

Multilingual Dialogue Generation with Shared-Private Memory

Chen Chen¹, Lisong Qiu², Zhenxin Fu², Junfei Liu³, and Rui Yan^{2*}

¹ School of Software and Microelectronics of Peking University, Beijing, China,

² Institute of Computer Science and Technology of Peking University, Beijing, China

³ National Engineering Research Center for Software Engineering (Peking University), Beijing, China

{chenchen, qiulisong, fuzhenxin, liujunfei, ruiyan}@pku.edu.cn

Abstract. Existing dialog systems are all monolingual, where features shared among different languages are rarely explored. In this paper, we introduce a novel multilingual dialogue system. Specifically, we augment the sequence to sequence framework with improved shared-private memory. The shared memory learns common features among different languages and facilitates a cross-lingual transfer to boost dialogue systems, while the private memory is owned by each separate language to capture its unique feature. Experiments conducted on Chinese and English conversation corpora of different scales show that our proposed architecture outperforms the individually learned model with the help of the other language, where the improvement is particularly distinct when the training data is limited.

Keywords: Multilingual dialogue system · Memory network · Seq2Seq · Multi-task learning.

1 Introduction

Dialogue systems have long been an interest to the community of natural language processing due to their wide range of applications. These systems can be classified as task-oriented and non-task-oriented where task-oriented dialogue systems accomplish a specific task and non-task-oriented dialogue systems are designed to chat in open domain as chatbots [1]. In particular, the sequence-to-sequence (**Seq2Seq**) framework [2], which learns to generate responses according to the given queries can achieve promising performance and grow popular [3].

Building a current state-of-the-art generation-based dialogue system requires large-scale conversational data. However, the difficulty of collecting conversational data in different languages varies greatly [4, 5]. For example, it is difficult for minority languages to collect enough dialogue corpora to build a dialogue generation model as other majority languages (e.g., English and Chinese) do. Herein, we investigate to move the frontier of dialogue generation forward from

* Corresponding author: Rui Yan (ruiyan@pku.edu.cn).

a different angle. More specifically, we find that some common features, e.g., dialogue logic, are shared in different languages but with different linguistic forms. Leveraging a multi-task framework for cross-lingual transfer learning can alleviate the problems caused by the scarcity of resources [6–8]. Through common dialogue features shared among different languages, the logic knowledge of different languages can be transferred and the robustness of the conversational model can be improved. However, to the best of our knowledge, no existing study has ever tackled multilingual generation-based dialogue systems.

This paper proposes a multi-task learning architecture for multilingual open-domain dialogue system that leverages the common dialogue features shared among different languages. Inspired by [16], we augment the `Seq2Seq` framework by adding a architecture-improved key-value memory layer between the encoder and decoder. Concretely, the memory layer consists of two parts, where the key memory is used for query addressing and the value memory stores the semantic representation of the corresponding response. To capture both shared and private features in different languages, the memory layer is further divided into shared and private memory separately. Though proposed for open-domain dialogue system, the multilingual shared-private memory architecture can be adapted flexibly and used for other tasks.

Experiments conducted on Weibo and Twitter conversational corpora of different sizes show that our proposed multilingual architecture outperforms existing techniques on both automatic and human evaluation metrics. Especially when the training data is scarce, the dialogue capability can be enhanced significantly with the help of the multilingual model.

To this end, the main contributions of our work are summarized into four folds: 1) To the best of our knowledge, the proposed work is the first to provide a solution for multilingual dialogue systems. 2) We improve the traditional key-value memory structure to expand its capacity, with which we extend the `Seq2Seq` model to capture dialogue features. 3) Based on the memory augmented dialogue model, a multi-task learning architecture with shared-private memory is proposed to achieve the transfer of dialogue features among different languages. 4) We empirically demonstrate the efficiency of multi-task learning in dialogue generation task and investigate some characteristics of this framework.

