

Proactive Retrieval-based Chatbots based on Relevant Knowledge and Goals

Yutao Zhu¹, Jian-Yun Nie¹, Kun Zhou², Pan Du¹, Hao Jiang³, and Zhicheng Dou⁴

¹ Université de Montréal, Montréal, Québec, Canada

² School of Information, Renmin University of China, Beijing, China

³ Huawei Poisson Lab., Hangzhou, Zhejiang, China

⁴ Gaoling School of Artificial Intelligence, Renmin University of China, Beijing, China

yutao.zhu@umontreal.ca, {nie, pandu}@iro.uimontreal.ca, francis_kun_zhou@163.com

jianghao66@huawei.com, dou@ruc.edu.cn

ABSTRACT

A proactive dialogue system has the ability to proactively lead the conversation. Different from the general chatbots which only react to the user, proactive dialogue systems can be used to achieve some goals, e.g., to recommend some items to the user. Background knowledge is essential to enable smooth and natural transitions in dialogue. In this paper, we propose a new multi-task learning framework for retrieval-based knowledge-grounded proactive dialogue. To determine the relevant knowledge to be used, we frame knowledge prediction as a complementary task and use explicit signals to supervise its learning. The final response is selected according to the predicted knowledge, the goal to achieve, and the context. Experimental results show that explicit modeling of knowledge prediction and goal selection can greatly improve the final response selection. Our code is available at <https://github.com/DaoD/KPN/>.

CCS CONCEPTS

• **Computing methodologies** → **Discourse, dialogue and pragmatics**.

KEYWORDS

Proactive Dialogue; Retrieval-based Chatbot; Multi-task Learning

ACM Reference Format:

Yutao Zhu, Jian-Yun Nie, Kun Zhou, Pan Du, Hao Jiang, and Zhicheng Dou. 2021. Proactive Retrieval-based Chatbots based on Relevant Knowledge and Goals. In *Proceedings of the 44th International ACM SIGIR Conference on Research and Development in Information Retrieval (SIGIR '21)*, July 11–15, 2021, Virtual Event, Canada. ACM, New York, NY, USA, 5 pages. <https://doi.org/10.1145/3404835.3463011>

1 INTRODUCTION

From Microsoft Xiaoice, Apple Siri, to Google Assistant, dialogue systems have been widely applied in our daily life. In general, these systems are designed to make responses in reaction to the user's

requirements, such as play music, set a clock, or show the weather forecast. These systems are perceived just as tools by users as they only react passively. Users may be bored quickly. The problem is even more severe in a chit-chat style dialogue system. **Proactive conversation** offers a solution to this problem [11, 12, 14]. A proactive dialogue system can lead the dialogue proactively to achieve some goals. For example, it can drive the conversation to educate the kids about some topics, to comfort a person, to place ads unobtrusively, or to recommend items to users. Various application scenarios are emerging, yet there has not been a standard definition of what proactive conversation should be. Variations have been observed in the way that a goal is defined: by a sentence [17] or by some entities that should be covered in the conversation [12], and whether the goal should be generated dynamically [11] or pre-defined [14]. We are still in the early stage of exploration in which people test different approaches in different settings. This study is a contribution to this exploration.

In this work, we follow the setting given by Wu et al. [12], where the goal is specified by a set of entities (topics) and the background knowledge about these entities is provided in a knowledge graph. The goal is to lead the conversation smoothly to mention the required entities. The knowledge graph helps to generate paths of conversation that appear natural. Despite its simplicity, this setting has many potential applications in practice, in particular in conversational recommendation systems [4, 5, 9], where some items can be set in advance for recommendation. An example of proactive conversation in the movie domain is shown in Figure 1. The goal is defined by two entities (topics): the movie *McDull: Rise of the Rice Cooker* and the star *Bo Peng*. The system is asked to cover both entities during the conversation. By exploiting the knowledge graph, the system aims to naturally transit from one conversation topic to another and eventually fulfill the pre-defined goal.

