

# Building Sequential Inference Models for End-to-End Response Selection

Jia-Chen Gu<sup>1</sup>, Zhen-Hua Ling<sup>1</sup>, Yu-Ping Ruan<sup>1</sup> and Quan Liu<sup>1,2</sup>

<sup>1</sup>National Engineering Laboratory for Speech and Language Information Processing,  
University of Science and Technology of China, Hefei, China

<sup>2</sup>iFLYTEK Research, Hefei, China

gujc@mail.ustc.edu.cn, zhling@ustc.edu.cn, ypruan@mail.ustc.edu.cn, quanliu@ustc.edu.cn

## Abstract

This paper presents an end-to-end response selection model for Track 1 of the 7th Dialogue System Technology Challenges (DSTC7). This task focuses on selecting the correct next utterance from a set of candidates given a partial conversation. We propose an end-to-end neural network based on enhanced sequential inference model (ESIM) for this task. Our proposed model differs from the original ESIM model in the following four aspects. First, a new word representation method which combines the general pre-trained word embeddings with those estimated on the task-specific training set is adopted in order to address the challenge of out-of-vocabulary (OOV) words. Second, an attentive hierarchical recurrent encoder (AHRE) is designed which is capable to encode sentences hierarchically and generate more descriptive representations by aggregation. Third, a new pooling method which combines multi-dimensional pooling and last-state pooling is used instead of the simple combination of max pooling and average pooling in the original ESIM. Last, a modification layer is added before the softmax layer to emphasize the importance of the last utterance in the context for response selection. In the released evaluation results of DSTC7, our proposed method ranked second on the Ubuntu dataset and third on the Advising dataset in subtask 1 of Track 1.

## Introduction

Building dialogue systems that can converse naturally with humans is a challenging yet intriguing problem of artificial intelligence. Recently, human-computer conversation has attracted increasing attention due to its promising potentials and alluring commercial values. According to the applications, dialogue systems can be roughly divided into two categories: (1) task-oriented systems and (2) non-task-oriented systems (also known as chatbots). Task-oriented systems aim to assist the user to complete certain tasks (e.g. booking accommodations and restaurants). Non-task-oriented systems aim to engage users in human-computer conversations in the open domain and attract lots of research attentions because they target on unstructured dialogues without a priori logical representation for the information exchanging during the conversation.

Existing approaches to dialogue response generation includes generation-based methods (Shang, Lu, and Li 2015; Serban et al. 2016) and retrieval-based methods (Zhou et al. 2016; Wu et al. 2017). Generation-based models maximize the probability of generating a response given the previous dialogue. This approach enables the incorporation of rich context when mapping between consecutive dialogue turns. Retrieval-based methods select a proper response for the current conversation from a repository with response selection algorithms, and have the advantage of producing informative and fluent responses. Track 1 of the 7th Dialogue System Technology Challenges (DSTC7) is a kind of retrieval-based task which selects the correct response from a large set of candidates. The set used in this track contains more candidates than many other datasets. Some candidates are also similar which increases the difficulty of making right decisions.

The techniques of word embeddings and sentence embeddings are important to response selection as well as many other natural language processing (NLP) tasks. The context and the response must be projected to a vector space appropriately in order to capture the relationships between them, which are essential for following procedures. Recently there has been a growing interest in models for word-level (Mikolov et al. 2013; Pennington, Socher, and Manning 2014; Dong and Huang 2018) and sentence-level (Wang, Hamza, and Florian 2017; Chen et al. 2017) representations using neural networks, which helped classification or inference algorithms to achieve better performance in many NLP tasks.

Another key technique to the response selection task lies in context-response matching. Modeling the semantic matching degree between two sentences is challenging. The enhanced sequential inference model (ESIM) (Chen et al. 2017) was proposed to measure the relationship between a pair of sentences in natural language inference (NLI) tasks. This model described the interactions between two sentences by sequential encoding and attention-based alignment. Considering the good performance and decomposable implementation of ESIM, it is adopted as our baseline model for response selection.