2 Related Works

2.1 Dialogue Systems

Building a dialogue system is a challenging task in natural language processing (NLP). The focus in previous decades was on template-based models [9]. However, recent generation-based dialogue systems are of growing interest due to their effectiveness and scalability. Ritter et al. [10] proposed a response generation model using statistical machine-translation methods. This idea was further developed by [11], who represented previous utterances as a context vector and incorporated the context vector into response generation. Many methods are

applied in dialogue generation. Attention helps the generation-based dialogue system by aligning the context and the response. [12] improved the performance of a recurrent neural network dialogue model via a dynamic attention mechanism. In addition, some works concentrate on many aspects of the dialogue generation, including diversity, coherence, personality, knowledgeable and controllability [13]. In these approaches, the corpora used are always in the same language. These systems are referred to as monolingual dialogue systems. As far as we know, this study is the first to explore the use of multilingual architecture to better suit the generation-based dialogue system.

2.2 Memory Networks

Memory networks [14, 15] are a class of neural network models that are augmented with external memory resources. Valuable information can be stored and reused in memory networks through the memory components. Based on the end-to-end memory network architecture [15], [16] proposed a key-value memory network architecture for question answering. The memory stores facts in a key-value structure so that the model can learn to use keys to address relevant memories with respect to the question and return corresponding values for answering. [17], [18] and [19] built goal-oriented dialogue systems based on memory-augmented neural networks. Compared with the above models, our memory components are not trained based on specific knowledge bases, but self-tuning in the training process, which makes the model more flexible. We further divide each memory module into several blocks to improve its capability.

2.3 Multi-task learning

Multi-task learning (MTL) is an approach to learn multiple related tasks simultaneously. It improves generalization by leveraging the domain-specific information contained in the training signals of related tasks [20]. [21] confirmed that NLP models benefit from the MTL approach. Many recent deep-learning approaches to multilingual issues also used MTL as part of their model.

In the context of deep learning, MTL is usually done with either hard or soft parameter sharing of hidden layers: hard parameter sharing method explicitly shares hidden layers between tasks while keeping several task-specific output layers; soft parameter sharing method usually employs regularization techniques to encourage the parameters in different tasks to be similar [22]. Hard parameter sharing is the most commonly used approach to MTL in neural networks. [7] learned a model that simultaneously translated sentences from one source language to multiple target languages. [23] propose an adversarial multi-task learning framework for text classification. [8] demonstrated a single deep learning model that jointly learned large-scale tasks from various domains including multiple translation tasks, an English parsing task, and an image captioning task. However to date, no multilingual dialogue-generation system based on multi-task learning framework has been built.

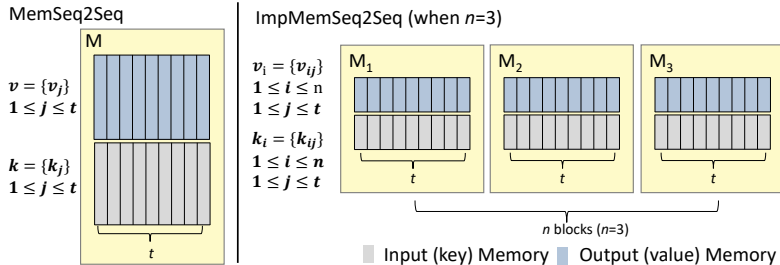


Fig. 1. Memory structure of MemSeq2Seq and ImpMemSeq2Seq. M represents memory module. k and v represent input and output memory respectively.

3 Model

In this section, we first review the vanilla Seq2Seq, then propose the key-value memory augmented Seq2Seq models, and extended them with shared-private memory components to implement the multilingual dialogue systems.

3.1 Preliminary background knowledge

A Seq2Seq model maps input sequences to output sequences. It consists of two key components: an encoder, which encodes the source input to a fix-sized context vector using the Recurrent Neural Network (RNN), and a decoder, which generates the output sequence with another RNN based on the context vector.