In this work, we focus on the retrieval-based method due to its higher fluency. Although knowledge has been incorporated in some existing approaches to proactive conversation [12], it has been simply embedded in the response selection process [1, 6, 7], which is optimized globally according to the loss on the final response selection. Although the end-to-end training could be reasonable with a very large amount of training data, in practice, the limited training data may lead to sub-optimal solutions: when a wrong response is selected by the system, it is hard to tell if it is due to a poor knowledge prediction or a bad response selection, thus hard to optimize.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

SIGIR '21, July 11–15, 2021, Virtual Event, Canada

© 2021 Copyright held by the owner/author(s). Publication rights licensed to ACM.

ACM ISBN 978-1-4503-8037-9/21/07...\$15.00

<https://doi.org/10.1145/3404835.3463011>




Figure 1: An example of proactive dialogue. The system is asked to exploit the background knowledge to lead the dialogue and accomplish the goal.

To tackle this problem, we design an explicit *Knowledge Prediction* (KP) module to select the relevant piece of knowledge to use. This module is combined with a *Response Selection* (RS) module, and both form a multi-task learning framework, called **Knowledge Prediction Network** (KPN). The two tasks are jointly learned. The KP module first tracks the state of goal achievement, *i.e.*, which part of the goal has been achieved, and then leverages the dialogue context to predict which knowledge should be used in the current turn. The RS module then relies on the selected knowledge to help select the final answer. Different from the existing methods, we explicitly optimize KP using automatically generated weak-supervision signals to help better learn to predict the relevant knowledge. Experimental results show that the explicitly trained KP process can indeed select the most relevant piece of knowledge to use, and this leads to superior performance over the state-of-the-art methods.

Our main contributions are two-fold: (1) We propose a multi-task learning framework for knowledge-grounded proactive dialogue, in which the knowledge prediction task is explicitly trained in a weakly supervised manner. (2) We show experimentally that our model significantly outperforms the existing methods, demonstrating the great importance of knowledge selection in proactive conversation.

2 KNOWLEDGE PREDICTION NETWORK

Problem Formalization We follow the task definition formulated by Wu et al. [12]. For a dataset \mathcal{D} , each sample is represented as (c, g, k, r, y) (as shown in Figure 1), where $c = \{u_1, \dots, u_L\}$ represents a conversation context with $\{u_i\}_{i=1}^L$ as utterances; g represents the goal containing some entities that the dialogue should talk about (e.g., “Bo Peng”); $k = (k_1, \dots, k_M)$ are knowledge triplets where k_i is in form of SPO (Subject, Predicate, Object); r is a response candidate; $y \in \{0, 1\}$ is a binary label. The task is to learn a matching model $s(c, g, k, r)$ with \mathcal{D} to measure the suitability of a response candidate r .

In this work, we propose a multi-task learning framework KPN that contains *response selection* (RS) and *knowledge prediction* (KP) as two distinct tasks, as illustrated in Figure 2. The predicted knowledge and updated goal from the KP task are used as input to the RS task. The loss functions in the two tasks are combined for training the model jointly. Different from the existing work that fuses the two tasks together and trains the whole model by only the final RS loss (\mathcal{L}_{rs}), we propose using a KP loss (\mathcal{L}_{kp}) to supervise the

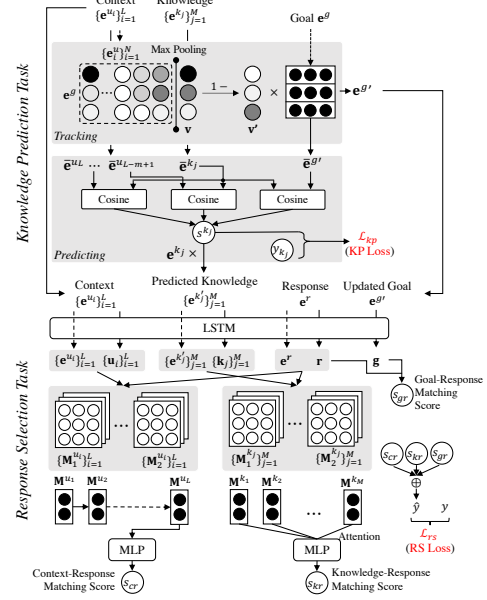


Figure 2: The structure of KPN. The predicted knowledge e^{k_i} and updated goal $e^{g'}$ in the knowledge prediction task will be used as input to the response selection task.

