. We observe the following for our models:

- **Comparing Task 1.a.i with 2.a.i and 1.a.ii with 2.a.ii:** Models that use both video and text data largely perform better than the models that just use the same text data. This is expected as the video features are highly relevant to answering the question asked. Based on findings presented by Hori et al. (2018) (which we discuss in Section 2), and due to the limited size of the dataset, we did not expect a large performance boost by adding the video features. As we can observe in Table 3, the results are in line with these expectations.
- **Comparing Task 1.a.i with 1.a.ii and 2.a.i with 2.a.ii:** During the data collection phase, since the *questioner* wrote the summary of the video at the end, we expected the summary to be extremely helpful in answering the *questioner*'s questions. In line with this, we see that models that use the summary description encoder perform better than the models that use the caption description on almost all the metrics, including the BLEU-4 metric, the

Model with (+) / without (-) inputs	FiLM	No FiLM
Attention + I3D + VGGish + Caption	0.0986	0.0953
-Caption	0.0951	0.0960
-Caption -VGGish	0.0985	0.0974
-I3D -VGGish	–	0.0967
-I3D -Caption	0.0960	0.0960
Attention + I3D + VGGish + Summary	0.1052	0.0989
-VGGish	0.1045	0.1022
-I3D -VGGish	–	0.1007
-I3D	0.1004	0.0999

Table 4: Model ablation Study comparing BLEU-4 on the validation set: The best model makes use of all modalities and the video summary. Applying FiLM to audio and video features consistently outperforms unconditioned feature extraction. Video features (I3D) are more important than audio (VGGish). Combining all multi-modal components (e.g., text, audio and video) helps improve performance only when using FiLM blocks.

METEOR metric, and the CIDEr metric, all of which are highly popular metrics in the machine translation community. Any deviations from this trend have a very small margin. Evaluation metrics for such dialogue systems do not correlate much with human evaluation, as we discussed in Section 2, which makes this difference challenging to interpret. However, all of our models significantly improve upon the competitive baseline model.

Since the official test set has not been released publicly, results reported on the official test set have been provided by the challenge organizers. For the prototype test set and for the ablation study presented in Table 4, we use the same code¹ for evaluation metrics as used by Hori et al. (2018) for fairness and comparability. We attribute the significant performance gain of our model over the baseline to a combination of several factors as described below:

- **Stronger encoder-decoder pair:** We use the HRED framework to compute utterance descriptors for questions and answers independently at the lower hierarchy level, instead of encoding their concatenation. This formulation removes the need for a separate question encoder and relies less on the long term memory of the utterance encoder. It improves validation performance by 0.015 BLEU-4 over a simple LSTM baseline.
- **Scheduled sampling:** In our experiments, we found that models trained with scheduled sampling performed better (about 0.004 BLEU-4 on validation set) than the ones trained using teacher-forcing for the AVSD dataset. Hence, we use scheduled sampling for all the results we report in this paper.
- **Regularization:** Due to the limited size of the AVSD dataset, we used dropout with retention probability be-

tween 0.6 and 0.9 for all our models during training to prevent over-fitting.

- **Feature extraction with FiLM:** We extract I3D and VGGish features using FiLM. The ablation study in Table 4 demonstrates that using FiLM provides a performance boost due to the extraction of features more relevant for the current question.

Our primary architectural differences over the baseline model are: not concatenating the question, answer pairs before encoding them, the auxiliary decoder module, and using the Time-Extended FiLM module for feature extraction. These, combined with using scheduled sampling and running hyperparameter optimization over the validation set to select hyperparameters, give us the observed performance boost.

We observe that our models generate fairly relevant responses to questions in the dialogues, and models with audio-visual inputs respond to audio-visual questions (e.g. “is there any voices or music ?”) correctly more often.