This paper introduces the end-to-end response selection method developed by us for subtask 1 of Track 1 in DSTC7. We propose to improve the original ESIM model for

Speaker	Utterances in a dialogue
A	Hmm, perhaps I should return to windows...
B	what kind of cd-rom do you have??
A	It's a external USB drive.... Any ideas?
B	what cd-rom do you have??
A	An external LG USB
Candidates:	
1. perhaps it's not running?	
2. is jack running properly?	
...	
x. sry didnt see your answer, write my name so i see it.	
...	
100. odd sized card?. why did you run mklable?	
Answer:	
sry didnt see your answer, write my name so i see it.	

Table 1: Dialogue example of subtask 1.

response selection from the following four aspects.

- A new word representation method which combines the general pre-trained word embeddings with those estimated on the task-specific training set is adopted in order to address the challenge of out-of-vocabulary (OOV) words.
- An attentive hierarchical recurrent encoder (AHRE) is designed which encodes sentences hierarchically and generates sentence representations by aggregation.
- A new pooling method which combines multi-dimensional pooling and last-state pooling is used instead of the simple combination of max pooling and average pooling in the original ESIM.
- A modification layer is added before the softmax layer to emphasize the importance of the last utterance in the context for response selection.

As shown in the released challenge results, our proposed model ranked second on the Ubuntu dataset and ranked third on the Advising dataset in subtask 1 of Track 1. In the following sections, we first introduce the task descriptions of Track 1 in DSTC7, and present the details of our proposed model. Then the model configurations, training settings and evaluation results are shown. Furthermore, the experimental results are analyzed by ablation tests. Finally we draw conclusions and give an overview of our future work.

## Task Description

The DSTC7 Track 1 organizers provided two datasets (Kummerfeld et al. 2018). One is Ubuntu Dialogue Corpus which contains dialogues between Ubuntu users for the purpose of solving an Ubuntu users posted problem and the other is Advising Data which consists of dialogues between a student and a advisor for the purpose of guiding the student to pick courses.

The task is divided into 5 subtasks and a participant may participate in one, several, or all the subtasks. Participants are required to meet different goals for different subtasks such as selecting the next utterance from the given 100 candidates or 120k candidates, selecting the next utterance

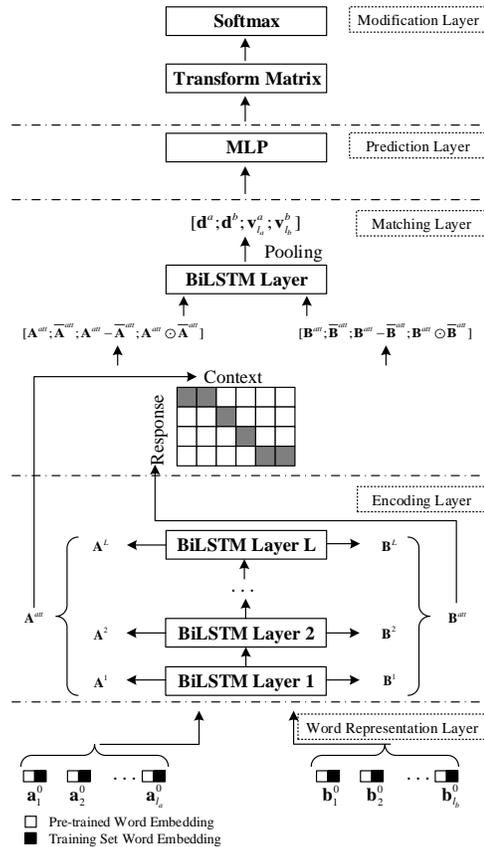


Figure 1: Diagram of our proposed model.

with the set of paraphrases, selecting the next utterance with a candidate pool which might not include the correct next utterance, and selecting the next utterance with a model incorporating external knowledge. Each subtask has its corresponding dataset and each dialogue in it has its corresponding response candidates together with the correct answer. Due to limited time and manpower, we only participate subtask 1 of this track, which aims to select the next utterance from a candidate set of 100 utterances. An example dialogue and its candidates is shown in Table 1.

## Model Description

Our proposed model is composed of five components: Word Representation Layer, Encoding Layer, Matching Layer, Prediction Layer and Modification Layer. Figure 1 shows the diagram of the model architecture. Details about each layer are described in this section.

### Word Representation Layer

One challenge of modeling dialogue is the large number of out-of-vocabulary words. To address this issue, we adopt an algorithm (Dong and Huang 2018) which combines the general pre-trained word embedding vectors with those generated on the task-specific training set to enhance word representations.

