Open-Domain Neural Dialogue Systems

Yun-Nung Chen
National Taiwan University
Taipei, Taiwan
y.v.chen@ieee.org

Jianfeng Gao
Microsoft AI & Research
Redmond, WA
jfgao@microsoft.com

1 Tutorial Overview

Until recently, the goal of developing open-domain dialogue systems that not only emulate human conversation but fulfill complex tasks, such as travel planning, seemed elusive. However, we start to observe promising results in the last few years as the large amount of conversation data is available for training and the breakthroughs in deep learning and reinforcement learning are applied to dialogue. In this tutorial, we start with a brief introduction to the history of dialogue research. Then, we describe in detail the deep learning and reinforcement learning technologies that have been developed for two types of dialogue systems. First is a task-oriented dialogue system that can help users accomplish tasks, ranging from meeting scheduling to vacation planning. Second is a social bot that can converse seamlessly and appropriately with humans. In the final part of the tutorial, we review attempts to developing open-domain neural dialogue systems by combining the strengths of task-oriented dialogue systems and social bots. The tutorial material is available at http://opendialogue.miulab.tw.

2 Outline

1. Introduction & Background [15 min.]
   - Brief history of dialogue research
   - Challenges of developing dialogue agents
   - Task-oriented dialogue systems
   - Social chat bots
   - How to evaluate dialogue systems
   - Neural network basics
   - Reinforcement learning (RL) basics
2. Task-Oriented Dialogue System [75 min.]
   - Natural language understanding (NLU)
     - Domain and intent classification
     - Slot tagging
     - Joint semantic frame parsing
     - Contextual language understanding
     - Structural language understanding
   - Dialogue management (DM) – Dialogue state tracking (DST)
     - Neural belief tracker
     - Multichannel tracker
   - Dialogue management (DM) – Dialogue policy optimization
     - Dialogue RL signal
     - Deep Q-network for learning policy
     - Hierarchical RL for learning policy
   - Natural language generation (NLG)
     - Rule-based NLG
     - Learning-based NLG
     - Structural NLG
     - Contextual NLG
   - End-to-end task-oriented dialogue systems
     - Joint learning of NLU and DM
     - Supervised learning for dialogues
     - Memory networks for dialogues
     - RL-based *InfoBot*
     - LSTM-based dialogue control
     - RL-based task-completion bots
3. Social Chat Bots [75 min.]
   - Neural response generation models
   - Making the response diverse
   - Making the response consistent
   - Deep reinforcement learning for response generation
   - Image-grounded response generation
   - Knowledge-grounded response generation
   - Generative seq2seq for task-oriented dialogues
   - Combining task-oriented bots and social bots
4. Challenges & Conclusions [15 mins]

Proceedings of the The 8th International Joint Conference on Natural Language Processing, pages 6–10,
Taipei, Taiwan, November 27 – December 1, 2017. ©2017 AFNLP

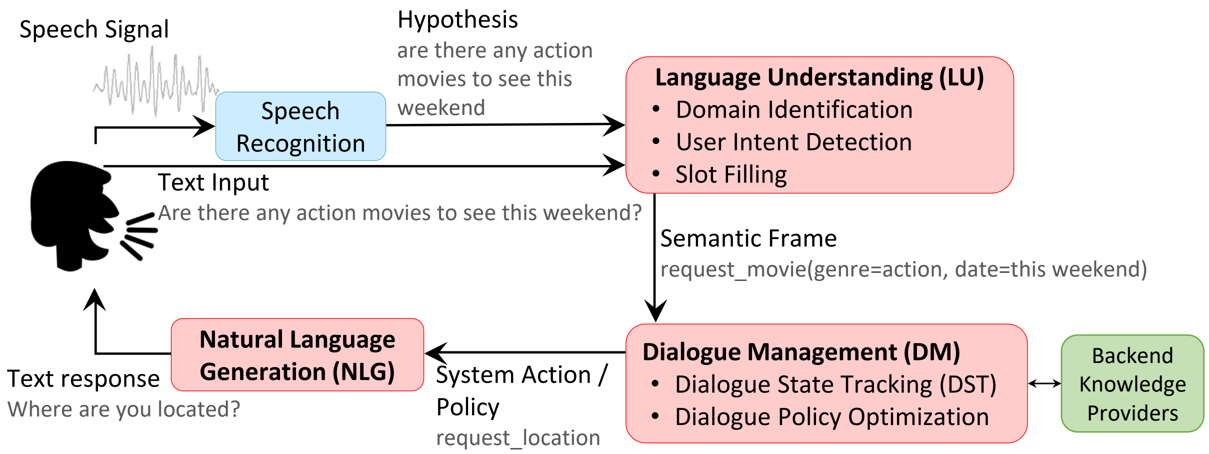

Figure 1: Pipeline framework of task-oriented dialog system.