knowledge prediction process directly. The overall loss is as follows:

$$\mathcal{L} = \lambda \mathcal{L}_{kp} + \mathcal{L}_{rs}, \quad (1)$$

where λ is a hyperparameter (set as 0.3 in our experiment) to control the influence of the KP loss. The joint learning process allows us to better tell if a wrong response is obtained due to a wrong prediction of knowledge or a wrong selection of response. Details of the two tasks are presented in Sections 2.1 and 2.2.

The processes of KP and RS are based on the following basic representations: an utterance u_i in the context, a goal g , a knowledge triplet k_j (concatenated as a word sequence), and a response r are first represented as matrices e^u , e^g , e^k , and e^r respectively through a pre-trained embedding table. They will be used in different ways in the KP and RS processes.

2.1 Knowledge Prediction (KP) Task

It is widely believed that knowledge can help select suitable responses. However, not all knowledge triplets are useful in selecting responses for a conversation turn. Therefore, predicting whether a knowledge triplet should be used is a critical step.

Goal Tracking To decide what to say in a response, one has to know what part of the goal is still uncovered. KPN achieves this by a goal tracking process (shown in the *Tracking* part of Figure 2). The basic idea is to match the goal and the context, then the mismatched entities are considered as uncovered. Concretely, we concatenate all utterances in the context as a long sequence $\{e_i^u\}_{i=1}^N$, where N is the total number of words in all the utterances, and then match it with the goal (e^g) by cosine similarity: $m_{ij} = \cos(e_i^u, e_j^g)$. Then max-pooling is applied to extract the strongest matching signals: $v_i = \text{ReLU}(\text{Maxpooling}(m_{i,:}))$. The obtained values (v) represent the degree of coverage of the entities in the goal, while $v' = 1 - v$ represents the remaining part that should be covered

in the following dialogue. Finally, the vector \mathbf{v}' is used to update the representation of the goal: $\mathbf{e}^{g'} = \mathbf{v}' \cdot \mathbf{e}^g$. This goal tracking method is simple but effective, and more sophisticated designs can be investigated as future work.

Knowledge Predicting The knowledge prediction process is shown in the *Predicting* part of Figure 2. The relevance of a piece of knowledge is determined by both the state of the goal and the current topic of the dialogue. The former determines the target, while the latter determines the starting point. Ideally, the relevant knowledge should pave a way leading from the current topic to the desired goal. Usually, the current topic is contained in the last several utterances, thus we leverage them to predict the relevant knowledge. Given the updated goal $\mathbf{e}^{g'}$, the last m utterances $\{\mathbf{e}^{u_i}\}_{i=L-m+1}^L$ (where L is the number of utterances in the context, and m is a hyperparameter set as 3 in our experiments), and the j -th piece of knowledge \mathbf{e}^{k_j} , we first compute their sentence-level representations by mean-pooling over word dimensions: $\bar{\mathbf{e}}^{u_i} = \text{mean}(\mathbf{e}^{u_i})$, $\bar{\mathbf{e}}^{g'} = \text{mean}(\mathbf{e}^{g'})$, and $\bar{\mathbf{e}}^{k_j} = \text{mean}(\mathbf{e}^{k_j})$. Then we use cosine similarity $s_{g'}^{k_j} = \cos(\bar{\mathbf{e}}^{g'}, \bar{\mathbf{e}}^{k_j})$, $s_i^{k_j} = \cos(\bar{\mathbf{e}}^{u_i}, \bar{\mathbf{e}}^{k_j})$ to measure their relevance, where $i \in [L - m + 1, L]$, and we obtain $(m + 1)$ scores $[s_{g'}^{k_j}, s_{L-m+1}^{k_j}, \dots, s_L^{k_j}]$. The relevance scores with respect to both the goal and the context topic are aggregated by a multi-layer perceptron (MLP) with a sigmoid activation function (σ), which is then used to update the representation of the j -th knowledge triplet:

$$s^{k_j} = \sigma(\text{MLP}([s_0^{k_j}, s_{L-m+1}^{k_j}, \dots, s_L^{k_j}])), \quad \mathbf{e}^{k_j} = s^{k_j} \mathbf{e}^{k_j}, \quad (2)$$

where s^{k_j} is the predicted probability of the j -th knowledge triplet to be used in the current turn.

