Improving Multi-turn Response Selection Models with Complementary Last-Utterance Selection by Instance Weighting

Kun Zhou¹, Wayne Xin Zhao^{2,3(✉)}, Yutao Zhu⁴,
Ji-Rong Wen^{2,3}, and Jingsong Yu¹

¹ School of Software and Microelectronics, Peking University, Beijing, China
franciszhou@pku.edu.com, yjs@ss.pku.edu.cn

² Gaoling School of Artificial Intelligence,
Renmin University of China, Beijing, China
batmanfly@gmail.com, jrwen@ruc.edu.cn

³ Beijing Key Laboratory of Big Data Management and Analysis Methods,
Beijing, China

⁴ Université de Montréal, Montréal, Québec, Canada
yutao.zhu@umontreal.ca

Abstract. Open-domain retrieval-based dialogue systems require a considerable amount of training data to learn their parameters. However, in practice, the negative samples of training data are usually selected from an unannotated conversation data set at random. The generated training data is likely to contain noise and affect the performance of the response selection models. To address this difficulty, we consider utilizing the underlying correlation in the data resource itself to derive different kinds of supervision signals and reduce the influence of noisy data. More specially, we consider a main-complementary task pair. The main task (*i.e.*, our focus) selects the correct response given the last utterance and context, and the complementary task selects the last utterance given the response and context. The key point is that the output of the complementary task is used to set instance weights for the main task. We conduct extensive experiments in two public datasets and obtain significant improvement in both datasets. We also investigate the variant of our approach in multiple aspects, and the results have verified the effectiveness of our approach.

Keywords: Dialog system · Instance weighting · Noise reduction

1 Introduction

Recent years have witnessed remarkable progress in retrieval-based open-domain conversation systems [3, 6]. In the past few years, various methods have been proposed for response selection [1, 3, 16, 22]. A key problem in response selection is how to measure the matching degree between a conversation context and a response candidate. Many efforts have been made to construct an effective matching model with neural architectures [16, 22].

Fig. 1. The case of response and last-utterance selection model.

To construct the training data, a widely adopted approach is pairing a positive response with several randomly selected utterances as negative responses, since the labeling of true negative responses is very time-consuming. Although such method does not require labeled negative data, it is likely to bring noise during the random sampling process for negative responses. In real-world datasets, a randomly selected response is likely to be “false negative”, in which the sampled response can reply to the last-utterance but is considered as a negative response. For example, the general utterance “OK!” or “It’s great.” can safely respond to many conversations. As shown in existing studies [1, 7, 15], the noise from random sampling will severely affect the performance of the matching model.

However, we do not have any labeled data related to true negative samples. To address this difficulty, we find inspiration from the recent progress made in complementary learning [14, 17]. We design a main-complementary task pair. As shown in Fig. 1, the left side is the main task (*i.e.*, our focus) which selects the correct response given the last utterance and context, while the right side is the complementary task which selects the last utterance given the response and context. To implement such a connection, we derive a weighted margin-based optimization objective for the main task. This objective is general to work with various matching models. It elegantly utilizes different prospects in utterance selection, either last-utterance selection or response selection. The main task is assisted by the complementary task, and finally, its performance is improved.

To summarize, the major novelty lies in that the proposed approach can capture different supervision signals from different perspectives, and it is effective to reduce the influence of noisy data. The approach is general and flexible to apply to various deep matching models. We conduct extensive experiments on two public data sets, and experimental results on both data sets indicate that the models learned with our approach can significantly outperform their counterparts learned with other strategies.

2 Related Work

Recently, data-driven approaches for chatbots [3, 9] have achieved promising results. Existing work can be categorized into generation-based methods [6, 9, 11, 20] and retrieval-based methods [3, 18, 21]. The first group of approaches learn response generation from the data. Based on the sequence-to-sequence structure

with attention mechanism [11], multiple extensions have been made to tackle the “safe response” problem and generate informative responses [6, 20]. The retrieval-based methods try to find the most reasonable response from a large repository of conversational data [3, 16]. Recent work pays more attention to context-response matching for multi-turn response selection [16, 18, 22].

