

A General Planning-Based Framework for Goal-Driven Conversation Assistant

Zhuoxuan Jiang,¹ Jie Ma,¹ Jingyi Lu,^{2*} Guangyuan Yu,^{3*} Yipeng Yu,¹ Shaochun Li¹

¹IBM Research - China, Beijing, China

²Cornell University, New York, USA

³Rice University, Houston, USA

{jzxjiang, bjmajie, yuyyp, lishaoc}@cn.ibm.com, jl3663@cornell.com, guangyuan.yu@rice.edu

Abstract

We propose a general framework for goal-driven conversation assistant based on Planning methods. It aims to rapidly build a dialogue agent with less handcrafting and make the more interpretable and efficient dialogue management in various scenarios. By employing the Planning method, dialogue actions can be efficiently defined and reusable, and the transition of the dialogue are managed by a Planner. The proposed framework consists of a pipeline of Natural Language Understanding (intent labeler), Planning of Actions (with a World Model), and Natural Language Generation (learned by an attention-based neural network). We demonstrate our approach by creating conversational agents for several independent domains.

Introduction

With the popularity of intelligent assistants, today more and more human affairs can be transferred to the assistants. For example, in the tedious and simple task of restaurant reservation, the user gives the assistant with restaurant information, preferred time slots and the number of persons, and then the assistant could make calls to the restaurants and book a table through conversations. We regard this kind of assistants which can help their users to complete missions with goals as *Conversational Assistant*. In the example mentioned before, the time point and number of persons are the goals.

Traditionally, each turn of a task-oriented dialogue system is built as a pipeline which receives an utterance as input from users and then outputs a response (Peng et al. 2017). The pipeline usually includes: Natural Language Understanding (NLU), Dialogue State Tracking, Dialogue Policy Selection, and Natural Language Generation (NLG) (Zhao and Eskenazi 2016). In practice, it is difficult to integrate the goals with the state tracking and policy selection components without thorough handcrafting, especially in a general framework. Because the goals are important information for specific task to influence the key dialogue management. On the other hand, studies of end-to-end methods are rising recently (Wen et al. 2017; Li et al. 2017). Although these methods need little slot filling for handcrafting, the certainty

*This work was done when Jingyi Lu and Guangyuan Yu were interns at IBM Research - China.

Copyright © 2019, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

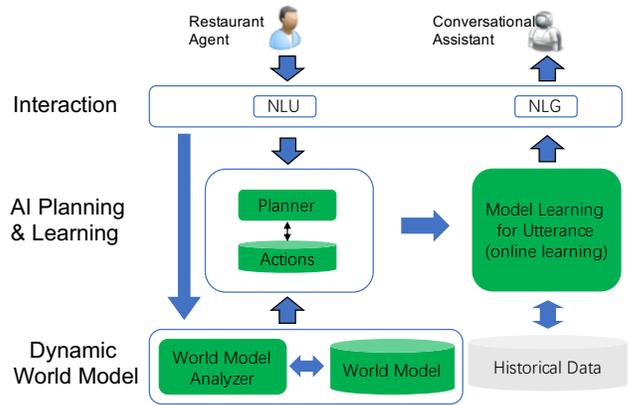


Figure 1: Architecture of the Framework

and interpretability of deep neural network are still questioned in practical applications (Chen et al. 2017).

In this paper, we propose a general framework for building the goal-driven conversational assistants which support various domains. Considering that the AI planning method is suitable to general goal-based tasks (Ghallab, Nau, and Traverso 2016), and the method can be clearly explained for each dialogue action, we integrate the AI planning with NLU and NLG components to constitute the whole planning-based framework for conversation. Noted that the world model required by the planning method is unavoidably task-specific and we show a general methodology that can rapidly build the model. We demonstrate the proposed framework in two different domains: restaurant reservation and financial products sales.

System Overview

As depicted in Figure 1, the proposed framework includes several core modules such as NLU (intent labeler), AI Planning, World Model and generative model learning for NLG.

Natural Language Understanding

In order to understand the opposite utterance, we should define several intents for a specific task. To facilitate this process, IBM Watson AI cloud platform ¹ is leveraged. Only

¹<https://www.ibm.com/watson/ai-assistant/>

when some examples for each intent are given, then the platform can automatically train the model to calculate sentence similarity. For example, if a sentence is posted to the platform by an API call, then the API returns an identified intent list with confidence values ranking from large to small. We refer to the API tutorial document on the platform for implementing the NLU component.