## Encoding Layer

Recurrent neural networks (RNN) (Mikolov et al. 2010) have been proven to be good at modeling chronological relationship in language sequences and multi-layer RNNs have achieved good performance in many NLP tasks such as neural machine translation (NMT) (Bahdanau, Cho, and Bengio 2014) and natural language inference (NLI) (Chen et al. 2017). Encoding the sequences with deep neural networks can help capture deeper and more useful information. Typically, the outputs of the top RNN layer are regarded as the final sentence representations and the other layers are neglected. However, the lower layers can also provide useful sentence descriptions, such as part-of-speech tagging and syntax-related ones (Hashimoto et al. 2017).

To make full use of the representations at all hidden layers, we propose a new sentence encoder called attentive hierarchical recurrent encoder (AHRE). This encoder is motivated by the method of embeddings from language models (ELMo) (Peters et al. 2018) which combines the internal states of multi-layer RNNs. More specifically, an AHRE learns a linear combination of the vectors stacked above each input word, which improves the performance of just using the top RNN layer in our experiments.

Let  $\mathbf{A}^0 = [\mathbf{a}_1^0, \dots, \mathbf{a}_{l_a}^0]$  and  $\mathbf{B}^0 = [\mathbf{b}_1^0, \dots, \mathbf{b}_{l_b}^0]$  denote sequences of word representations of context and response respectively.  $l_a$  and  $l_b$  are token numbers in these two sequences. Both  $\mathbf{a}_i^0 \in \mathbb{R}^l$  and  $\mathbf{b}_j^0 \in \mathbb{R}^l$  are  $l$ -dimensional embedding vectors given by the word representation layer mentioned above. Furthermore, bidirectional LSTMs (BiLSTM) (Hochreiter and Schmidhuber 1997) are employed as our basic building blocks. In an  $L$ -layer RNN, the  $l^{th}$  layer takes the output of the  $l-1^{th}$  layer as its input. We denote the calculations as follows,

$$\mathbf{a}_i^l = \text{BiLSTM}(\mathbf{a}_i^{l-1}, i), i \in \{1, \dots, l_a\}, l \in \{1, \dots, L\}, \quad (1)$$

$$\mathbf{b}_j^l = \text{BiLSTM}(\mathbf{b}_j^{l-1}, j), j \in \{1, \dots, l_b\}, l \in \{1, \dots, L\}. \quad (2)$$

The weights for these two BiLSTMs are shared in our implementation. Due to limit space, we skip the descriptions on the basic chain LSTMs and readers can refer to (Hochreiter and Schmidhuber 1997) for details.

Finally we get a set of  $L$  representations  $\{\mathbf{a}^1, \dots, \mathbf{a}^L\}$  and  $\{\mathbf{b}^1, \dots, \mathbf{b}^L\}$  through the  $L$ -layer RNNs. Typically  $\mathbf{a}^L$  or  $\mathbf{b}^L$ , i.e. the outputs of the top layer, are used as the final encoded vectors. Here, we propose to combine the set of representations to get enhanced representations  $\mathbf{a}_i^{att}$  and  $\mathbf{b}_j^{att}$  by learning attention weights of all layers. Mathematically, we have

$$\mathbf{a}_i^{att} = \sum_{l=1}^L w_l \mathbf{a}_i^l, i \in \{1, \dots, l_a\}, \quad (3)$$

$$\mathbf{b}_j^{att} = \sum_{l=1}^L w_l \mathbf{b}_j^l, j \in \{1, \dots, l_b\}, \quad (4)$$

where  $w_l$  are softmax-normalized weights shared between context and response which need to be estimated during the training process. Our representations differ from those of traditional encoder in that ours not only considers the

top layer representations but also takes the lower layer representations which may be informative into account. As a result, the representations given by our encoder are expected to capture and fuse multi-level characteristics of sentences.

## Matching Layer

Interactions between context and response is important to provide information for deciding the matching degree between them. Our model follows the matching part of ESIM (Chen et al. 2017) which collects local information between two sentences by attention-based alignment and is fully computationally decomposable.