Given a source sequence of words (query) $q = \{x_1, x_2, \dots, x_{n_q}\}$ and a target sequence of words (response) $r = \{y_1, y_2, \dots, y_{n_r}\}$, a basic Seq2Seq based dialogue system automatically generates response r conditioned on query q by maximizing the generation probability $p(r|q)$. Specifically, the encoder encodes q to a context vector c , and the decoder generates r word by word with c as input. The objective function of Seq2Seq can be written as

$$h_t = \mathbf{f}(x_t, h_{t-1}), \quad c = h_{n_q}, \quad (1)$$

$$p(r|q) = p(y_1|c) \prod_{t=2}^{n_r} p(y_t|c, y_1, \dots, y_{t-1}), \quad (2)$$

where h_t is the hidden state at time t and f is a non-linear transformation. Moreover, gated recurrent units (GRU) and the attention mechanism proposed by [24] are used in this work.

3.2 Key-Value Memory Augmented Seq2Seq

Inspired by the end-to-end memory network[16], we introduce the MemSeq2Seq model which adds a key-value memory layer between the encoder and decoder to learn dialogue features, and the ImpMemSeq2Seq which divides the memory of MemSeq2Seq into blocks to expand model capacity.

MemSeq2Seq The MemSeq2Seq augments the Seq2Seq with a key-value memory layer between the encoder and decoder. The memory component consists of two parts: input (key) and output (value) memories. The input memory is used for query representation addressing, while the output memory stores the representation of the corresponding response information. The model retrieves information from the value memory with the weights computed as the similarity between the query representation and the key memory, with the goal of selecting values that are most relevant to the query.

Formally, we first encode a query q to a context vector c , and then calculate the similarity $p = \{p_1, \dots, p_t\}$ between c and each item of the key memory using softmax weight. Later, the model computes a new context vector c^* , which is a weighted sum of the value memory according to p .

$$p_j = \text{softmax}(c \cdot k_j), \quad (3)$$

$$c^* = \sum_{j=1}^t p_j v_j \quad (1 \leq j \leq t), \quad (4)$$

where k_j and v_j are items in the key and value memory, and t is the number of key and value items. During training, all items in memory and parameters in the Seq2Seq are jointly learned to maximize the likelihood of generating the ground-truth responses conditioned on the queries in the training set.

ImpMemSeq2Seq In the MemSeq2Seq, the key-value pairs in memory are limited, which are linear with the number of items in memory. To expand capacity, we further divide the entire memory into several individual blocks and accordingly split the input vector into several segments to compute the similarity scores. After division, similarity to multi-head attention mechanism [25], different representation subspaces at different positions are individually projected and the number of key-value pairs becomes the number of slot combinations in these blocks, while one key still corresponds to one value.

The model first split a context vector c into n segments, then compute new context segments c_i^* by memory blocks independently, and the final new context vector c^* is the concatenation of c_i^* . The formula is as follows.

$$c_1, \dots, c_n = \text{split}(c), \quad c_i^* = M_i(c_i), \quad (5)$$

$$c^* = \text{concat}(c_1^*, c_2^*, \dots, c_n^*), \quad (6)$$

where M_i represents the calculation in i^{th} memory block.

The ImpMemSeq2Seq calculates the weight p with a finer granularity, which makes the addressing more precise and flexible. Besides, with a parallel implementation, the memory layer becomes more efficient.

3.3 Seq2Seq with Shared-Private Memory

The models introduced in the previous sections can be extended for monolingual tasks. Specifically, we augment the MemSeq2Seq and ImpMemSeq2Seq for