3 Task-Oriented Dialogue Systems

The architecture of a task-oriented dialogue system is illustrated in Figure 1 (Tur and De Mori, 2011). It consists of three components, natural language understanding (NLU), dialogue management (DM), and natural language generation (NLG) (Rudnicky et al., 1999; Zue et al., 2000; Zue and Glass, 2000).

Natural Language Understanding NLU traditionally consists of domain identification and intent prediction, which are framed as utterance classification problems, and slot filling, framed as a sequence tagging task.

With the advances on deep learning, recent development has been focused on neural approaches. Ravuri and Stolcke (2015) proposed an RNN architecture for intent determination. Xu and Sarikaya (2013) incoporated features generated using neural approaches into the CRF framework for slot filling. Yao et al. (2013) and Mesnil et al. (2015) later employed soly the RNN-based sequence labeling model for slot filling. Such an architecture has been further extended to jointly model intent detection and slot filling in multiple domains (Hakkani-Tür et al., 2016; Jaech et al., 2016). End-to-end memory networks have shown to provide a good mechanism for integrating global knowledge context and local dialogue context into these models (Chen et al., 2016a,b). In addition, the importance of the NLU module is investigated in Li et al. (2017a), showing that different types of errors from NLU can degrade the whole system's performance in a reinforcement learning setting.

Dialogue Management DM plays two roles, tracking the dialogue state and performing the dialogue policy (i.e., telling the agent how to act given the dialogue state.)

The state-of-the-art dialogue managers monitor the dialogue progress (state) using neural dialogue state tracking models (Henderson et al., 2013). Recent work shows that that Neural Dialog Managers provide conjoint representations between the utterances, slot-value pairs as well as knowledge graph representations (Wen et al., 2016; Mrkšić et al., 2016; Liu and Lane, 2017), and thus make the deployment of large-scale dialogue systems for complex domain much easier.

A partially observable Markov decision process (POMDP) has been shown to be an effective mathematical framework for dialogue policy learning since it can model the uncertainty such as those caused by speech recognition errors and semantic decoding errors (Williams and Young, 2007; Young et al., 2013). Under POMDP, dialogue policy is trained using reinforcement learning (RL) where the agent learns how to act based on the reward signals recieved from the environment (Sutton and Barto, 1998).

Natural Language Generation NLG approaches can be grouped into two categories, one focuses on generating text using templates or rules (linguistic) methods, the other uses corpus-based statistical methods (Oh and Rudnicky, 2002).

The RNN-based models have been applied to language generation for both social bots and task-orientated dialogue systems (Sordoni et al., 2015; Vinyals and Le, 2015; Wen et al., 2015b). The RNN-based NLG can learn from unaligned

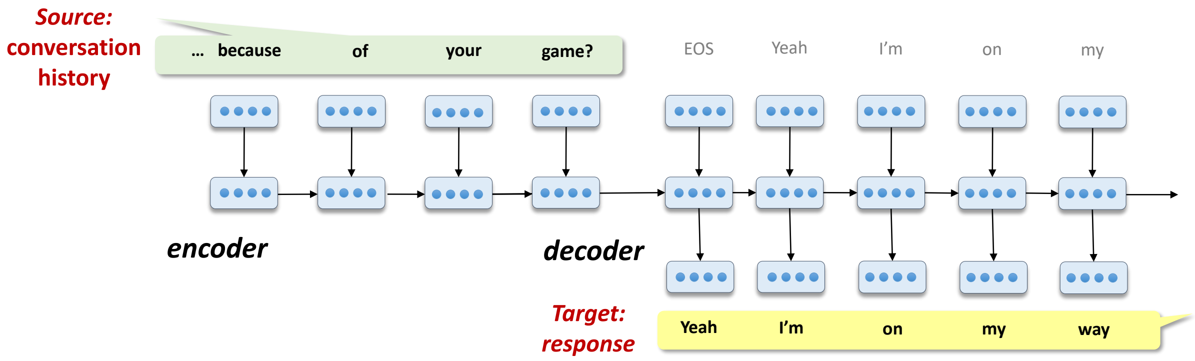

Figure 2: Illustration of a sequence-to-sequence model for chit-chat dialogues.