First, a soft alignment is conducted by computing the attention weight between each representation tuple  $\{\mathbf{a}_i^{att}, \mathbf{b}_j^{att}\}$  as

$$e_{ij} = (\mathbf{a}_i^{att})^T \cdot \mathbf{b}_j^{att}. \quad (5)$$

Then, local inference is determined by the attention weights  $e_{ij}$  computed above to obtain the local relevance between a context and a response. For a word in the context, its relevant representation carried by the response is identified and composed using  $e_{ij}$  as

$$\bar{\mathbf{a}}_i^{att} = \sum_{j=1}^{l_b} \frac{\exp(e_{ij})}{\sum_{k=1}^{l_b} \exp(e_{ik})} \mathbf{b}_j^{att}, i \in \{1, \dots, l_a\}, \quad (6)$$

where  $\bar{\mathbf{a}}_i^{att}$  is a weighted summation of  $\{\mathbf{b}_j^{att}\}_{j=1}^{l_b}$ . Intuitively, the contents in  $\{\mathbf{b}_j^{att}\}_{j=1}^{l_b}$  that are relevant to  $\mathbf{a}_i^{att}$  are selected to form  $\bar{\mathbf{a}}_i^{att}$ . The same calculation is performed for each word in the response as

$$\bar{\mathbf{b}}_j^{att} = \sum_{i=1}^{l_a} \frac{\exp(e_{ij})}{\sum_{k=1}^{l_a} \exp(e_{kj})} \mathbf{a}_i^{att}, j \in \{1, \dots, l_b\}. \quad (7)$$

To further enhance the collected information, we compute the differences and the element-wise products between  $\{\mathbf{A}^{att}, \bar{\mathbf{A}}^{att}\}$  and between  $\{\mathbf{B}^{att}, \bar{\mathbf{B}}^{att}\}$ . The difference and element-wise product are then concatenated with the original vectors to get the enhanced representations as follows,

$$\mathbf{M}_a = [\mathbf{A}^{att}; \bar{\mathbf{A}}^{att}; \mathbf{A}^{att} - \bar{\mathbf{A}}^{att}; \mathbf{A}^{att} \odot \bar{\mathbf{A}}^{att}], \quad (8)$$

$$\mathbf{M}_b = [\mathbf{B}^{att}; \bar{\mathbf{B}}^{att}; \mathbf{B}^{att} - \bar{\mathbf{B}}^{att}; \mathbf{B}^{att} \odot \bar{\mathbf{B}}^{att}]. \quad (9)$$

Then, BiLSTMs are employed to compose the enhanced local matching information  $\mathbf{M}_a$  and  $\mathbf{M}_b$  as

$$\mathbf{v}_i^a = \text{BiLSTM}(\mathbf{M}_a, i), i \in [1, \dots, l_a], \quad (10)$$

$$\mathbf{v}_j^b = \text{BiLSTM}(\mathbf{M}_b, j), j \in [1, \dots, l_b]. \quad (11)$$

where BiLSTMs have  $d$  hidden units along each direction and  $\{\mathbf{v}_i^a, \mathbf{v}_j^b\} \in \mathbb{R}^{1 \times 2d}$ .

Instead of using max pooling and average pooling in the original ESIM model, we combine multi-dimensional pooling (Shen et al. 2017) and last-state pooling to derive the final matching feature vectors from the sequences of  $\mathbf{v}_i^a$  and  $\mathbf{v}_j^b$ .

Multi-dimensional attention differs from general attention in that the logit for an input vector is not a scalar but a

vector with dimensions equal to the dimensions of the input vector. This allows each dimension of the input vector to have a scalar logit, and we can perform attention in each dimension separately. In our model, for  $\mathbf{v}_i^a$ , its logit  $l(\mathbf{v}_i^a)$  is calculated by two linear transformations with an exponential linear units (ELU) activation function in between, i.e.,

$$l(\mathbf{v}_i^a) = \text{ELU}(\mathbf{v}_i^a \mathbf{W}_1^M + \mathbf{b}_1^M) \mathbf{W}_2^M + \mathbf{b}_2^M, \quad (12)$$

where  $\{\mathbf{W}_1^M, \mathbf{W}_2^M\} \in \mathbb{R}^{2d \times 2d}$  and  $\{\mathbf{b}_1^M, \mathbf{b}_2^M\} \in \mathbb{R}^{1 \times 2d}$ . Further, we have

$$\mathbf{d}_i^a = \text{softmax}(l(\mathbf{v}_i^a)) \odot \mathbf{v}_i^a, \quad (13)$$

$$\mathbf{d}^a = \sum_{i=1}^{l_a} \mathbf{d}_i^a. \quad (14)$$

The calculations of Eq. (12)-(14) are also applied to  $\mathbf{v}_j^b$  to get  $\mathbf{d}^b$ . Finally, we combine the multi-dimensional pooling introduced above and last-state pooling to form the matching feature vector as

$$\mathbf{f} = [\mathbf{d}^a; \mathbf{d}^b; \mathbf{v}_{l_a}^a; \mathbf{v}_{l_b}^b]. \quad (15)$$