Instance weighting is a semi-supervised approach proposed by Grandvalet et al. [2]. The key idea is to utilize weighted margin-based optimization to train the model with a weight function to produce a reward for each instance. Then, researchers used this method to promote the model in noisy training data [8], and extended this method to other tasks [1, 4]. A recent work showed that the instance weighting strategy can be extended to different machine learning models and validated the improvement in different tasks.

Our work is inspired by the work of using new learning strategies to distinguish the noise in training data [7, 10, 15]. Shang et al. [10] and Lison et al. [7] utilized instance weighting strategy in open domain dialog systems via simple methods. Wu et al. [15] altered the negative sampling strategy and utilized a sequence-to-sequence model to distinguish false negative samples. Feng et al. [1] proposed three co-teaching mechanisms to reduce noise.

Different from aforementioned works, we utilize the last-utterance selection task as the complementary task to assist the response selection task by computing the instance weights. This complementary task is similar to the main task since it just exchanges the last utterance with the response. Our method is similar to a dual-learning approach and the difference is that the complementary model is not optimized together with the main model but only provides the instance weights to assist the main task. Besides, the two tasks own the same neural architecture, but leverage different supervision signals from the data.

3 Preliminaries

We denote a conversation as $\{u_1, \dots, u_j, \dots, u_n\}$, where each utterance u_j is a conversation sentence. A dialogue system is built to give the next utterance u_{n+1} to reply u_n . We refer to the last known utterance (*i.e.*, u_n) as *last-utterance*, and the utterance to be predicted (*i.e.*, u_{n+1}) as *response*.

We assume a training set represented by $\mathcal{D} = \{\langle U_{qi}, q_i, r_i, y_i \rangle\}_{i=1}^N$, where U_{qi} denotes the previous utterances $\{u_1, \dots, u_{n-1}\}$. q_i and r_i denote the last-utterance and response respectively. y_i is a label indicating whether r_i is an appropriate response to the entire conversation context consisting of U_{qi} and q_i .

A retrieval-based dialogue system is designed to select the correct response r from a candidate response pool $\mathcal{R}$ based on the context (namely U_q and q). This is also commonly called *multi-turn response selection task* [16, 18]. Formally, we usually solve this task by learning a matching model between last utterance and response given the context to compute the conditional probability of $\Pr(y = 1|q, r, U_q)$, which indicates the probability that r can appropriately reply to q . For simplification, we omit U_q and represent the probability by $\Pr(y = 1|q, r)$.

A commonly adopted loss for the matching model is the Cross-Entropy as:

$$L_{CE} = - \sum_{i=1}^N [y_i \cdot \log (\Pr(y_i|q_i, r_i)) + (1 - y_i) \cdot \log (1 - \Pr(y_i|q_i, r_i))]. \quad (1)$$

This is indeed a binary classification task. The optimization loss drives the probability of the positive utterance to be one and the negative utterance to be zero.

Fig. 2. The overall sketch of our approach. Our approach contains a main task (Loss Optimization Module) and a complementary task (Instance Weight Calculation Module). Last-utterance selection model M_{utte} is utilized to calculate the instance weight, while response selection model M_{res} is utilized to calculate the loss for optimization.

4 Approach

In this section, we present the proposed approach to learning matching models for multi-turn response selection. Our idea is to assign different weights to training instances, so that we can force the model to focus on confident training instances. An overall illustration of the proposed approach is shown in Fig. 2. In our approach, a general weight-enhanced margin-based optimization objective is given, where the weights indicate the reliability level of different instances. We design a complementary task that is to predict last-utterance for automatically setting these weights of training instances used in the main task.

4.1 A Pairwise Weight-Enhanced Optimization Objective

Previous methods treat all sampled responses equally, which is easily influenced by the noise in training data. To address this problem, we propose a general weighted-enhanced optimization objective. We consider a pairwise setting: each training instance consists of a positive response and a negative response for a last utterance, denoted by r^+ and r^- . For convenience, we assume each positive response is paired with a single negative sample.