Weakly Supervised Knowledge Prediction To make a correct prediction of knowledge, the common method is tuning the knowledge prediction process according to the final response selection error. The process is thus implicitly supervised [6, 7, 12]. To further improve the learning of the knowledge prediction, besides the response selection loss, we introduce a weakly supervised knowledge prediction loss to train it explicitly.

In practice, it is difficult to have manual labels for knowledge triplets in each dialogue turn. To address this problem, we propose a method to generate weak labels automatically. For each knowledge SPO triplet, we adopt an entity linking method to link it to the response: if the *object entity* appears in the ground-truth response, we label it as 1, otherwise as 0¹. We assume this weak label can indicate whether such a piece of knowledge is used in the ground-truth response. With the weak labels y_{k_j} , we can compute a binary cross-entropy loss, which we call KP loss, as follows:

$$\mathcal{L}_{kp} = -\frac{1}{|\mathcal{D}|} \sum (y_{k_j} \log s^{k_j} + (1 - y_{k_j}) \log(1 - s^{k_j})). \quad (3)$$

2.2 Response Selection (RS) Task

Response selection (RS) is the main task. As shown in Figure 2, KPN considers the interactions between response and three types of information, *i.e.*, the context, the knowledge, and the remaining goal. The former two can be modeled in the same way: similar to

¹For long descriptive entities (*i.e.*, non-factoid sentences such as the *Comment* entity about *Bo Peng* in the Knowledge Graph in Figure 1), if more than 70% part is covered by the ground-truth response, we label it as one. We do not use the subject entity (*e.g.*, “Bo Peng”), because it is shared by many triplets, thus is less accurate as the label.

existing work [3, 13, 16], we compute matching matrices based on both the input representations (\mathbf{e}^{u_i} , \mathbf{e}^{k_j} and \mathbf{e}^r) and their sequential representations obtained by LSTM [2]. As a result, we denote the obtained matrices as \mathbf{M}^u and \mathbf{M}^k and apply a CNN with max-pooling to extract the matching features \mathbf{v}^u and \mathbf{v}^k .

(1) **Context-Response Matching** The matching features between the context and response are aggregated by an LSTM and the corresponding final state is fed into an MLP to compute the matching score s_{cr} . We use LSTM because it can model the dependency and the temporal relationship of utterances in the context.

(2) **Knowledge-Response Matching** Different from the context, we assume knowledge triplets to be independent. Thus, we use an attention-based method to aggregate the matching features:

$$\alpha_i = \text{ReLU}(\text{MLP}(\mathbf{v}_i^k)), \quad \mathbf{h}_2 = \sum_{i=1}^{k_M} \frac{e^{\alpha_i}}{\sum_j e^{\alpha_j}} \mathbf{v}_i^k, \quad s_{kr} = \text{MLP}(\mathbf{h}_2). \quad (4)$$

This way, a knowledge triplet that is more related to the response will have a higher weight in the aggregated features and contributes more in computing the final matching score.

(3) **Goal-Response Matching** As the goal is a single sequence of tokens, which is much easier to model, we compute the goal-response matching score s_{gr} by an MLP based on their LSTM representations at the last time step.

The final matching score is then computed as: $\hat{y} = (s_{cr} + s_{kr} + s_{gr})/3$. We use the binary cross-entropy loss to compute the errors:

$$\mathcal{L}_{rs} = -\frac{1}{|\mathcal{D}|} \sum (y \log \sigma(\hat{y}) - (1 - y) \log(1 - \sigma(\hat{y}))). \quad (5)$$

3 EXPERIMENTS

3.1 Datasets and Baseline Models

We experiment on datasets DuConv and DuRecDial. **DuConv** [12] is built for knowledge-grounded proactive human-machine conversation. The dialogues are about movies and stars. The total number of training, validation, and test samples is 898,970, 90,540, and 50,000. **DuRecDial** [8] is created as a conversational recommendation dataset, which contains dialogues between a seeker and a recommender. The domain of dialogue includes movie, music, food, etc. The number of training, validation, and test samples is 342,340, 38,060, and 55,270. The negative responses are randomly sampled with a 1:9 positive/negative ratio in both datasets.

