Team PAssion at SereTOD-EMNLP 2022: End-to-End Task-Oriented Dialog System with Improved Prompting Scheme

Weichen Lu, Yifei Wang, Weizhen Zhao, Buxian Chen, Xiaojie Chen, Jinsong Pan, Wentao Liang, Yongquan Lai
Ping An Property & Casualty Insurance Company, Shenzhen, China

Abstract

We present the system description for our submission to the SereTOD-EMNLP 2022 competition. The task of Track 2 is to develop a task-oriented dialog system that can predict the user intent, the system intent, and the system response. We choose the official baseline model as a starting point, which is a pre-trained language model (PLM) based method that solves the problem in an end-to-end manner. Then we focus on finding a good prompting scheme for the PLM, since the prompting scheme is crucial for the PLM-based method. Motivated by making good use of the local knowledge base (KB), an effective prompting scheme that explicitly describes the connection between the user goal and the local KB is proposed. Moreover, the unlabeled dialogues are used to fine-tune the PLM before training with the labeled dialogues. As a result, our method ranks second in the competition with a final score of 2.7.

1 Introduction

The goal of the task-oriented dialog system in this competition(Ou et al., 2022)(Liu et al., 2022b) is to solve the user's requirements through multiple rounds of dialog. With the development of large pre-trained language model (PLM), many methods based on large PLM have been proposed for task-oriented dialog systems, such as Simple TOD (Hosseini-Asl et al., 2020), UBAR(Yang et al., 2021), MTTOD(Lee, 2021), and MGA(Liu et al., 2022a). These methods take the historical utterances, dialog state, current utterance, and knowledge base as input to the PLM, and then decode the dialog state, system action, and response step by step. We can see that the main improvement of these PLM based methods is on the input side. The input of SimpleTOD consists of the historical utterances of the user and the system. UBAR adds dialog states and system actions to the input. MGA proposes to input only the previous round of dialog state, the system response, and the current round of user utterance, as this is based on the dialog state definition, which should include all historical information, thus saving memory and computational resources. Motivated by the good use of local KB, an effective prompting scheme is proposed to more explicitly construct the relationship between the user's potential requests and local KB information. In addition, more historical utterances are included in the input. Note that the history of the service response is excluded to avoid error propagation.

2 Method

2.1 Architecture

As shown in Fig. 1, we build the task-oriented dialog system based on large pre-trained language model in an end-to-end manner. In the beginning, by inputting the history of user utterances, the history of entity names, the current user utterance with constraints, and the current entity names, the language model predicts the user intent. Then, the information extracted from the local KB is added. Finally, the system's intent and response are predicted progressively.

2.2 Leverage unlabeled data.

There is about nine times more unlabeled data than labeled data, so it is intuitive to think about using the unlabeled data. The most intuitive method for dealing with the unlabeled data is to fine-tune the large language model on the unlabeled data. By observing the data, we found that the speaker with the smaller ID usually represents the system, and based on this regularity, we unified the unlabeled data so that the user spoke first. Then we fine-tune the large language model based on the unlabeled data.

2.3 Effective prompting scheme.

In this section, the effective prompting scheme is presented. As shown in Fig. 1, we first add the

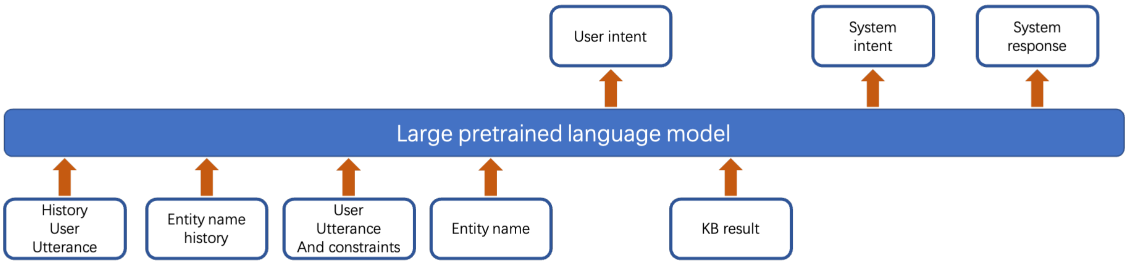

例：十八的[EOS_U]嗯喂就是，我上次叫你们给我改的套餐是十八块钱的怎么嗯又变成二十八的了[EOS_U]二十八的，十八的，套餐[EOS_L]嗯，那我查话费我查的是二十八的说我的四g套餐是%(二十八的-业务费用)(二十八的套餐-业务费用)(四g套餐-业务费用)(十八的-业务费用)(套餐-业务费用)[EOS_U]二十八的，十八的，套餐[EOS_E]提供信息[EOS_UI](十八的，套餐-流量总量-一百兆)(十八的，套餐-业务费用-十八块钱，十八)(二十八的-业务费用-二十八)[EOS_K]通知[EOS_SI]二十八，您的这个是十八的啊，没问题已经改好了[EOS_S]

Figure 1: Model architecture

user's historical utterances to include more contextual information. It is worth noting that we only add the user's historical utterances and not the system's historical utterances. Keep in mind that in the inference phase, the system's utterances are generated by the model, so adding the system's historical utterances could lead to error propagation. Second, the entity names of the utterances are obtained by looking them up in the local KB. Third, the entity and attribute pairs are added at the end of the user's current utterance by looking up the local KB, which may prompt the user's requests more explicitly. Accordingly, the triples of entity, property, and value are constructed as a sequence by looking up the local KB. In this way, it is more like a multiple choice problem to satisfy the user's requirements in the response. Moreover, the constraints according to the user's main intent are removed in the output compared with the official baseline method, since it is difficult to predict these constraints accurately.