### Prediction Layer

The matching feature vector  $\mathbf{f}$  is fed into a multi-layer perception (MLP) classifier. An MLP is a feedforward neural network estimated in a supervised way using examples of features together with known labels. Here, the MLP is designed to predict whether a pair of context and response match appropriately through the matching feature  $\mathbf{f}$ . Finally, the MLP returns a score  $s_1$  before softmax to denote the degree of matching.

### Modification Layer

At this layer, the matching score given by the prediction layer is further modified to emphasize the effect of the last utterance in the context. We denote the length of the last utterance  $\mathbf{u}$  as  $l_u$  and its output after AHRE as  $\mathbf{u}^{att}$ . A last-state pooling is employed over it to get its representation  $\mathbf{u}_{l_u}^{att}$ . A transform matrix is applied to compute another matching score  $s_2$  and the final score  $s$  is the combination of  $s_1$  and  $s_2$  with a scalar weight

$$s_2 = (\mathbf{u}_{l_u}^{att})^T \cdot \mathbf{M} \cdot \mathbf{b}_{l_b}^{att}, \quad (16)$$

$$s = s_1 + w \cdot s_2, \quad (17)$$

where  $\mathbf{M}$  and  $w$  are both parameters need to be estimated during training. Finally, a softmax layer is applied to the score  $s$  to predict the correct answer among all candidates. All model parameters are estimated in an end-to-end way by minimizing the multi-class cross-entropy loss on training set.

## Experiments

### Dataset

There were two datasets provided by the subtask 1. Both of them provided 100k training dialogues and each was equipped with 100 candidates. They are different in the development dataset size, test dataset size and vocabulary

size. Specifically, the Ubuntu dialogue has 5k development dialogues and the vocabulary size is 113k, while the Advising dialogue has only 0.5k development dialogues and the vocabulary size is only 5k.

### Training details

Adam method (Kingma and Ba 2014) was employed for optimization with a minibatch size of 2. The initial learning rate was 0.001 and was exponentially decayed by 0.96 every 5000 steps. The word embeddings were concatenations of 300-dimensional fixed GloVe embeddings (Pennington, Socher, and Manning 2014) and 100-dimensional embeddings estimated on the training set using Word2Vec (Mikolov et al. 2013) algorithm. The word embeddings were not updated during training. All hidden states of LSTM had 200 dimensions. The number of BiLSTM layers in AHRE was 3. The MLP at the prediction layer had a hidden unit size of 256 with ReLU (Nair and Hinton 2010) activation. We set the maximum context length as 160. Zeros were padded if the length was less than 160, otherwise the last 160 words were kept. We used the development dataset to select the best model for testing.

All codes were implemented using TensorFlow framework (Abadi et al. 2016) and were released to help replicate our results<sup>1</sup>.

### Evaluation metrics

Both datasets in the task were designed for selecting the best answer among a set of candidates for each given conversation. Recalls of the selected top- $k$  responses from 100 available candidates for each conversation (i.e.,  $R_{100}@k$ ) were employed as metrics to evaluate our model performance.

We also used mean reciprocal rank (MRR) to evaluate our model performance, which is a statistic measure for evaluating any process that produces a list of possible responses to a sample of queries, ordered by probability of correctness. The reciprocal rank of a query response is the multiplicative inverse of the rank of the first correct answer, and MRR is the average of the reciprocal ranks of results for a query set  $Q$ . It can be formulated as

$$MRR = \frac{1}{|Q|} \sum_{i=1}^{|Q|} \frac{1}{rank_i}, \quad (18)$$

where  $rank_i$  refers to the rank position of the first relevant document for the  $i$ -th query.

The average of  $R_{100}@10$  and MRR was adopted by the challenge organizers to get the ranks of all participants.