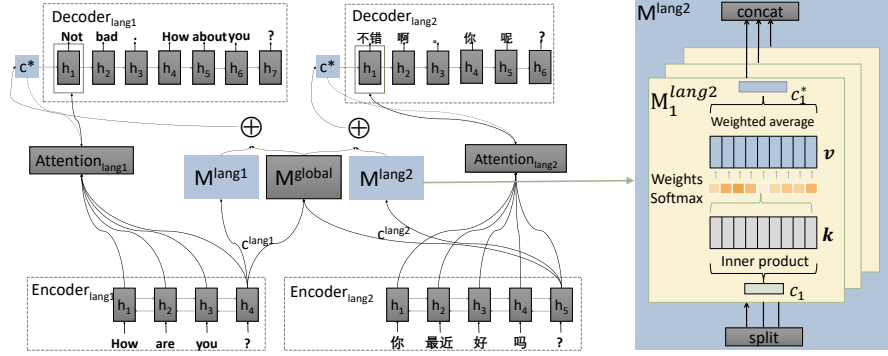


Fig. 2. SPImpMem model for multilingual dialogue system. M^{global} represents the shared memory. Superscript $lang1$, $lang2$ represent two different languages respectively.

multilingual tasks and named the extensions SPMem and SPImpMem, respectively. According to multi-task learning, dialogue systems in two different languages can be simultaneously trained. By sharing representations between two dialogue tasks, the model facilitates the cross-lingual transfer of dialogue capability.

Our multilingual model consists of four modules: an encoder, decoders, a private memory for each language and shared memory occupied by all languages. Figure 2 gives an illustration of SPImpMem. The SPMem is a special case where the number of memory blocks n is set to 1. More specifically, given a input query q , the encoder of its language first encode it into a context vector c , and then the model feeds c to both its private and shared memory. The private memory is occupied by the language corresponding to the input. The shared memory is expected to capture common features of conversations among different languages. By matching and addressing the shared and private memory components, we obtain two output vectors that are then concatenated as a new context vector c^* . The returned vector is supposed to contain features from both its own language and other languages involved in the multilingual model, which is then fed to the decoder of its language.

Given the first language conversational corpus $(\mathbf{q}_i^1, \mathbf{r}_i^1)_{i=1}^{T_1}$ and the second language conversational corpus $(\mathbf{q}_i^2, \mathbf{r}_i^2)_{i=1}^{T_2}$, the parameters Θ are learned by minimizing the negative log-likelihood between the generated $\tilde{\mathbf{r}}$ and reference \mathbf{r} , that is equivalent to maximizing the conditional probability of responses \mathbf{r}_1 and \mathbf{r}_2 given Θ , \mathbf{q}_1 and \mathbf{q}_2 :

$$J = \frac{1}{T_1} \sum_{i=1}^{T_1} \log p(\mathbf{r}_i^1 | \mathbf{q}_i^1, \Theta_{s_1}, \Theta_{M_1}, \Theta_{M_g}) + \frac{1}{T_2} \sum_{i=1}^{T_2} \log p(\mathbf{r}_i^2 | \mathbf{q}_i^2, \Theta_{s_2}, \Theta_{M_2}, \Theta_{M_g}) \quad (7)$$

where Θ_S is a collection of parameters for the encoders and decoders; Θ_M is the parameters of memory contents; T is the size of corpus; and subscripts 1, 2 and g represent $lang1$, $lang2$, and $global$ in Figure 2 respectively.

4 Experimental Settings

4.1 Datasets

We conducted experiments on open-domain single-turn Chinese (Zh) and English (En) conversational corpora. The Chinese corpus consists of 4.4 million conversations and the English corpus consists of 2.1 million conversations [10]. The conversations are scraped from Sina Weibo¹ and Twitter² respectively.

The experiments include two parts: balanced and unbalanced tests, which are discriminated by the relative size of training data for each language. In the balanced tests, the sizes of the Chinese and English corpus are comparable. We empirically set the dataset size to 100k, 400k, 1m and the whole (4.4m-Zh, 2.1m-En) to evaluate the model performance in different data scales. The unbalanced tests consist of training data of (1m-Zh, 100k-En) and (100k-Zh, 1m-En) respectively. Subsets used are sampled randomly. All the experiments have the same validation and testing data with size 10k.