3 Experiments

3.1 Implementation

The large pre-trained language model used in the experiments is GPT-2(Radford et al., 2019). And the initial model we used was trained by UER(Zhao et al., 2019) in Huggingface Transformers(Wolf et al., 2020). The sequence of the user's historical utterances longer than 384 is truncated. During training, the AdamW is used as an optimizer with an initial learning rate of 5e-5. All the model in these experiments is trained for 20 epochs, with the learning rate decreasing linearly to 0.

3.2 Evaluation

As illustrated in Section 2, we mainly tried to use many different prompting schemes to improve performance. The input pattern used in the official baseline method is referred to as the baseline, where the sequence consists of the entity name history, the current user utterance, the predicted entity name, the information extracted from the local KB, the system intent, and the system response. In general, the baseline is mainly improved in 5 aspects:

- FTU: fine-tuning on the unlabeled data.
- HISU: adding the historical utterances of the user.
- UC: the entity and attribute pairs are added at the end of the user's current utterance of user by looking it up in the local KB.
- KC: the triples of entity, attribute and value are constructed as a sequence by looking them up in the local KB.
- WOUC: the constraints according to the main user intent are removed.

As shown in Table 1, fine-tuning on the unlabeled data improves BLUE, which is consistent with intuition. Adding the user's historical utterances improves all metrics, as they can give more clues to the model. Then, the entity and attribute pairs are added at the end of the current user's utterance, and this strategy leads the model to make progress in predicting user intent and system intent by 1.42 and 0.9 respectively. This result shows that the entity and attribute pairs may have the ability to prompt the intent of the user more explicitly. Note that this strategy does not improve the success rate. The

Table 1: Quantitative comparison of each sequence pattern

Sequence Pattern	User intent(F1)	System intent(F1)	BLEU	Success
Baseline	64.42	57.45	4.38	29.72
Baseline+FTU	63.96	58.35	4.77	26.10
Baseline+FTU+HISU	65.00	59.10	5.425	30.52
Baseline+FTU+HISU+UC	66.42	60.00	5.651	30.12
Baseline+FTU+HISU+UC+KC+WOUC	**70.00**	**60.03**	**6.144**	**41.36**

reason is considered as the mismatch of the entity and attribute pairs and the sequence extracted from the local KB, which is the same as the official baseline method. Therefore, we reformat the sequence extracted from the local KB as triples of entity, attribute, and value. The success rate increases significantly from 30.12 to 41.36 when UC and KC are used together. At the same time, the constraints related to the main user intent are removed, compare with the official baseline method, since we find it is difficult to predict these constraints accurately. After removing the constraints of the user's main intent, the user intent F1 is also improved. In addition, it is difficult to significantly improve the BLUE score because the response of the system can be very diverse.

4 Conclusion

In this work, a large pretrained language model based task oriented dialog system is built for the EMNLP SereTOD Track 2 competition. GPT2 is selected as the pre-trained language model in this work and fine-tuned on the unlabeled data. Then, the user's historical utterances are added to provide more contextual information. Motivated by the good use of local KB, the entity and attribute pairs are added at the end of the user's current utterance by looking up the local KB. At the same time, the triples of entity, attribute and value are extracted from the local KB and used as a complementary part. With these improvements, our method ranks second in the competition with a final score of 2.7. In the future, there is still much room for improvement in the performance of the method by making progress in data cleaning, well-designed rules, leveraging unlabeled data, and model integration.

References

Ehsan Hosseini-Asl, Bryan McCann, Chien-Sheng Wu, Semih Yavuz, and Richard Socher. 2020. A simple language model for task-oriented dialogue. *Advances in Neural Information Processing Systems*, 33:20179–20191.

Yohan Lee. 2021. Improving end-to-end task-oriented dialog system with a simple auxiliary task. In *Findings of the Association for Computational Linguistics: EMNLP 2021*, pages 1296–1303.

Hong Liu, Yucheng Cai, Zhijian Ou, Yi Huang, and Junlan Feng. 2022a. Revisiting markovian generative architectures for efficient task-oriented dialog systems. *arXiv preprint arXiv:2204.06452*.

Hong Liu, Hao Peng, Zhijian Ou, Juan-Zi Li, Yi Huang, and Junlan Feng. 2022b. Information extraction and human-robot dialogue towards real-life tasks: A baseline study with the mobilecs dataset. *arXiv preprint arXiv:2209.13464*.

Zhijian Ou, Junlan Feng, Juanzi Li, Yakun Li, Hong Liu, Hao Peng, Yi Huang, and Jiangjiang Zhao. 2022. A challenge on semi-supervised and reinforced task-oriented dialog systems. *arXiv preprint arXiv:2207.02657*.

Alec Radford, Jeff Wu, Rewon Child, David Luan, Dario Amodei, and Ilya Sutskever. 2019. Language models are unsupervised multitask learners.

Thomas Wolf, Lysandre Debut, Victor Sanh, Julien Chaumond, Clement Delangue, Anthony Moi, Pierric Cistac, Tim Rault, Rémi Louf, Morgan Funtowicz, et al. 2020. Transformers: State-of-the-art natural language processing. In *Proceedings of the 2020 conference on empirical methods in natural language processing: system demonstrations*, pages 38–45.

Yunyi Yang, Yunhao Li, and Xiaojun Quan. 2021. Ubar: Towards fully end-to-end task-oriented dialog system with gpt-2. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 35, pages 14230–14238.

Zhe Zhao, Hui Chen, Jinbin Zhang, Xin Zhao, Tao Liu, Wei Lu, Xi Chen, Haotang Deng, Qi Ju, and Xiaoyong Du. 2019. Uer: An open-source toolkit for pre-training models. *EMNLP-IJCNLP 2019*, page 241.