We conduct an ablation study on the effectiveness of different components (eg., text, video and audio) and present it in Table 4. Our experiments show that:

1. Overall, models with audio, video, and summary performance better than models using audio, video, and caption.
2. Using FiLM blocks for feature extraction consistently helps in improving performance across all settings.
3. Video features (I3D) are more important than audio (VGGish) for performance on the AVSD dataset.
4. Caption and summary are the most important components. Interestingly, when not using FiLM, the performance decreases when combining text components with audio/video (0.1007 vs 0.0989 or 0.0967 vs 0.0953). By contrast, performance increases when using FiLM blocks to combine all multimodal components (0.1007 vs 0.1052 or 0.0967 vs 0.0986). We conclude that our FiLM implementation reduces overfitting compared to averaging or concatenation alternatives, when combining information from very different multiple modalities.

6 Conclusions

We presented FA-HRED, a state-of-the-art dialogue model for conversations about videos. We evaluated the model on the official AVSD test set, where it achieves a relative improvement of more than 16% over the baseline model on BLEU-4 and more than 33% on CIDEr. The challenging aspect of multi-modal dialogue is fusing modalities with varying information density. On AVSD, it is easiest to learn from the input text, while video features remain largely opaque to the decoder. FA-HRED uses a generalization of FiLM to video that conditions video feature extraction on a question. However, similar to related work, absolute improvements of incorporating video features into dialogue are consistent but small. Thus, while our results indicate the suitability of our FiLM generalization, they also highlight that applications at the intersection between language and video are currently constrained by the quality of video features, and emphasizes the need for larger datasets.