4.2 Evaluation Metrics

Three different metrics are used in our experiments:

- **Word overlap based metric:** Following previous work [11], we employ BLEU [26] as an evaluation metric to measure word overlaps in a given response compared to a reference response.
- **Distinct-1 & Distinct-2:** Distinct-1 and Distinct-2 are the ratios of distinct unigrams and bigrams in generated responses respectively [27] which measure the diversity of the generated responses.
- **Three-scale human annotation:** We adopt human evaluation following [28]. Four human annotators were recruited to judge the quality of 500 generated responses from different models. All of the responses are pooled and randomly permuted. The criteria are as follows: **+2**: the response is relevant and natural; **+1**: the response is a correct reply, but contains little errors; **0**: the response is irrelevant, meaningless, or has serious grammatical errors.

4.3 Implementation Details

The Adam algorithm is adopted for optimization during training. All embeddings are set to 630-dimensional and hidden states 1024d. Considering both efficiency and memory size, we restrict both the source and target vocabulary to 60k and the batch size to 32. Chinese word segmentation is performed on Chinese conversational data. For the single block memory components, the number of cells in the memory block is set to 1024 empirically, and the dimension of each cell is adjusted according to the encoder. The memory block is further divided

¹ <http://weibo.com>

² <http://www.twitter.com>

into 32 parts in our improved memory model. In multilingual models, the number of blocks for shared and private memory component are the same. To prevent the multilingual model from favoring one certain language, we switched sentences of different languages individually by batch during training.

4.4 Comparisons

We compare our framework with the following methods:

- **Seq2Seq**. A **Seq2Seq** model with attention mechanism.
- **MemSeq2Seq**. The key-value memory augmented **Seq2Seq** model in Section 3.2.
- **ImpMemSeq2Seq**. The improved memory augmented **Seq2Seq** model with the memory block decomposed into several blocks as in Section 3.2.
- **SPMem**. The proposed multilingual model with shared-private memory which is extended from **MemSeq2Seq**. It is a special case of **SPImpMem** where the number of memory blocks n is set to 1.
- **SPImpMem**. The proposed multilingual model with shared-private memory component which is extended from **ImpMemSeq2Seq** as in Section 3.3.

5 Results and Analysis

We present the evaluation results of balanced test and unbalanced test in Table 1 and 2 respectively. Table 1 contains evaluation results of monolingual dialogue systems with **Seq2Seq**, **MemSeq2Seq** and **ImpMemSeq2Seq** as baseline. Table 2 can be viewed in conjunction with the data in Table 1.

5.1 Monolingual Models

From Table 1, we observe that the performance of the **MemSeq2Seq** model only slightly outperforms the **Seq2Seq** model. However, with memory decomposed into several parts, the **ImpMemSeq2Seq** model surpasses the basic **Seq2Seq** model. Therefore, we conclude from the comparisons that our modification of the memory components improves the capability of the model. Another observation is that in English a good conversation model can be trained with less data. Hence it does not get a significant performance gain in English as the size of data increases.

5.2 Multilingual Models

Balanced Test From the experimental results shown in Table 1, we observe that the proposed multilingual model outperforms the monolingual baselines on English corpus of different sizes. For the Chinese corpus, the promotion decreases when the size of training data increases, and thus it can only be seen on data of small sizes (i.e., 100k and 400k). Similar results can also be observed in [29]. There are several interpretations of the phenomena: 1) By the shared memory

Datasets		Monolingual (baseline)			Multilingual	
		Seq2seq	MemSeq2Seq	ImpMemSeq2Seq	SPMem	SPImpMem
100k	Zh	0.485	0.478	0.492	0.524	0.549
	En	0.779	0.780	0.825	0.831	0.805
400k	Zh	2.024	2.144	2.142	2.565	2.317
	En	1.001	0.937	0.974	0.976	1.034
1m	Zh	3.135	3.131	3.268	2.732	2.827
	En	1.027	1.100	1.099	1.135	1.146
all (4.4m-Zh, 2.1m-En)	Zh	3.600	3.383	3.755	2.955	2.765
	En	1.082	1.140	1.208	1.254	1.336

Table 1. BLEU-4 scores of the **balanced test**. The results of monolingual experiments are included for comparison.