### Results

The results of our model on Ubuntu dataset and Advising dataset are summarized in Table 2. We tuned our single models on the development datasets and submitted the final results for subtask 1 of the track using ensemble models. The ensemble models were built by averaging the outputs of three single models with identical architectures and different random initializations.

<sup>1</sup><https://github.com/JasonForJoy/DSTC7-ResponseSelection>

Development / Test	Dataset	$R_{100}@1$	$R_{100}@10$	$R_{100}@50$	MRR
Development	Ubuntu(single)	0.521	0.817	0.982	0.616
	Ubuntu(ensemble)	0.534	0.825	0.982	0.631
	Advising(single)	0.206	0.556	0.906	0.323
	Advising(ensemble)	0.26	0.626	0.93	0.377
Test	Ubuntu(ensemble)	0.608	0.853	0.984	0.691
	Advising-Case 1(ensemble)	0.42	0.766	0.972	0.538
	Advising-Case 2(ensemble)	0.194	0.582	0.908	0.32

Table 2: Evaluation results on Ubuntu dataset and Advising dataset of subtask 1.

	$R_{100}@1$	$R_{100}@10$	$R_{100}@50$	MRR
Our model (single)	0.521	0.817	0.982	0.616
- Modification layer	0.514	0.804	0.981	0.611
- Attentive hierarchical recurrent encoder	0.506	0.799	0.977	0.602
- Multi-dimensional and last-state pooling	0.5	0.791	0.974	0.598
- Fixed word embedding	0.488	0.776	0.969	0.591

Table 3: Results of ablation tests using our single model on the Ubuntu development set of subtask 1.

It should be noticed that the test set originally released for the Advising dataset had some dependency with the training set which we denoted as *Advising-Case 1* in Table 2. The *Advising-Case 2* test set was further released to better evaluate model performance for unseen conversations and was used for system ranking.

According to the evaluation results released by challenge organizers, our proposed method ranked second on the Ubuntu dataset and third on the Advising dataset in subtask 1 of Track 1 among all 20 participants.

## Analysis

### Dataset comparison

From the evaluation results on the two different datasets shown in Table 2, we can see that there were significant recall and MRR differences between the two datasets although the same model architectures were shared. We have mentioned above that these two datasets were different in the sizes of development set, test set and vocabulary. Although the Ubuntu dataset had a much larger vocabulary, its development/test set performances were better than the Advising dataset. Meanwhile, our model showed a good generalization ability on the Ubuntu dataset because the evaluation results on test set were better than that on development set, showing its less dependency on the training set. However, the response selection performance on the Advising dataset was much worse. One possible reason is that the Advising dataset had a much small development set for model selection. Another reason is that there were some symbols such as *EECS 351* and *Classes 280* which increased the difficulty of representation and modeling.

### Ablation tests

We further investigated the effects of different parts in our proposed model by removing them one by one. A

	Layer 1	Layer 2	Layer 3
Weights	0.5324	0.2067	0.2609

Table 4: Weights of each layer in AHRE

single model built on the Ubuntu dataset was adopted for this investigation and the development set performances are as shown in Table 3. First, we can see that removing the modification layer degrades the recalls and MRR. This confirmed the positive effect of emphasizing the last utterance in the context for response selection. Second, we replaced the proposed AHRE with a simple single-layer BiLSTM at the encoding layer of our model and we can also see the performance degradation. Meanwhile, we also reported the learned weights of each layer in AHRE as shown in Table 4. Furthermore, we replaced the multi-dimensional pooling and last-state pooling at the matching layer with max pooling and average pooling employed in original ESIM. The results shown that our proposed pooling strategy was more appropriate for the response selection task. Finally, the word embedding were updated instead of being fixed during the training process, which also led to a performance degradation. Actually, the model described by the last row in Table 3 was the original ESIM. Comparing the first row and the last row in this table, we can see that significant performance improvement has been achieved by applying all our proposed techniques.

## Conclusion

In this paper, we have introduced our end-to-end model proposed for the response selection task in DSTC7. This model improves the original ESIM model from several aspects, including concatenated and fixed word representations, AHRE for sentence encoding, multi-dimensional

and last-state pooling for context-response matching, and score calculation with emphasis on the last utterance in the context. In the released evaluation results of DSTC7, our proposed method ranked second on the Ubuntu dataset and third on the Advising dataset in subtask 1 of Track 1 among all 20 participants. Ablation tests also confirm the effectiveness of our proposed methods. Our future work includes to explore the methods for other subtasks and to design a more domain-general framework that can alleviate domain-dependency of models.

