



- Introduction and Background
- Modular Dialogue System
 - Spoken/Natural Language Understanding (SLU/NLU)
 - Dialogue State Tracking (DST)
 - Dialogue Policy
 - Natural Language Generation (NLG)
- End-to-End Learning for Dialogue Systems
- Conclusion



Introduction & Background

Language Empowering Intelligent Assistants



Apple Siri (2011)

Google Now (2012)

Microsoft Cortana (2014)



Amazon Alexa/Echo (2014)

Facebook M & Bot (2015)

Google Home (2016)

Dialogue System

- Task-Oriented
 - Personal assistant, achieve a certain task
 - Combination of <u>rules</u> and statistical components
 - POMDP for spoken dialog systems (Williams and Young, 2007)
 - Learning End-to-End Goaloriented Dialog (Antoni and Weston, 2016)
 - An End-to-End Trainable Task-oriented Dialogue System (Wen el al., 2016)

Chit-Chat

- No specific goal, focus on conversation flow
- Work using variants of seq2seq model
 - A Neural Conversation Model (Vinyals and Le, 2015)
 - Deep Reinforcement Learning for Dialogue Generation (Li et al., 2016)
 - Conversational Contextual Cues: The Case of Personalization & History for Response Ranking (Al-Rfou et al., 2016)

Pipelined Task-Oriented Dialogue System



⁸ Part II

Modular Dialogue System

- Introduction and Background
- Modular Dialogue System
 - Spoken/Natural Language Understanding (SLU/NLU)
 - Dialogue State Tracking (DST)
 - Dialogue Policy
 - Natural Language Generation (NLG)
- End-to-End Learning for Dialogue Systems
- Conclusion

Semantic Frame Representation

- Requires a domain ontology
- Contains core content (intent, a set of slots with fillers)



Language Understanding (LU)



LU – Domain/Intent Classification

Mainly viewed as an utterance classification task

 Given a collection of utterances u_i with labels c_i, D= {(u₁,c₁),...,(u_n,c_n)} where c_i ∈ C, train a model to estimate labels for new utterances u_k.

find me a cheap taiwanese restaurant in oakland

...

Movies Restaurants Sports Weather Music

...

Find_movie Buy_tickets Find_restaurant Book_table Find_lyrics

Language Understanding - Slot Filling





Figure credited by Milica Gašić

Language Understanding

- Intent Classification (Ravuri and Stolcke, 2015)
 intent h_0 h_1 h_2 h_n w_0 w_1 w_2 w_n
- IOB Sequence Labeling for Slot Filling (Hakkani-Tur et al., 2016)



Joint Semantic Frame Parsing



- Introduction and Background
- Modular Dialogue System
 - Spoken/Natural Language Understanding (SLU/NLU)
 - Dialogue State Tracking (DST)
 - Dialogue Policy
 - Natural Language Generation (NLG)
- End-to-End Learning for Dialogue Systems
- Conclusion

Dialogue State Tracking (DST)

sample problem

- S: where would you like to fly from?
- U: [Boston/0.45]; [Austin/0.30]
- S: sorry, did you say you wanted to fly from Boston?
- U: [No/0.37] + [Aspen / 0.7]



Dialogue State Tracking (DST)

- 18
- Maintain a probabilistic distribution instead of a 1-best prediction for <u>better robustness</u>



Dialogue State Tracking (DST)

 Maintain a probabilistic distribution instead of a 1-best prediction for <u>better robustness to SLU errors or</u> ambiguous input

Slot	Value
# people	5 (0.5)
time	5 (0.5)

Slot	Value
# people	3 (0.8)
time	5 (0.8)



- Introduction and Background
- Modular Dialogue System
 - Spoken/Natural Language Understanding (SLU/NLU)
 - Dialogue State Tracking (DST)
 - Dialogue Policy
 - Natural Language Generation (NLG)
- End-to-End Learning for Dialogue Systems
- Conclusion

Dialogue Policy Optimization



Optimized dialogue policy selects the best action that can maximize the future reward. Correct rewards are a crucial factor in dialogue policy training

Reward for RL \cong Evaluation for SDS

- Dialogue is a special RL task
 - Human involves in <u>interaction</u> and <u>rating</u> (evaluation) of a dialogue
 - Fully human-in-the-loop framework
- Rating: correctness, appropriateness, and adequacy

- Expert rating	high quality, <mark>high</mark> cost
- User rating	unreliable quality, medium cost
- Objective rating	Check desired aspects, low cost