Datasets		Multilingual	
		SPMem	SPImpMem
100k-Zh, 1m-En	Zh	1.442	1.484
	En	1.083	1.075
1m-Zh, 100k-En	Zh	2.607	2.690
	En	0.800	0.839

Table 2. BLEU-4 scores of the **unbalanced test**. Numbers in bold mean that it achieves the best performance among all models trained with this dataset.

component in the proposed multilingual model, common features are learned and transferred through both languages. Thus, when one language corpus is insufficient, some common features from other languages are helpful. 2) With the scale of corpus increasing, the monolingual model is already capable enough so that noisy information from other languages may hinder the original system.

Nevertheless, the contrary behaviors of the multilingual model on Chinese and English corpus remain suspended. As the scale of training data grows, the performance of **SPMem** and **SPImpMem** on English corpus outperforms the monolingual baselines while the performance decreases on Chinese corpus. This may result from the various qualities of different corpora which further influence the features in the shared memory blocks. The Chinese monolingual model whose parameters are originally well estimated are hindered by the noise from the shared memory. However, the English monolingual model that is relatively poorly trained benefit from the shared features. In a word, the higher quality corpus needs multilingual training less. Our model focuses on the scenario that the corpus of one language is scarce.

Unbalanced Test Since models benefit a lot from the multilingual model when training data is scarce in Table 2, we present more detailed evaluation results of models trained with the 100k datasets in Table 4 and Table 3. It is clear that, with the help of another rich resource language corpus, the multilingual model improves the performance of language with limited training data on automatic evaluation metrics except for Distinct-1. The improvements remain true even when comparing the unbalanced test results with the balanced test results, which are strengthened by the other language corpus with the same size. According to the human evaluation results, **SPMem** and **SPImpMem** generate more informative and interesting responses (+2 responses) but perform much worse

Dataset (100k-En)	BLEU-1	BLEU-2	BLEU-3	BLEU-4	Distinct-1	Distinct-2	0	+1	+2	Kappa
Seq2Seq	8.513	3.331	1.566	0.779	0.019	0.077	0.36	0.60	0.04	0.54
MemSeq2Seq	8.613	3.378	1.576	0.780	0.017	0.061	0.38	0.59	0.03	0.58
ImpMemSeq2Seq	9.178	3.623	1.682	0.825	0.018	0.070	0.37	0.58	0.05	0.47
SPMem (with 100k-zh)	8.844	3.520	1.651	0.831	0.017	0.071	0.56	0.41	0.03	0.43
SPImpMem (with 100k-zh)	9.048	3.573	1.637	0.805	0.020	0.080	0.52	0.44	0.04	0.58
SPMem (with 1m-zh)	10.669	3.686	1.580	0.800	0.007	0.124	0.31	0.62	0.07	0.46
SPImpMem (with 1m-zh)	9.118	3.648	1.701	0.839	0.025	0.103	0.52	0.33	0.15	0.51

Table 3. Evaluation results of models trained with the 100k English dataset.

Dataset(100k-Zh)	BLEU-1	BLEU-2	BLEU-3	BLEU-4	Distinct-1	Distinct-2	0	+1	+2	Kappa
Seq2Seq	9.463	3.035	1.168	0.485	0.026	0.113	0.47	0.45	0.08	0.74
MemSeq2Seq	9.210	2.859	1.101	0.478	0.018	0.073	0.43	0.43	0.14	0.74
ImpMemSeq2Seq	9.682	3.041	1.164	0.492	0.021	0.094	0.54	0.34	0.12	0.65
SPMem (with 100k-en)	10.329	3.132	1.213	0.524	0.026	0.114	0.52	0.33	0.15	0.70
SPImpMem (with 100k-en)	9.978	3.136	1.242	0.549	0.026	0.114	0.32	0.53	0.15	0.76
SPMem (with 1m-en)	11.940	4.193	2.211	1.442	0.017	0.148	0.55	0.27	0.18	0.82
SPImpMem (with 1m-en)	11.991	4.179	2.236	1.484	0.019	0.164	0.53	0.23	0.24	0.85

Table 4. Evaluation results of models trained with the 100k Chinese dataset.