data by jointly optimizing sentence planning and surface realization, and language variation can be achieved by sampling from output candidates (Wen et al., 2015a). Moreover, Wen et al. (2015b) improved the prior work by adding a gating mechanism to control the dialogue act during generation in order to avoid repetition.

End-to-End Task-Oriented Dialogue System Awaring the representation power of deep neural networks, there are more and more attempts to learning dialogue systems in an end-to-end fashion using different learning frameworks, including supervised learning and reinforcement learning (Yang et al., 2017).

Wen et al. (2016) and Bordes and Weston (2016) introduced a network-based end-to-end trainable task-oriented dialogue system. The authors treated training a dialogue system as learning a mapping from dialogue histories to system responses, and applied an encoder-decoder model.However, the system is trained in a supervised fashion that requires a lot of training data. Thus, the agent cannot learn a robust dialogue policy since it never explore the unknown space that is not covered by the limited training data.

Zhao and Eskenazi (2016) presented an end-to-end reinforcement learning (RL) approach to dialogue state tracking and policy learning. They show some promising results when applying the agent to the task of guessing the famous person a user is thinking of. Dhingra et al. (2017) proposed an end-to-end differentiable KB-Infobot for efficient information access. Li et al. (2017b) presented an end-to-end neural dialogue system for task completion. The agent can handle a wide variety of question types, including user-initated request.

4 Social Chat Bots

Social bots are of growing importance in facilitating smooth interaction between humans and their electronic devices. Recently, researcher have begun to explore data-driven generation of conversational responses within the framework of nerual machine translation (NMT) in the form of encoder-decoder or seq2seq models (Sordoni et al., 2015; Vinyals and Le, 2015; Li et al., 2016a), as illustrated in Figure 2.

However, the generated responses are often too general to carry meaningful information, such as "I don't know.", which can serve as a response to any user questions. A mutual information based model was proposed to address the issue, a mutual information model is proposed by Li et al. (2016a), and is later improved by using deep reinforcement learning (Li et al., 2016c). Furthermore, Li et al. (2016b) presented a persona-based model to address the issue of speaker consistency in neural response generation.

Although task-oriented dialogue systems and social bots are originally developed for different purposes, there is a trend of combining both as a step towards building an open-domain dialogue agent.

For example, on the one hand, Ghazvininejad et al. (2017) presented a fully data-driven and knowledge-grounded neural conversation model aimed at producing more contentful responses without slot filling. On the other hand, Zhao et al. (2017) proposed a task-oriented dialogue agented based on the encoder-decoder model with chatting capability.

5 Instructors

Yun-Nung (Vivian) Chen is currently an assistant professor at the Department of Computer Science, National Taiwan University. She earned her Ph.D. degree from Carnegie Mellon University, where her research interests focus on spoken dialogue system, language understanding, natural language processing, and multi-modal speech applications. She received the Google Faculty Research Awards 2016, two Student Best Paper Awards from IEEE SLT 2010 and IEEE ASRU 2013, a Student Best Paper Nominee from Interspeech 2012, and the Distinguished Master Thesis Award from ACLCLP. Before joining National Taiwan University, she worked in the Deep Learning Technology Center at Microsoft Research Redmond. More information about her can be found at http://vivianchen.idv.tw.

Jianfeng Gao is Partner Research Manager in Business AI at Microsoft AI and Research. From 2014 to 2017, he was Partner Research Manager and Principal Researcher at Deep Learning Technology Center at MSR Redmond. He leads the development of AI solutions to Predictive Sales and Marketing. Gao also works on deep learning for text and image processing and leads the development of AI systems for dialogue, machine reading comprehension, and question answering. From 2006 to 2014, he was a principal researcher at NLP Group at MSR Redmond, working on Web search, query understanding and reformulation, ads prediction, and statistical machine translation. From 2005 to 2006, Gao was a research lead in Natural Interactive Services Division at Microsoft, working on Project X, an effort of developing natural user interface for Windows. From 1999 to 2005, Gao was a research lead in Natural Language Computing Group at MSR Asia, developing the first Chinese speech recognition system released with Microsoft Office, the Chinese/Japanese Input Method Editors which were the leading products in the market, and the natural language platform for Windows Vista. More information can be found at https://www.microsoft.com/en-us/research/people/jfgao/.