We compare our model against two groups of baseline methods: **DuRetrieval** [12] is the only retrieval-based model specifically designed for proactive dialogue. It uses a Transformer-based encoder for context and response representation. The conversation goal is used as an additional piece of knowledge. All knowledge triplets are represented by a bi-GRU with attention mechanism.

The other group of methods are not proposed for proactive dialogue but for general knowledge-grounded dialogue. As they also incorporate knowledge into dialogue generation, we replace our knowledge selecting module in the KP task by theirs to make a comparison. **MemNet** [1] uses a memory network that performs “read” and “write” on the knowledge by matrix multiplication. **PostKS** [6] trains a knowledge prediction process to make the prior probability (using only the context) of the knowledge prediction close to the posterior probability (using both context and response). **NKD** [7] is

similar to MemNet, but it first concatenates the context and knowledge representations and then uses an MLP to compute the weight for each piece of knowledge.

3.2 Evaluation

All models are evaluated in two scenarios.

On test set Similar to the existing work [12, 15], we evaluate the performance of each model by **Hits@1**, **Hits@3**, and Mean Reciprocal Rank (**MRR**) for selecting the correct response when it is mixed up with several other candidates. Hits@ k measures the ratio of the ground-truth response among the top k results.

Practical application Following [12], we also evaluate the performance of the models in a more practical scenario, where each ground-truth utterance is mixed up with 49 utterances retrieved by Solr². The task is to rank the ground-truth response as high as possible. This test simulates a practical scenario where the model is acting as a reranker for the candidate list returned by an upstream retrieval system. We use several metrics to evaluate the model from different perspectives. BLEUs are used to measure the quality (similarity) of the response w.r.t. the ground-truth. To evaluate the model’s ability to incorporate knowledge into dialogues, we compute the **knowledge precision/recall/F1** score used in previous studies [6, 10, 12], which measure how much knowledge (either correct or wrong) has been used in the responses. We also compute a more meaningful **knowledge accuracy** to measure if the selected response uses the same piece of knowledge as that involved in the ground-truth response. Similarly, **goal accuracy** measures if a goal in the ground-truth is correctly covered by the selected response.

3.3 Experimental Results

The evaluation results are shown in Table 1. Based on the results, we can observe: (1) KPN outperforms all baselines significantly by achieving the highest scores on all evaluation metrics. (2) Compared with DuRetrieval, KPN improves Hits@1, Hits@3, and MRR by a large margin. This strongly indicates that KPN has a better capability of selecting correct responses. (3) In the practical application scenario, according to the results on BLEU, we can conclude that KPN can select responses that are more similar to the golden responses. (4) On knowledge prediction, as a comparison, we also provide the evaluation result of the ground-truth. We find that our method outperforms other knowledge prediction models (MemNet, PostKS, and NKD) on knowledge P/R/F1 and accuracy. This demonstrates that the explicit supervised knowledge prediction is more effective than the implicit ones used in the other methods. Nevertheless, there is still a big gap between our results and the ground-truth, showing that the process could be much improved. **Reliability of the Weak Labels** As we use an entity linking method to automatically generate weak labels for knowledge prediction, to evaluate the reliability of these labels, we randomly select 100 samples comprising 1,437 knowledge triplets from the validation set of DuConv, and ask three human annotators to label which triplet is necessary to select the current response. The result indicates that 90.26% of the generated labels are consistent with

²<https://lucene.apache.org/solr/>. If the number of retrieved results is less than 49, we use random samples to pad.

Table 1: Evaluation results. “KLG” stands for knowledge, and “Acc” for accuracy. “GT” for ground-truth. “+X” means that the knowledge prediction module is replaced by X. The improvement obtained by KPN over DuRetr. is statistically significant with p -value < 0.01 in t-test.