Planning of Conversation

After obtaining the identified intents, we can plan a sequence of dialogue actions which can lead to goals with the shortest path. This means we should define two subcomponents for the AI planning algorithm: Dynamic World Model and Actions. To the planning algorithm, we use the classical STRIPS because it is efficient enough and easy to develop for our tasks.

Dynamic World Model This part is related to store and manage the states of both the agent and the environment. The states in world model is essential for the planning algorithm, which are preconditions and post-conditions. For example in the task of restaurant reservation, the defined states include user profiling, user goals and etc. We also implement a world model analyzer to update all the states based on identified intents under each dialogue turn.

Action Definition Each action for planning should include preconditions and postconditions based on the Planning Domain Definition Language (PDDL) (Norvig and Russell 1995). Here all the conditions are consistent to the states in world model. Therefore, we only need to define the dialogue actions which can update the states. Taking the restaurant reservation as an example, defined actions seem like *say-Name*, *opening* and *callBackLater*.

Natural Language Generation

For each dialogue turn to generate a response, we leverage an attention generation model which can learn from historical data (Bordes, Boureau, and Weston 2017), given both the planned action and previous utterances happened in the same dialogue as input. That is, each action corresponds to an independent generative model, because this way can guarantee the quality of generated utterance. To solve the cold-start issue, some seed utterances can be created to train initial models. Then with the increasing accumulation of historical data, the models can be timely upgraded by online learning.

Case Studies

We apply our framework to the tasks of restaurant reservation and financial products pre-sales. Figure 2 shows the demo interface of the former task. The left part collects and exhibits the input information with a process diagram. The right part is the interaction window. More detailed information of the demos can refer to the videos.

Conclusion and Future Work

In this paper, we propose a general framework for goal-driven conversational assistant based on AI planning methods. Compared with flow-based conversational systems, the

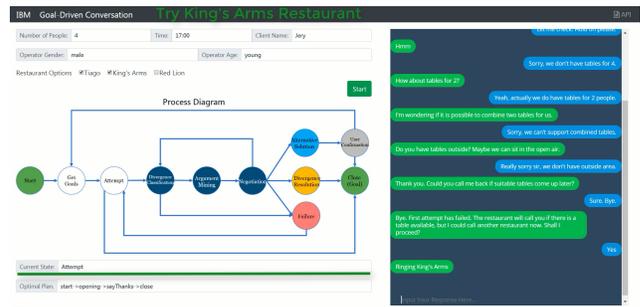


Figure 2: Demo Interface for Restaurant Reservation

framework makes the dialogue build-time more efficient with less handcrafting, and more interpretable in terms of the dialogue action selections than end-to-end methods. In a summary, the handcrafting of our framework for a specific task only includes the definition of goals, intents, states in world model, actions with preconditions and post-conditions, and several initial seed utterances for generating a response if in cold-start situation. Then the planning algorithm can select the optimal goal-driven action for each dialogue turn. Our framework is easy to implement and can be applied to various tasks. For the future work, we can improve the NLG by learning personalized information. For the planning of actions, a parallel method of action selection can be enabled by reinforcement learning from historical data during the actual use of our demons.

References

- Bordes, A.; Boureau, Y.-L.; and Weston, J. 2017. Learning end-to-end goal-oriented dialog. In *ICLR*.
- Chen, H.; Liu, X.; Yin, D.; and Tang, J. 2017. A survey on dialogue systems- recent advances and new frontiers. *ACM SIGKDD Explorations Newsletter* 19(2):25–35.
- Ghallab, M.; Nau, D.; and Traverso, P. 2016. *Automated Planning and Acting*. Cambridge University Press.
- Li, X.; Chen, Y.-N.; Li, L.; Gao, J.; and Celikyilmaz, A. 2017. End-to-end task-completion neural dialogue systems. In *IJCNLP*, 733–743.
- Norvig, P., and Russell, S. J. 1995. *Artificial Intelligence: A Modern Approach*. Prentice Hall.
- Peng, B.; Li, X.; Li, L.; Gao, J.; Celikyilmaz, A.; Lee, S.; and Wong, K.-F. 2017. Composite task-completion dialogue policy learning via hierarchical deep reinforcement learning. In *EMNLP*, 2231–2240.
- Wen, T.-H.; Vandyke, D.; Mrkšić, N.; Gašić, M.; Rojas-Barahona, L. M.; Su, P.-H.; Ultes, S.; and Young, S. 2017. A network-based end-to-end trainable task-oriented dialogue system. In *EACL*, 438–449.
- Zhao, T., and Eskenazi, M. 2016. Towards end-to-end learning for dialog state tracking and management using deep reinforcement learning. In *SIGDIAL*, 1–10.