References

Antoine Bordes and Jason Weston. 2016. Learning end-to-end goal-oriented dialog. *arXiv preprint arXiv:1605.07683*.

Yun-Nung Chen, Dilek Hakkani-Tur, Gokhan Tur, Asli Celikyilmaz, Jianfeng Gao, and Li Deng. 2016a. Knowledge as a teacher: Knowledge-guided structural attention networks. *arXiv preprint arXiv:1609.03286*.

Yun-Nung Chen, Dilek Hakkani-Tür, Gokhan Tur, Jianfeng Gao, and Li Deng. 2016b. End-to-end memory networks with knowledge carryover for multi-turn spoken language understanding. In *Proceedings of the Interspeech*.

Bhuwan Dhingra, Lihong Li, Xiujun Li, Jianfeng Gao, Yun-Nung Chen, Faisal Ahmed, and Li Deng. 2017. Towards end-to-end reinforcement learning of dialogue agents for information access. pages 484–495.

Marjan Ghazvininejad, Chris Brockett, Ming-Wei Chang, Bill Dolan, Jianfeng Gao, Wen-tau Yih, and Michel Galley. 2017. A knowledge-grounded neural conversation model. *arXiv preprint arXiv:1702.01932*.

Dilek Hakkani-Tür, Gokhan Tur, Asli Celikyilmaz, Yun-Nung Chen, Jianfeng Gao, Li Deng, and Ye-Yi Wang. 2016. Multi-domain joint semantic frame parsing using bi-directional RNN-LSTM. In *Proceedings of the Interspeech*, San Francisco, CA.

Matthew Henderson, Blaise Thomson, and Steve Young. 2013. Deep neural network approach for the dialog state tracking challenge. In *Proceedings of the SIGDIAL 2013 Conference*, pages 467–471.

A. Jaech, L. Heck, and M. Ostendorf. 2016. Domain adaptation of recurrent neural networks for natural language understanding. In *Proceedings of the Interspeech*, San Francisco, CA.

Jiwei Li, Michel Galley, Chris Brockett, Jianfeng Gao, and Bill Dolan. 2016a. A diversity-promoting objective function for neural conversation models. In *Proceedings of NAACL-HLT*, pages 110–119.

Jiwei Li, Michel Galley, Chris Brockett, Georgios P Spithourakis, Jianfeng Gao, and Bill Dolan. 2016b. A persona-based neural conversation model. In *Proceedings of ACL*.

Jiwei Li, Will Monroe, Alan Ritter, Michel Galley, Jianfeng Gao, and Dan Jurafsky. 2016c. Deep reinforcement learning for dialogue generation. In *Proceedings of EMNLP*.

Xiujun Li, Yun-Nung Chen, Lihong Li, Jianfeng Gao, and Asli Celikyilmaz. 2017a. Investigation of language understanding impact for reinforcement learning based dialogue systems. *arXiv preprint arXiv:1703.07055*.

Xuijun Li, Yun-Nung Chen, Lihong Li, Jianfeng Gao, and Asli Celikyilmaz. 2017b. End-to-end task-completion neural dialogue systems. In *Proceedings of IJCNLP*.

Bing Liu and Ian Lane. 2017. An end-to-end trainable neural network model with belief tracking for task-oriented dialog. In *Proceedings of Interspeech*.

Grégoire Mesnil, Yann Dauphin, Kaisheng Yao, Yoshua Bengio, Li Deng, Dilek Hakkani-Tur, Xiaodong He, Larry Heck, Gokhan Tur, Dong Yu, et al. 2015. Using recurrent neural networks for slot filling in spoken language understanding. *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, 23(3):530–539.

Nikola Mrkšić, Diarmuid Ó Séaghdha, Tsung-Hsien Wen, Blaise Thomson, and Steve Young. 2016. Neural belief tracker: Data-driven dialogue state tracking. *arXiv preprint arXiv:1606.03777*.