## Acknowledgments

This work was partially funded by the National Nature Science Foundation of China (Grant No. U1636201).

## References

- Abadi, M.; Barham, P.; Chen, J.; Chen, Z.; Davis, A.; Dean, J.; Devin, M.; Ghemawat, S.; Irving, G.; Isard, M.; et al. 2016. Tensorflow: a system for large-scale machine learning. In *OSDI*, volume 16, 265–283.
- Bahdanau, D.; Cho, K.; and Bengio, Y. 2014. Neural machine translation by jointly learning to align and translate. *arXiv preprint arXiv:1409.0473*.
- Chen, Q.; Zhu, X.; Ling, Z.-H.; Wei, S.; Jiang, H.; and Inkpen, D. 2017. Enhanced lstm for natural language inference. In *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, volume 1, 1657–1668.
- Dong, J., and Huang, J. 2018. Enhance word representation for out-of-vocabulary on ubuntu dialogue corpus. *arXiv preprint arXiv:1802.02614*.
- Hashimoto, K.; Tsuruoka, Y.; Socher, R.; et al. 2017. A joint many-task model: Growing a neural network for multiple nlp tasks. In *Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing*, 1923–1933.
- Hochreiter, S., and Schmidhuber, J. 1997. Long short-term memory. *Neural computation* 9(8):1735–1780.
- Kingma, D. P., and Ba, J. 2014. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*.
- Kummerfeld, J. K.; Gouravajhala, S. R.; Peper, J.; Athreya, V.; Gunasekara, C.; Ganhotra, J.; Patel, S. S.; Polymenakos, L.; and Lasecki, W. S. 2018. Analyzing assumptions in conversation disentanglement research through the lens of a new dataset and model. *arXiv preprint arXiv:1810.11118*.
- Mikolov, T.; Karafiát, M.; Burget, L.; Černocký, J.; and Khudanpur, S. 2010. Recurrent neural network based language model. In *Eleventh Annual Conference of the International Speech Communication Association*.
- Mikolov, T.; Sutskever, I.; Chen, K.; Corrado, G. S.; and Dean, J. 2013. Distributed representations of words and phrases and their compositionality. In *Advances in neural information processing systems*, 3111–3119.
- Nair, V., and Hinton, G. E. 2010. Rectified linear units improve restricted boltzmann machines. In *Proceedings of the 27th international conference on machine learning (ICML-10)*, 807–814.
- Pennington, J.; Socher, R.; and Manning, C. 2014. Glove: Global vectors for word representation. In *Proceedings of the 2014 conference on empirical methods in natural language processing (EMNLP)*, 1532–1543.
- Peters, M.; Neumann, M.; Iyyer, M.; Gardner, M.; Clark, C.; Lee, K.; and Zettlemoyer, L. 2018. Deep contextualized word representations. In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long Papers)*, volume 1, 2227–2237.
- 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.
- Shang, L.; Lu, Z.; and Li, H. 2015. Neural responding machine for short-text conversation. In *Proceedings of the 53rd Annual Meeting of the Association for Computational Linguistics and the 7th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*, volume 1, 1577–1586.
- Shen, T.; Zhou, T.; Long, G.; Jiang, J.; Pan, S.; and Zhang, C. 2017. Disan: Directional self-attention network for rnn/cnn-free language understanding. *arXiv preprint arXiv:1709.04696*.
- Wang, Z.; Hamza, W.; and Florian, R. 2017. Bilateral multi-perspective matching for natural language sentences. In *Proceedings of the 26th International Joint Conference on Artificial Intelligence*, 4144–4150. AAAI Press.
- Wu, Y.; Wu, W.; Xing, C.; Zhou, M.; and Li, Z. 2017. Sequential matching network: A new architecture for multi-turn response selection in retrieval-based chatbots. In *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, volume 1, 496–505.
- Zhou, X.; Dong, D.; Wu, H.; Zhao, S.; Yu, D.; Tian, H.; Liu, X.; and Yan, R. 2016. Multi-view response selection for human-computer conversation. In *Proceedings of the 2016 Conference on Empirical Methods in Natural Language Processing*, 372–381.