The basic idea is to minimize the *Weighted Margin-based Loss* in a pairwise way, which is defined as:

$$L_{WM} = \sum_{i=0}^N w_i \cdot \max \{ \Pr(y = 1 | r_i^-, q_i) - \Pr(y = 1 | r_i^+, q_i) - \gamma, 0 \}, \quad (2)$$

where w_i is the weight for the i -instance consisting of r_i^+ and r_i^- . $\gamma \geq 0$ is a parameter to control the threshold of difference. $\Pr(y = 1 | r_i^+, q_i)$ and $\Pr(y = 1 | r_i^-, q_i)$ denote the conditional probabilities of an utterance being an appropriate and inappropriate response for q . When the probability of a negative response is larger than a positive one, we penalize it by summing the difference into the loss. This objective is general to work with various matching methods.

4.2 Instance Weighting with Last-Utterance Selection Model

A major difficulty in setting weights (shown in Eq. 1) is that there is no external supervision information. Inspired by the recent progress made in self-supervised learning and co-teaching [1, 7], we leverage supervision signals from the data itself. Since response selection aims to select a suitable response from a candidate response pool, we devise a complementary task (*i.e.*, last-utterance selection) that is trained with an assistant signal for setting the weights.

Last-Utterance Selection. Similar to response selection, here q^- can be sampled negative utterances. The complementary task captures data characteristics from a different perspective, so that the learned complementary model can be used to set weights by providing evidence on instance importance.

Instance Weighting. After learning the last-utterance selection model, we now utilize it to set weights for training instances. The basic idea is if an utterance is a proper response, it should well match the real last-utterance q^+ . On the contrary, for a true negative response, it should be uninformative to predict the last-utterance. Therefore, we introduce a new measure Δ to compute the degree that an utterance is a true positive response as:

$$\Delta_r = \Pr(y = 1 | q^+, r) - \Pr(y = 1 | q^-, r), \quad (3)$$

where $\Pr(y = 1 | q^+, r)$ and $\Pr(y = 1 | q^-, r)$ are the conditional probabilities of q^+ and q^- learned by the last-utterance selection model. In this way, a false negative response tends to yield a large Δ value, since it is able to reply to q^+ and contains useful information to discriminate between q^+ and q^- . With this measure, we introduce our solution to set the weights defined in Eq. 2. Recall that a training instance is a pair of positive and “negative” utterances, and we want to assign a weighted score indicating how much attention the response selection model should pay. Intuitively, a good training instance should be able to provide

useful information to discriminate between positive and negative responses. We define the instance weighting formula as:

$$w_i = \min \{ \max \{ \Delta_{r_i^+} - \Delta_{r_i^-} + \epsilon, 0 \}, 1 \}, \quad (4)$$

where ϵ is a parameter to adjust the mean value of weights, and we constrain the weight w_i to be less equal to 1. From this formula, we can see that a large weight w_i tends to correspond to a large $\Delta_{r_i^+}$ (a more informative positive response) and a small $\Delta_{r_i^-}$ (a less discriminative negative utterance).

4.3 Complete Learning Approach and Optimization

In this part, we present the complete learning approach.

Instantiation of the Deep Matching Models. We instantiate matching models for response selection. Our learning algorithm can work with any deep matching models. Here, we consider two recently proposed attention-based matching models, namely SMN [16] and DAM [22]. The SMN model is an RNN-based model. It first constructs semantic representations for context and response by GRU. Then, the matching features are captured by word-level and sequence-level similarity matrix. Finally a convolution neural network is adopted to distill important matching information as a matching vector and an utterance-level GRU is used to compute the matching score. The DAM model is a deep attention-based model which constructs semantic representation for context and response by a multi-layer transformer. Then, the word-level matching features are captured by cross-attention and self-attention layers. Finally a 3D-convolution is adopted to compute the matching score. These two models are selected due to their state-of-the-art performance on multi-turn response selection. Besides, previous studies have also adapted them with techniques such as weak-supervised learning [16] and co-teach learning [1].