		GT	DuRetr.	KPN			
				Full	+MemNet	+PostKS	+NKD
DuConv	Hits@1	-	50.12	66.94	52.54	39.98	56.42
	Hits@3	-	75.68	87.52	78.79	65.70	81.54
	MRR	-	63.13	78.30	67.90	57.09	70.77
	BLEU1	1.00	0.47	0.56	0.50	0.48	0.50
	BLEU2	1.00	0.32	0.42	0.34	0.33	0.35
	KLG. P	38.24	30.11	33.45	29.24	28.55	29.40
	KLG. R	9.20	7.24	8.05	7.03	6.87	7.07
	KLG. F1	14.83	11.68	12.97	11.34	11.07	11.40
	KLG. Acc.	100.00	53.64	57.82	50.90	50.42	52.94
	Goal Acc.	100.00	58.90	77.58	72.36	69.44	74.62
DuRecDial	Hits@1	-	77.38	91.50	75.34	82.45	82.74
	Hits@3	-	89.02	98.86	93.92	96.60	97.03
	MRR	-	84.07	95.18	85.00	89.58	89.96
	BLEU1	1.00	0.46	0.61	0.51	0.53	0.53
	BLEU2	1.00	0.39	0.51	0.39	0.41	0.41
	KLG. P	52.64	43.42	52.55	41.04	43.70	42.87
	KLG. R	3.76	3.10	3.76	2.93	3.12	3.07
	KLG. F1	7.02	5.79	7.01	5.48	5.83	5.72
	KLG. Acc.	100.00	94.90	95.35	94.32	94.90	94.81
	Goal Acc.	100.00	78.34	84.96	82.58	83.12	83.93

human annotations³. This demonstrates the high reliability of the labels automatically generated by our entity linking method.

We carried out detailed **Ablation Study** and **Influence of Hyperparameter**, showing that both the goal and knowledge strongly impact the final results. Due to space limit, these experiments are presented in our Github page.

4 CONCLUSION

In this paper, we proposed a new approach to retrieval-based proactive dialogue. In our model, we define two tasks for response selection and knowledge prediction. An interactive matching structure is applied to model the matching between the knowledge and the response. In order to make a good prediction of knowledge, explicit supervision signals are used, which are derived from the ground-truth responses. Experimental results demonstrated that our model can achieve better performance than the baseline models in which the two tasks are mixed up. In particular, it is shown that training the knowledge prediction explicitly is very effective. This work is a first demonstration of the importance of modeling knowledge and goals explicitly in proactive dialogue.

ACKNOWLEDGMENTS

We thank Wenquan Wu and Zhen Guo for the insightful suggestions. This work was supported by a Discovery grant of the Natural Science and Engineering Research Council of Canada, National Natural Science Foundation of China (No. 61872370 and No. 61832017), Beijing Outstanding Young Scientist Program (NO. BJJWZYJH012019100020098), and Shandong Provincial Natural Science Foundation (Grant ZR2019ZD06).

³The Fleiss Kappa is 0.698 that indicates the annotators achieve a substantial agreement