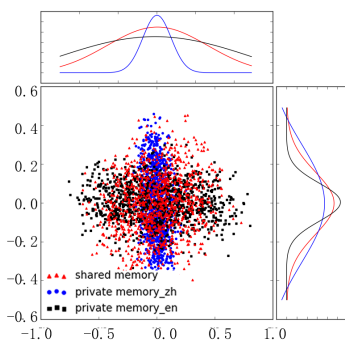


Fig. 3. The figure shows a 2-dimensional PCA projection of the input block in the memory networks. The two private memory blocks are differently oriented, and the shared memory block tends to be the mixture of them. The curves located above and right show more details of the distributions along two axes.

on +1 responses for grammatical errors. Fleiss’ Kappa on all models are larger than 0.4, which proves the correlation of the human evaluation. Therefore, some features captured by the shared memory from one language can be efficiently utilized by other languages.

5.3 Model Analysis

To illustrate the information stored in the memory components, Figure 3 visualizes the first input block of each memory, namely two private and one shared memory components. From the scatter diagram and the fitting results of the Gaussian distribution, we observe some characters in the memory layer. Tuned explicitly by each separate language, the two private memory blocks learn and store different features that appear to distribute differently in the two dimensions after principal component analysis (PCA) projection. Nevertheless, the

shared memory that is jointly updated by the two languages is likely to keep some common features of each private memory block.

6 Conclusion

This paper proposes a multi-task learning architecture with share-private memory for multilingual open-domain dialogue generation. The private memory is occupied by each separate language, and the shared memory is expected to capture and transfer common dialogue features among different languages by exploiting non-parallel corpora. To expand the capacity of vanilla memory network, the entire memory is further divided into individual blocks. Experimental results show that our model outperforms separately learned monolingual models when the training data is limited.

Acknowledgments

We thank the anonymous reviewers for their insightful comments on this paper. This work was supported by the National Key Research and Development Program of China (No. 2017YFC0804001), the National Science Foundation of China (NSFC No. 61876196 and 61672058).

References

1. Chen, H., Liu, X., Yin, D., Tang, J.: A Survey on Dialogue Systems: Recent Advances and New Frontiers. In *Sigkdd Explorations*, 19(2), pp.25-35,(2017)
2. Sutskever, I., Vinyals, O., Le, Q. V.: Sequence to Sequence Learning with Neural Networks. In: *Neural Information Processing Systems*, pp.3104-3112, (2014)
3. Serban, I. V., Sordoni, A., Bengio, Y., Courville, A. C., Pineau, J.: Building end-to-end dialogue systems using generative hierarchical neural network models. In: *National Conference on Artificial Intelligence*, 3776-3783. (2016).
4. Serban, I. V., Lowe, R., Henderson, P., Charlin, L., Pineau, J.: A Survey of Available Corpora for Building Data-Driven Dialogue Systems. *arXiv: Computation and Language* (2015).
5. Li, Y., Su, H., Shen, X., Li, W., Cao, Z., Niu, S.: DailyDialog: A Manually Labeled Multi-turn Dialogue Dataset. In: *International Joint Conference on Natural Language Processing*, pp.986-995, (2017)
6. Heigold, G., Vanhoucke, V., Senior, A. W., Nguyen, P., Ranzato, M., Devin, M., Dean, J. : Multilingual acoustic models using distributed deep neural networks. In: *International Conference on Acoustics, Speech, and Signal processing*, (2013)
7. Dong, D., Wu, H., He, W., Yu, D., Wang, H.: Multi-Task Learning for Multiple Language Translation. In: *International Joint Conference on Natural Language Processing*, (2015)
8. Kaiser, L., Gomez, A. N., Shazeer, N., Vaswani, A., Parmar, N., Jones, L., Uszkoreit, J.: One model to learn them all. *arXiv preprint arXiv:1706.05137*, (2017)
9. Wallace, R. S.: The anatomy of ALICE. In: *Parsing the Turing Test* Springer, Dordrecht, pp.181-210 (2009).