Learning and Optimization. Given a matching model, we first pre-train it with the cross-entropy in Eq. 1. This step aims to obtain a basic model that will be further fine-tuned by our approach. For each instance consisting of a positive and a negative response, the last-utterance selection model computes the Δ value for each response by Eq. 3. Then, the weights are derived by Eq. 4 and utilized in the fine-tuning process by Eq. 2. The gradient will back-propagate to optimize the parameters in the response selection model (the gradient to last-utterance selection model is obstructed). This training approach encourages the model to focus on more confident instances with the supervision signal from the complementary task.

Discussions. In addition to the measure defined in Eq. 4, we consider using other alternatives to compute w_i , such as Jaccard similarity and embedding cosine similarity between positive and negative responses. Indeed, it is also possible to replace our multi-turn last-utterance selection model with a single-turn

last-utterance selection model to reduce the influence of the context information. Currently, we do not fine-tune the last-utterance selection model, since there is no significant improvement from this strategy in our early experiments. More details will be discussed in Sect. 5.3.

5 Experiment

In this section, we first set up the experiments, and then report the results and analysis.

5.1 Experimental Setup

Construction of the Datasets. To evaluate the performance of our approach, we use two public open-domain multi-turn conversation datasets. The first dataset is Douban Conversation Corpus (Douban) which is a multi-turn Chinese conversation data set crawled from Douban group¹. This dataset consists of one million context-response pairs for training, 50,000 pairs for validation, and 6,670 pairs for test. Another dataset is E-commerce Dialogue Corpus (ECD) [19]. It consists of real-world conversations between customers and customer service staff in Taobao². There are one million context-response pairs in the training set, and 10,000 pairs in both the validation set and the test set. For both datasets, the negative responses in the training set and the validation set are randomly sampled and the ratio of the positive and the negative is 1:1³. In the test set, each context has 10 response candidates retrieved from an index whose appropriateness regarding to the context is judged by human annotators.

Task Setting. We implement our method as Sect. 4.3. We select DAM [22] and SMN [16] as response selection models. We only select DAM [22] as our last-utterance selection model not only due to its strong feature extraction ability, but also for guaranteeing the gain only comes from the response selection model. The pre-training process follows the setting in [16, 22]. During the instance weighting, we choose 50 as the size of the mini-batches. We use Adam optimizer [5] with the learning rate as 1e-4. All gradients are clipped by 1.0 to stabilize the training process. We tune γ in $\{0, 1/8, 2/8, 3/8, 4/8\}$, and finally choose 2/8 for Douban dataset, 4/8 for ECD dataset. And we test ϵ in $\{0, 1/4, 2/4, 3/4\}$, and find 2/4 is the best choice for both datasets.

Following the works [16, 22], we use *Mean Average Presion* (MAP), *Mean Reciprocal Rank* (MRR) and *Precision at position 1* (P@1) as evaluation metrics.

¹ <https://www.douban.com/group/explore>.

² <https://www.taobao.com/>.

³ In the released training data of ECD, negative ones are automatically collected by ranking the response corpus based on conversation history augmented messages using Apache Lucene. Because retrieval negative samples from the index will bring more noisy data, we reconstruct the negative responses by random sampling from the training data. We also conduct experiments on the original training data and witness less promote than our rebuilt training data.

Baseline Models. We combine our approach with SMN and DAM to validate the effect. Besides, we compare our models with a number of baseline models: *SMN* [16] and *DAM* [22]: We utilize the pre-training results of the two models as baselines to validate the promotion of our proposed method.

Single-turn models: MV-LSTM [12] and match-LSTM [13] are the typical single-turn matching models. They concatenate all utterances in contexts as a long document for matching.

Multi-view [21]: It measures the matching degree between a context and a response candidate in both a word view and an utterance view.