REFERENCES

- [1] Marjan Ghazvininejad, Chris Brockett, Ming-Wei Chang, Bill Dolan, Jianfeng Gao, Wen-tau Yih, and Michel Galley. 2018. A Knowledge-Grounded Neural Conversation Model. In *Proceedings of the Thirty-Second AAAI Conference on Artificial Intelligence, (AAAI-18), the 30th innovative Applications of Artificial Intelligence (IAAI-18), and the 8th AAAI Symposium on Educational Advances in Artificial Intelligence (EAAI-18)*, New Orleans, Louisiana, USA, February 2-7, 2018, Sheila A. McIlraith and Kilian Q. Weinberger (Eds.). AAAI Press, 5110–5117.
- [2] Sepp Hochreiter and Jürgen Schmidhuber. 1997. Long Short-Term Memory. *Neural Comput.* 9, 8 (1997), 1735–1780. <https://doi.org/10.1162/neco.1997.9.8.1735>
- [3] Kai Hua, Zhiyuan Feng, Chongyang Tao, Rui Yan, and Lu Zhang. 2020. Learning to Detect Relevant Contexts and Knowledge for Response Selection in Retrieval-based Dialogue Systems. In *CIKM '20: The 29th ACM International Conference on Information and Knowledge Management, Virtual Event, Ireland, October 19-23, 2020*, Mathieu d'Aquin, Stefan Dietze, Claudia Hauff, Edward Curry, and Philippe Cudré-Mauroux (Eds.). ACM, 525–534. <https://doi.org/10.1145/3340531.3411967>
- [4] Dongyeop Kang, Anusha Balakrishnan, Pararth Shah, Paul A. Crook, Y-Lan Boureau, and Jason Weston. 2019. Recommendation as a Communication Game: Self-Supervised Bot-Play for Goal-oriented Dialogue. In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing, EMNLP-IJCNLP 2019, Hong Kong, China, November 3-7, 2019*, Kentaro Inui, Jing Jiang, Vincent Ng, and Xiaojun Wan (Eds.). Association for Computational Linguistics, 1951–1961. <https://doi.org/10.18653/v1/D19-1203>
- [5] Raymond Li, Samira Ebrahimi Kahou, Hannes Schulz, Vincent Michalski, Laurent Charlin, and Chris Pal. 2018. Towards Deep Conversational Recommendations. In *Advances in Neural Information Processing Systems 31: Annual Conference on Neural Information Processing Systems 2018, NeurIPS 2018, December 3-8, 2018, Montréal, Canada*, Samy Bengio, Hanna M. Wallach, Hugo Larochelle, Kristen Grauman, Nicolò Cesa-Bianchi, and Roman Garnett (Eds.). 9748–9758.
- [6] Rongzhong Lian, Min Xie, Fan Wang, Jinhua Peng, and Hua Wu. 2019. Learning to Select Knowledge for Response Generation in Dialog Systems. In *Proceedings of the Twenty-Eighth International Joint Conference on Artificial Intelligence, IJCAI 2019, Macao, China, August 10-16, 2019*, Sarit Kraus (Ed.). ijcai.org, 5081–5087. <https://doi.org/10.24963/ijcai.2019/706>
- [7] Shuman Liu, Hongshen Chen, Zhaochun Ren, Yang Feng, Qun Liu, and Dawei Yin. 2018. Knowledge Diffusion for Neural Dialogue Generation. In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics, ACL 2018, Melbourne, Australia, July 15-20, 2018, Volume 1: Long Papers*, Iryna Gurevych and Yusuke Miyao (Eds.). Association for Computational Linguistics, 1489–1498. <https://doi.org/10.18653/v1/P18-1138>
- [8] Zeming Liu, Haifeng Wang, Zheng-Yu Niu, Hua Wu, Wanxiang Che, and Ting Liu. 2020. Towards Conversational Recommendation over Multi-Type Dialogs. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics, ACL 2020, Online, July 5-10, 2020*, Dan Jurafsky, Joyce Chai, Natalie Schluter, and Joel R. Tetreault (Eds.). Association for Computational Linguistics, 1036–1049. <https://doi.org/10.18653/v1/2020.acl-main.98>
- [9] Jinghui Qin, Zheng Ye, Jianheng Tang, and Xiaodan Liang. 2020. Dynamic Knowledge Routing Network for Target-Guided Open-Domain Conversation. In *The Thirty-Fourth AAAI Conference on Artificial Intelligence, AAAI 2020, The Thirty-Second Innovative Applications of Artificial Intelligence Conference, IAAI 2020, The Tenth AAAI Symposium on Educational Advances in Artificial Intelligence, EAAI 2020, New York, NY, USA, February 7-12, 2020*. AAAI Press, 8657–8664.