Dialogue Reinforcement Signal

Typical Reward Function

- per turn penalty -1
- Large reward at completion if successful
- Typically requires domain knowledge
 - ✓ Simulated user
 - Amazon Mechanical Turk)
 - Real users







User Simulation

- User Simulation
 - Goal: generate natural and reasonable conversations to enable reinforcement learning for exploring the policy space
 - Approach
 - Rule-based crafted by experts (Li et al., 2016)
 - Learning-based (Schatzmann et al., 2006)



- Introduction and Background
- Modular Dialogue System
 - Spoken/Natural Language Understanding (SLU/NLU)
 - Dialogue State Tracking (DST)
 - Dialogue Policy
 - Natural Language Generation (NLG)
- End-to-End Learning for Dialogue Systems
- Conclusion

Natural Language Generation (NLG)

Mapping semantic frame into natural language inform(name=Seven_Days, foodtype=Chinese)

Seven Days is a nice Chinese restaurant

Template-Based NLG

Define a set of rules to map frames to NL

Semantic Frame	Natural Language
confirm()	"Please tell me more about the product your are looking for."
confirm(area=\$V)	"Do you want somewhere in the \$V?"
confirm(food=\$V)	"Do you want a \$V restaurant?"
confirm(food=\$V,area=\$W)	"Do you want a \$V restaurant in the \$W."

Pros: simple, error-free, easy to control **Cons:** time-consuming, poor scalability

Trainable Generator (Walker et al., 2002)

Divide the problem into pipeline



Statistical sentence plan generator (Stent et al., 2009)

Statistical surface realization (Dethlefs et al., 2013; Cuayáhuitl et al.,

2014; ...)

Pros: can model complex linguistic structures **Cons:** heavily engineered, require domain knowledge

RNN Language Generator (Wen et al., 2015)





- Introduction and Background
- Modular Dialogue System
 - Spoken/Natural Language Understanding (SLU/NLU)
 - Dialogue State Tracking (DST)
 - Dialogue Policy
 - Natural Language Generation (NLG)
- End-to-End Learning for Dialogue Systems
- Conclusion

E2E Joint NLU and DM (Yang et al., 2017)

Errors from DM can be propagated to NLU for better robustness



E2E Supervised Dialogue System (Wen et al., 2016)



Wen, et al., "A Network-based End-to-End Trainable Task-Oriented Dialogue System," arXiv.:1604.04562v2.

Info Access Bot (Info-Bot) (Dhingra et al., 2016)



Dhingra et al., "End-to-End Reinforcement Learning of Dialogue Agents for Information Access," arXiv.:1609.00777v2.

E2E Task-Completion Bot (TC-Bot) (Li et al., 2017)



Idea: supervised learning for each component and reinforcement learning for end-to-end training the neural dialogue system

E2E Task-Completion Bot (TC-Bot) (Li et al., 2017)

User goal: Two tickets for "the witch" tomorrow 9:30 PM at regal meridian 16, Seattle.

RULE BASED SYSTEM

User: Can I get 2 tickets for the witch? Agent: Which theater do you want? User: Regal meridian 16 Agent: What time would you like to see it? User: 9:30 pm Agent: What date would you like to see it? User: Tomorrow! Agent: How many people are you? User: Just two adults. Agent: What is your location? User: Seattle! Agent: What kind of video format do you like? User: Oh, any kind is okay. Agent: Do you have any price constraint? User: No, I don't care. Agent: Great - I was able to purchase 2 tickets for you to see the witch tomorrow at regal meridian 16 theater in Seattle at 9:30 pm. User: Thank vou. Agent: Thank you!

REINFORCEMENT LEARNING SYSTEM

User: Can I get 2 tickets for the witch? Agent: What date would you like to see it? User: Tomorrow Agent: Which theater do you prefer? User: Regal meridian 16 Agent: What is your location? User: Seattle Agent: What time would you like to see it? User: 9:30 pm. Agent: Great - I was able to purchase 2 tickets for you to see the witch tomorrow at regal meridian 16 theater in Seattle at 9:30 pm. User: Thanks. Agent: Thanks!

- Introduction and Background
- Modular Dialogue System
 - Spoken/Natural Language Understanding (SLU/NLU)
 - Dialogue State Tracking (DST)
 - Dialogue Policy
 - Natural Language Generation (NLG)
- End-to-End Learning for Dialogue Systems
- Conclusion

Concluding Remarks

Modular dialogue system