DL2R [18]: It represents each utterance in contexts by RNNs and CNNs, and the matching score is computed based on the concatenation of the representations.

In addition to these baseline models, we denote the model with our proposed weighting method as Model-WM.

Table 1. Results on two datasets. Numbers marked with * indicate that the improvement is statistically significant compared with the pre-trained baseline (t-test with p-value < 0.05). We copy the numbers from [16] for the baseline models. Because the first four baselines obtain similar results in Douban dataset, we only implement two of them in ECD dataset.

Dataset	Douban			ECD		
Models	MAP	MRR	P@1	MAP	MRR	P@1
MV-LSTM	0.498	0.538	0.348	0.613	0.684	0.525
Match-LSTM	0.500	0.537	0.345	-	-	-
Multiview	0.505	0.543	0.342	-	-	-
DL2R	0.488	0.527	0.330	0.604	0.661	0.489
SMN	0.530	0.569	0.378	0.666	0.739	0.591
SMN-WM	0.550*	0.589*	0.397*	0.670	0.749*	0.612*
DAM	0.551	0.598	0.423	0.683	0.756	0.621
DAM-WM	0.584*	0.636*	0.459*	0.686	0.771*	0.647*

5.2 Results and Analysis

We present the results of all comparison methods in Table 1. First, these methods show a consistent trend on both datasets over all metrics, i.e., DAM-WM > DAM > SMN-WM > SMN > other models. We can conclude that DAM and SMN are the best baselines in this task than other models because they can capture more semantic features from word-level and sentence-level matching information. Second, our method yields improvement in SMN and DAM on two datasets, and most of these promotions are statistically significant (t-test with p-value < 0.05). This proves the effectiveness of our instance weighting method.

Third, the promotion on Douban dataset by our approach is larger than that on ECD dataset. The difference may stem from the distribution of test sets of the two data. The test set of Douban is built from random sampling, while that

of the ECD dataset is constructed by a response retrieval system. Therefore, the negative samples are more semantically similar to the positive ones. It is difficult to yield improvement by our approach with SMN and DAM in ECD dataset. Fourth, our method yields less improvement in SMN than DAM. A possible reason is that DAM fits our method better than SMN because DAM is a deep attention-based network, which owns stronger learning capacity. Another possible reason is that DAM is less sensitive to noisy training data since we have observed that the convergence process of SMN is not as stable as DAM.

Table 2. Evaluation of DAM with different weighting strategies on Douban dataset.

Method	Models	MAP	MRR	P@1
Original	DAM	0.551	0.598	0.423
Heuristic	DAM-uniform	0.577	0.623	0.433
	DAM-random	0.549	0.594	0.399
	DAM-jaccard	0.572	0.622	0.438
	DAM-embedding	0.573	0.615	0.426
Model-based	DAM-DAM	0.580	0.627	0.438
	DAM-last-WM	0.578	0.625	0.439
	DAM-dual	0.579	0.621	0.430
Ours	DAM-WM	0.584	0.636	0.459

5.3 Variations of Our Method

In this section, we explore a series of variations of our method. We replace the multi-turn last-utterance selection with other models or replace the weight produced by Eq. 4 with other heuristic methods. In this part, our experiments are conducted on Douban dataset with DAM [22] as our base model.

Heuristic Method. We consider the following methods, which change the weight produced by Eq. 4 with heuristic methods.

DAM-uniform: we fix the weight as one and follow the same procedure of our learning approach, to validate the effectiveness of our dynamic weight strategy.

DAM-random: we replace the weight model as a random function to produce random values varied in $[0,1]$.

DAM-Jaccard: we use the Jaccard similarity between positive response and negative response as the weight.

DAM-embedding [7]: we use the cosine similarity between the representation of positive and negative response as the weight. For DAM model, we calculate the average hidden state of all the words in the response as its representation.

Model-Based Method. We consider the following methods, which change the computing approach of Δ in Eq.3 by substituting our complementary model with other similar models.