- [10] Lianhui Qin, Michel Galley, Chris Brockett, Xiaodong Liu, Xiang Gao, Bill Dolan, Yejin Choi, and Jianfeng Gao. 2019. Conversing by Reading: Contentful Neural Conversation with On-demand Machine Reading. In *Proceedings of the 57th Conference of the Association for Computational Linguistics, ACL 2019, Florence, Italy, July 28- August 2, 2019, Volume 1: Long Papers*, Anna Korhonen, David R. Traum, and Lluís Màrquez (Eds.). Association for Computational Linguistics, 5427–5436. <https://doi.org/10.18653/v1/p19-1539>
- [11] Jianheng Tang, Tiancheng Zhao, Chenyan Xiong, Xiaodan Liang, Eric P. Xing, and Zhiting Hu. 2019. Target-Guided Open-Domain Conversation. In *Proceedings of the 57th Conference of the Association for Computational Linguistics, ACL 2019, Florence, Italy, July 28- August 2, 2019, Volume 1: Long Papers*, Anna Korhonen, David R. Traum, and Lluís Màrquez (Eds.). Association for Computational Linguistics, 5624–5634. <https://doi.org/10.18653/v1/p19-1565>
- [12] Wenquan Wu, Zhen Guo, Xiangyang Zhou, Hua Wu, Xiyuan Zhang, Rongzhong Lian, and Haifeng Wang. 2019. Proactive Human-Machine Conversation with Explicit Conversation Goal. In *Proceedings of the 57th Conference of the Association for Computational Linguistics, ACL 2019, Florence, Italy, July 28- August 2, 2019, Volume 1: Long Papers*, Anna Korhonen, David R. Traum, and Lluís Màrquez (Eds.). Association for Computational Linguistics, 3794–3804. <https://doi.org/10.18653/v1/p19-1369>
- [13] Chunyuan Yuan, Wei Zhou, Mingming Li, Shangwen Lv, Fuqing Zhu, Jizhong Han, and Songlin Hu. 2019. Multi-hop Selector Network for Multi-turn Response Selection in Retrieval-based Chatbots. In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing, EMNLP-IJCNLP 2019, Hong Kong, China, November 3-7, 2019*, Kentaro Inui, Jing Jiang, Vincent Ng, and Xiaojun Wan (Eds.). Association for Computational Linguistics, 111–120. <https://doi.org/10.18653/v1/D19-1011>
- [14] Hao Yuan and Jinqi An. 2020. Multi-Hop Memory Network with Graph Neural Networks Encoding for Proactive Dialogue. In *ICCAI '20: 2020 6th International Conference on Computing and Artificial Intelligence, Tianjin, China, April 23-26, 2020*. ACM, 24–29. <https://doi.org/10.1145/3404555.3404605>
- [15] Saizheng Zhang, Emily Dinan, Jack Urbanek, Arthur Szlam, Douwe Kiela, and Jason Weston. 2018. Personalizing Dialogue Agents: I have a dog, do you have pets too?. In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics, ACL 2018, Melbourne, Australia, July 15-20, 2018, Volume 1: Long Papers*, Iryna Gurevych and Yusuke Miyao (Eds.). Association for Computational Linguistics, 2204–2213. <https://doi.org/10.18653/v1/P18-1205>
- [16] Yutao Zhu, Jian-Yun Nie, Kun Zhou, Pan Du, and Zhicheng Dou. 2021. Content Selection Network for Document-Grounded Retrieval-Based Chatbots. In *Advances in Information Retrieval - 43rd European Conference on IR Research, ECIR 2021, Virtual Event, March 28 - April 1, 2021, Proceedings, Part I (Lecture Notes in Computer Science, Vol. 12656)*, Djoerd Hiemstra, Marie-Francine Moens, Josiane Mothe, Raffaele Perego, Martin Potthast, and Fabrizio Sebastiani (Eds.). Springer, 755–769. https://doi.org/10.1007/978-3-030-72113-8_50
- [17] Yutao Zhu, Ruihua Song, Zhicheng Dou, Jian-Yun Nie, and Jin Zhou. 2020. ScriptWriter: Narrative-Guided Script Generation. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics, ACL 2020, Online, July 5-10, 2020*, Dan Jurafsky, Joyce Chai, Natalie Schluter, and Joel R. Tetreault (Eds.). Association for Computational Linguistics, 8647–8657. <https://doi.org/10.18653/v1/2020.acl-main.765>