10. Ritter, A., Cherry, C., Dolan, W. B.: Data-Driven Response Generation in Social Media. In *Empirical Methods in Natural Language Processing*, pp.583-593, (2011)
11. Sordoni A, Galley M, Auli M, et al.: A Neural Network Approach to Context-Sensitive Generation of Conversational Responses. In: North American Chapter of the Association for Computational Linguistics, pp.196-205, (2015)
12. Mei, H., Bansal, M., Walter, M. R.: Coherent Dialogue with Attention-Based Language Models. In: National Conference on Artificial Intelligence, pp.3252-3258, (2017)
13. Yan, R.: Chitty-Chitty-Chat Bot: Deep Learning for Conversational AI. In: International Joint Conferences on Artificial Intelligence, (2018)
14. Weston, J., Chopra, S., Bordes, A.: Memory Networks. In: International Conference on Learning Representations, (2015)
15. Sukhbaatar, S., Szlam, A., Weston, J., Fergus, R.: End-to-end memory networks. In: *Neural Information Processing Systems*, pp.2440-2448, (2015)
16. Miller, A. H., Fisch, A., Dodge, J., Karimi, A., Bordes, A., Weston, J.: Key-Value Memory Networks for Directly Reading Documents. In: *Empirical Methods in Natural Language Processing*, pp.1400-1409, (2016)
17. Bordes, A., Boureau, Y., Weston, J.: Learning End-to-End Goal-Oriented Dialog. In: International Conference on Learning Representations, (2017)
18. Madotto, A., Wu, C., Fung, P.: Mem2Seq: Effectively Incorporating Knowledge Bases into End-to-End Task-Oriented Dialog Systems. In: meeting of the Association for Computational Linguistics, pp.1468-1478, (2018)
19. Wu, C., Madotto, A., Winata, G. I., Fung, P.: End-to-End Dynamic Query Memory Network for Entity-Value Independent Task-Oriented Dialog. In: International Conference on Acoustics, Speech, and Signal Processing, (2018)
20. Caruana, R.: Multitask learning. In: *Machine learning*, 28(1), pp.41-75, (1997)
21. Collobert, R., Weston, J.: A unified architecture for natural language processing: Deep neural networks with multitask learning. In: *Proceedings of the 25th International Conference on Machine Learning*, pp.160-167, (2008)
22. Ruder, S.: An overview of multi-task learning in deep neural networks. arXiv preprint arXiv:1706.05098, (2017)
23. Liu, P., Qiu, X., Huang, X.: Adversarial Multi-task Learning for Text Classification. In: Meeting of the Association for Computational Linguistics. pp.1-10, (2017)
24. Bahdanau, D., Cho, K., Bengio, Y.: Neural Machine Translation by Jointly Learning to Align and Translate. In: International Conference on Learning Representations, (2015)
25. Vaswani A, Shazeer N, Parmar N, et al.: Attention is All you Need. In: *Neural Information Processing Systems*. pp.5998-6008. (2017)
26. Papineni, K., Roukos, S., Ward, T., Zhu, W.: Bleu: a Method for Automatic Evaluation of Machine Translation. In: Meeting of the Association for Computational Linguistics, pp.311-318, (2002)
27. Li J, Galley M, Brockett C, et al.: A Diversity-Promoting Objective Function for Neural Conversation Models. In: North American Chapter of the Association for Computational Linguistics, pp.110-119, (2016)
28. Wu, Y., Wu, W., Li, Z., Xu, C., Yang, D.: Neural Response Generation with Dynamic Vocabularies. In: National Conference on Artificial Intelligence, pp.5594-5601, (2018)
29. Firat, O., Cho, K., Bengio, Y.: Multi-Way, Multilingual Neural Machine Translation with a Shared Attention Mechanism. In: North American Chapter of the Association for Computational Linguistics, pp.866-875, (2016)