DAM-last-WM replaces the multi-turn last-utterance selection model with a single-turn last-utterance selection model. This method is used to prove the effectiveness of the context information U in the last-utterance selection model. *DAM-DAM* replaces the last-utterance selection model by a response selection model. We utilize DAM model to produce $Pr(y = 1|q^+, r)$ and $Pr(y = 1|q^-, r)$. *DAM-dual* is a prime-dual approach. The response selection model is the prime model and the last-utterance selection model is the dual model. The two approaches learn instance weights for each other as Eq. 2.

Result Analysis. Table 2 reports the results of these different variations of our method on Douban dataset. First, most of these variants outperform DAM model. It demonstrates that these instance weight strategies are effective in noisy data training. Among them, DAM-WM achieves the best results for all the three evaluation metrics. It indicates that our proposed method is more effective. Second, the improvement yielded by heuristic methods is less than model-based methods. A possible reason is that neural networks own stronger semantic capacity and the weights produced by these models can better distinguish noise in training data. Third, heuristic methods achieve worse performance than DAM-uniform. It indicates that Jaccard similarity and cosine similarity of representation are not proper instance weighting functions and bring a negative effect on response selection model.

Moreover, all these model-based methods receive similar results in all three metrics and outperform DAM model. It indicates that these methods are effective but not as powerful as our proposed method. For DAM-DAM model, a possible reason is that it cannot provide more useful signal for this task than our proposed method. For DAM-last-WM, its last-utterance selection model only utilizes the last utterance therefore it cannot select positive last-utterance confidently⁴, therefore the distinguish ratio becomes noisy and low confident. For DAM-dual model, we observe that the dual-learning approach does not improve the performance of the last-utterance selection task, the reason may be that the response selection task and last-utterance selection task are not an appropriate dual-task or the dual-learning approach is not proper. We will conduct further investigation to find an appropriate dual-learning approach for this task.

5.4 Case Study

Previously, we have shown the effectiveness of our method. In this section, we qualitatively analyze why our method can yield good performance.

⁴ The last-utterance selection model of DAM-WM obtains 0.846 in P@1 metric while the one of DAM-last-WM only obtains 0.526. The distribution of positive and negative in test data is 1:9.

We calculate the weights of all the instances in training data of Douban dataset, and select the instances with maximum and minimum weight (1.0 and 0.0) respectively. We present some of them in Table 3 and annotate them manually. The first case receives a weight of 0.0, which demonstrates that the case is identified as inappropriate negative case by our last-utterance selection model. The last case receives a weight of 1.0, and we can identify the positive and negative responses. This case study shows that our instance weighting method can identify the false negative samples and punish them with less weight.

Table 3. Samples with the maximum and minimum weight learned by our approach. Green checkmarks indicate that the response candidates are proper replies of the contexts from human annotated, while red cross marks indicate inappropriate replies. The first case receives a weight of 0.0 and the negative responses can respond to the contexts to some extent. The second case receives a weight of 1.0 and the negative responses are unrelated to contexts.

Weight	0.0	1.0
1st Utterance	Girls shouldn't be too thin, so I gain weight successfully	You can make a Urban Poster
2nd Utterance	I am 1.63 meters tall and about 94 kilos, is it too thin?	Nice idea
Last Utterance	It is just in the right places	Hello, online celebrity
Pos Response	I am small boned and look thinner, so the people around me always laugh at me. (✓)	I'm not online celebrity. (✓)
Neg Response	Haha, I think so. (✓)	If you carry too many things, please think over again. (✗)

6 Conclusion and Future Work

Previous studies mainly focus on the neural architecture for multi-turn retrieval-based dialog systems, but neglect the fundamental problem from noisy training data. In this paper, we proposed a novel learning approach that was able to effectively reduce the influence of noisy data. We utilized a complementary task to learn the weights for training instances that were used by the main task. The main task was furthermore fine-tuned according to a weight-enhanced margin-based loss. Such an approach can force the model to focus on more confident training instances. Experimental results on two public datasets have demonstrated the effectiveness of our proposed method. As future work, we will design other instance weighting methods to detect noise in open domain multi-turn response selection task. Furthermore, we will consider combining our approach with more learning paradigms such as dual-learning and adversarial-learning.