¹Sharma et al. (2017), <https://github.com/Maluuba/nlg-eval>

References

- Agrawal, A.; Batra, D.; and Parikh, D. 2016. Analyzing the behavior of visual question answering models. In *EMNLP*.
- 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. Technical Report TR2018-069, MERL.
- Antol, S.; Agrawal, A.; Lu, J.; Mitchell, M.; Batra, D.; Lawrence Zitnick, C.; and Parikh, D. 2015. VQA: Visual question answering. In *ICCV*.
- Bengio, S.; Vinyals, O.; Jaitly, N.; and Shazeer, N. 2015. Scheduled sampling for sequence prediction with recurrent neural networks. In *NIPS*.
- Carreira, J., and Zisserman, A. 2017. Quo vadis, action recognition? a new model and the kinetics dataset. In *CVPR*.
- Cho, K.; Van Merriënboer, B.; Gulcehre, C.; Bahdanau, D.; Bougares, F.; Schwenk, H.; and Bengio, Y. 2014. Learning phrase representations using RNN encoder-decoder for statistical machine translation. In *EMNLP*.
- Das, A.; Kottur, S.; Gupta, K.; Singh, A.; Yadav, D.; Moura, J. M.; Parikh, D.; and Batra, D. 2017a. Visual Dialog. In *CVPR*.
- Das, A.; Kottur, S.; Moura, J. M.; Lee, S.; and Batra, D. 2017b. Learning cooperative visual dialog agents with deep reinforcement learning. In *ICCV*.
- De Vries, H.; Strub, F.; Chandar, S.; Pietquin, O.; Larochelle, H.; and Courville, A. C. 2017. GuessWhat?! visual object discovery through multi-modal dialogue. In *CVPR*.
- Gemmeke, J. F.; Ellis, D. P. W.; Freedman, D.; Jansen, A.; Lawrence, W.; Moore, R. C.; Plakal, M.; and Ritter, M. 2017. Audio set: An ontology and human-labeled dataset for audio events. In *ICASSP*.
- Goyal, Y.; Khot, T.; Summers-Stay, D.; Batra, D.; and Parikh, D. 2017. Making the V in VQA matter: Elevating the role of image understanding in visual question answering. In *CVPR*.
- Hermann, K. M.; Kocisky, T.; Grefenstette, E.; Espeholt, L.; Kay, W.; Suleyman, M.; and Blunsom, P. 2015. Teaching machines to read and comprehend. In *NIPS*.
- Hershey, S.; Chaudhuri, S.; Ellis, D. P. W.; Gemmeke, J. F.; Jansen, A.; Moore, C.; Plakal, M.; Platt, D.; Saurous, R. A.; Seybold, B.; Slaney, M.; Weiss, R.; and Wilson, K. 2017. CNN architectures for large-scale audio classification. In *ICASSP*.
- Hochreiter, S., and Schmidhuber, J. 1997. Long short-term memory. *Neural computation* 9(8).
- Hori, C.; Alamri, H.; Wang, J.; Winchern, G.; Hori, T.; Cherian, A.; Marks, T. K.; Cartillier, V.; Lopes, R. G.; Das, A.; et al. 2018. End-to-end audio visual scene-aware dialog using multimodal attention-based video features. *arXiv:1806.08409*.
- Johnson, J.; Hariharan, B.; van der Maaten, L.; Fei-Fei, L.; Zitnick, C. L.; and Girshick, R. 2017. CLEVR: A diagnostic dataset for compositional language and elementary visual reasoning. In *CVPR*.
- Karpathy, A.; Toderici, G.; Shetty, S.; Leung, T.; Sukthankar, R.; and Fei-Fei, L. 2014. Large-scale video classification with convolutional neural networks. In *CVPR*.
- Kay, W.; Carreira, J.; Simonyan, K.; Zhang, B.; Hillier, C.; Vijayanarasimhan, S.; Viola, F.; Green, T.; Back, T.; Natsev, P.; et al. 2017. The kinetics human action video dataset. *arXiv:1705.06950*.
- Krizhevsky, A.; Sutskever, I.; and Hinton, G. E. 2012. Imagenet classification with deep convolutional neural networks. In *NIPS*.
- Liu, C.-W.; Lowe, R.; Serban, I.; Noseworthy, M.; Charlin, L.; and Pineau, J. 2016. How not to evaluate your dialogue system: An empirical study of unsupervised evaluation metrics for dialogue response generation. In *EMNLP*.
- Lowe, R.; Pow, N.; Serban, I.; and Pineau, J. 2015. The ubuntu dialogue corpus: A large dataset for research in unstructured multi-turn dialogue systems. In *SIGDIAL*.
- Mostafazadeh, N.; Brockett, C.; Dolan, B.; Galley, M.; Gao, J.; Spithourakis, G. P.; and Vanderwende, L. 2017. Image-grounded conversations: Multimodal context for natural question and response generation. In *IJCNLP*.
- Pennington, J.; Socher, R.; and Manning, C. 2014. Glove: Global vectors for word representation. In *EMNLP*.
- Perez, E.; Strub, F.; de Vries, H.; Dumoulin, V.; and Courville, A. 2018. FiLM: Visual reasoning with a general conditioning layer. In *AAAI*.
- Reddi, S. J.; Kale, S.; and Kumar, S. 2018. On the convergence of adam and beyond. In *ICLR*.
- Ritter, A.; Cherry, C.; and Dolan, W. B. 2011. Data-driven response generation in social media. In *EMNLP*.
- Serban, I. V.; Sordoni, A.; Bengio, Y.; Courville, A.; and Pineau, J. 2016. Building end-to-end dialogue systems using generative hierarchical neural network models. In *AAAI*.
- Serban, I.; Sordoni, A.; Lowe, R.; Charlin, L.; Pineau, J.; Courville, A.; and Bengio, Y. 2017. A hierarchical latent variable encoder-decoder model for generating dialogues. In *AAAI*.
- Shang, L.; Lu, Z.; and Li, H. 2015. Neural responding machine for short-text conversation. In *IJCNLP*.
- Sharma, S.; Asri, L. E.; Schulz, H.; and Zumer, J. 2017. Relevance of unsupervised metrics in task-oriented dialogue for evaluating natural language generation. *arXiv:1706.09799*.
- 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 *ECCV*.
- Sordoni, A.; Galley, M.; Auli, M.; Brockett, C.; Ji, Y.; Mitchell, M.; Nie, J.-Y.; Gao, J.; and Dolan, B. 2015. A neural network approach to context-sensitive generation of conversational responses. In *NAACL: HLT*.
- Szegedy, C.; Toshev, A.; and Erhan, D. 2013. Deep neural networks for object detection. In *NIPS*.
- Tapaswi, M.; Zhu, Y.; Stiefelhagen, R.; Torralba, A.; Urtasun, R.; and Fidler, S. 2016. MovieQA: Understanding stories in movies through question-answering. In *CVPR*.