Alice H Oh and Alexander I Rudnicky. 2002. Stochastic natural language generation for spoken dialog systems. *Computer Speech & Language*, 16(3):387–407.

Suman Ravuri and Andreas Stolcke. 2015. Recurrent neural network and lstm models for lexical utterance classification. In *Sixteenth Annual Conference of the International Speech Communication Association*.

Alexander I Rudnicky, Eric H Thayer, Paul C Constantinides, Chris Tchou, R Shern, Kevin A Lenzo, Wei Xu, and Alice Oh. 1999. Creating natural dialogs in the carnegie mellon communicator system. In *Eurospeech*.

Alessandro Sordoni, Michel Galley, Michael Auli, Chris Brockett, Yangfeng Ji, Margaret Mitchell, Jian-Yun Nie, Jianfeng Gao, and Bill Dolan. 2015. A neural network approach to context-sensitive generation of conversational responses. In *Proceedings of NAACL-HLT*.

Richard S Sutton and Andrew G Barto. 1998. *Reinforcement learning: An introduction*, volume 1. MIT press Cambridge.

Gokhan Tur and Renato De Mori. 2011. *Spoken language understanding: Systems for extracting semantic information from speech*. John Wiley & Sons.

Oriol Vinyals and Quoc Le. 2015. A neural conversational model. *arXiv preprint arXiv:1506.05869*.

Tsung-Hsien Wen, Milica Gašic, Dongho Kim, Nikola Mrkšic, Pei-Hao Su, David Vandyke, and Steve Young. 2015a. Stochastic language generation in dialogue using recurrent neural networks with convolutional sentence reranking. In *16th Annual Meeting of the Special Interest Group on Discourse and Dialogue*, page 275.

Tsung-Hsien Wen, Milica Gasic, Nikola Mrksic, Lina M Rojas-Barahona, Pei-Hao Su, Stefan Ultes, David Vandyke, and Steve Young. 2016. A network-based end-to-end trainable task-oriented dialogue system. *arXiv preprint arXiv:1604.04562*.

Tsung-Hsien Wen, Milica Gasic, Nikola Mrksic, Pei-Hao Su, David Vandyke, and Steve Young. 2015b. Semantically conditioned lstm-based natural language generation for spoken dialogue systems. *arXiv preprint arXiv:1508.01745*.

Jason D Williams and Steve Young. 2007. Partially observable markov decision processes for spoken dialog systems. *Computer Speech & Language*, 21(2):393–422.

Puyang Xu and Ruhi Sarikaya. 2013. Convolutional neural network based triangular CRF for joint intent detection and slot filling. In *2013 IEEE Workshop on Automatic Speech Recognition and Understanding (ASRU)*, pages 78–83. IEEE.

Xuesong Yang, Yun-Nung Chen, Dilek Hakkani-Tür, Paul Crook, Xiujun Li, Jianfeng Gao, , and Li Deng. 2017. End-to-end joint learning of natural language understanding and dialogue manager. In *Proceedings of ICASSP*.

Kaisheng Yao, Geoffrey Zweig, Mei-Yuh Hwang, Yangyang Shi, and Dong Yu. 2013. Recurrent neural networks for language understanding. In *INTERSPEECH*, pages 2524–2528.

Steve Young, Milica Gašić, Blaise Thomson, and Jason D Williams. 2013. POMDP-based statistical spoken dialog systems: A review. *Proceedings of the IEEE*, 101(5):1160–1179.

Tiancheng Zhao and Maxine Eskenazi. 2016. Towards end-to-end learning for dialog state tracking and management using deep reinforcement learning. In *Proceedings of SIGDIAL*.

Tiancheng Zhao, Allen Lu, Kyusong Lee, and Maxine Eskenazi. 2017. Generative encoder-decoder models for task-oriented spoken dialog systems with chatting capability. In *Proceedings of SigDial*.

Victor Zue, Stephanie Seneff, James R Glass, Joseph Polifroni, Christine Pao, Timothy J Hazen, and Lee Hetherington. 2000. JUPITER: a telephone-based conversational interface for weather information. *IEEE Transactions on speech and audio processing*, 8(1):85–96.

Victor W Zue and James R Glass. 2000. Conversational interfaces: Advances and challenges. *Proceedings of the IEEE*, 88(8):1166–1180.