Acknowledgement. This work was partially supported by the National Natural Science Foundation of China under Grant No. 61872369 and 61832017, the Fundamental

Research Funds for the Central Universities, and Beijing Outstanding Young Scientist Program under Grant No. BJJWZYJH012019100020098, and Beijing Academy of Artificial Intelligence (BAAI).

References

1. Feng, J., Tao, C., Wu, W., Feng, Y., Zhao, D., Yan, R.: Learning a matching model with co-teaching for multi-turn response selection in retrieval-based dialogue systems. [arXiv:1906.04413](#) (2019)
2. Grandvalet, Y., Bengio, Y.: Semi-supervised learning by entropy minimization. In: NIPS (2004)
3. Ji, Z., Lu, Z., Li, H.: An information retrieval approach to short text conversation. [arXiv:1408.6988](#) (2014)
4. Jiang, J., Zhai, C.: Instance weighting for domain adaptation in NLP. In: ACL (2007)
5. Kingma, D.P., Ba, J.: Adam: a method for stochastic optimization. CoRR abs/1412.6980 (2014)
6. Li, J., Galley, M., Brockett, C., Gao, J., Dolan, W.B.: A diversity-promoting objective function for neural conversation models. In: HLT-NAACL (2015)
7. Lison, P., Bibauw, S.: Not all dialogues are created equal: instance weighting for neural conversational models. In: SIGDIAL (2017)
8. Rebbapragada, U., Brodley, C.E.: Class noise mitigation through instance weighting. In: ECML (2007)
9. Serban, I., Sordoni, A., Bengio, Y., Courville, A.C., Pineau, J.: Building end-to-end dialogue systems using generative hierarchical neural network models. In: AAAI (2015)
10. Shang, M., Fu, Z., Peng, N., Feng, Y., Zhao, D., Yan, R.: Learning to converse with noisy data: generation with calibration. In: IJCAI (2018)
11. Vinyals, O., Le, Q.V.: A neural conversational model. CoRR abs/1506.05869 (2015)
12. Wan, S., Lan, Y., Xu, J., Guo, J., Pang, L., Cheng, X.: Match-SRNN: modeling the recursive matching structure with spatial RNN. [arXiv:1604.04378](#) (2016)
13. Wang, S., Jiang, J.: Learning natural language inference with LSTM. [arXiv:1512.08849](#) (2015)
14. Wang, Y., et al.: Iterative learning with open-set noisy labels. In: CVPR (2018)
15. Wu, Y., Wu, W., Li, Z., Zhou, M.: Learning matching models with weak supervision for response selection in retrieval-based chatbots. In: ACL (2018)
16. Wu, Y., Wu, W., Xing, C., Zhou, M., Li, Z.: Sequential matching network: a new architecture for multi-turn response selection in retrieval-based chatbots. In: ACL (2016)
17. Xia, Y., et al.: Dual learning for machine translation. In: NIPS (2016)
18. Yan, R., Song, Y., Wu, H.: Learning to respond with deep neural networks for retrieval-based human-computer conversation system. In: SIGIR (2016)
19. Zhang, Z., Li, J., Zhu, P., Zhao, H., Liu, G.: Modeling multi-turn conversation with deep utterance aggregation. In: COLING (2018)
20. Zhou, K., Zhang, K., Wu, Y., Liu, S., Yu, J.: Unsupervised context rewriting for open domain conversation. In: EMNLP-IJCNLP 2019, Hong Kong, China, 3–7 November 2019, pp. 1834–1844 (2019)
21. Zhou, X., et al.: Multi-view response selection for human-computer conversation. In: EMNLP (2016)
22. Zhou, X., et al.: Multi-turn response selection for chatbots with deep attention matching network. In: ACL (2018)