

Open-Domain Neural Dialogue Systems

YUN-NUNG (VIVIAN) CHEN 陳溫儂 HTTP://VIVIANCHEN.IDV.TW



Iron Man (2008)

What can machines achieve now or in the future?



Language Empowering Intelligent Assistants

Material: http://opendialogue.miulab.tw



Apple Siri (2011)



Google Now (2012) Google Assistant (2016)



Microsoft Cortana (2014)







Amazon Alexa/Echo (2014)

Facebook M & Bot (2015)

D15) Google Home (2016)

) Apple HomePod (2017)



Why and When We Need?

- "I want to chat"
- "I have a question"
- "I need to get this done"
- "What should I do?"

Turing Test (talk like a human) Information consumption Task completion



Why and When We Need?

- "I want to chat"
- "I have a question"
- "I need to get this done"
- "What should I do?"

Turing Test (talk like a human)

Information consumption

Task completion



Why and When We Need?

- "I want to chat"
- "I have a question"
- "I need to get this done"
- "What should I do?"

Turing Test (talk like a human) Information consumption

Task completion

- What is today's agenda?
- Which room is dialogue tutorial in?
- What does ISCSLP stand for?



Why and When We Need?

- "I want to chat"
- "I have a question"
- "I need to get this done"
- "What should I do?"

Turing Test (talk like a human)

Information consumption

Task completion

- Book me the train ticket from Taoyuan to Taipei
- Reserve a table at Din Tai Fung for 5 people, 7PM tonight
- Schedule a meeting with Vivian at 10:00 tomorrow



Why and When We Need?

- "I want to chat"
- "I have a question"
- "I need to get this done"
- "What should I do?"

Turing Test (talk like a human)

Information consumption

Task completion

- Is this bubble tea worth to try?
- Is the ISCSLP conference good to attend?



Why and When We Need?

- "I want to chat"
- "I have a question"
- "I need to get this done"
- "What should I do?"





Intelligent Assistants





Conversational Agents



¹² Task-Oriented Dialogue Systems



JARVIS – Iron Man's Personal Assistant



Baymax – Personal Healthcare Companion



Task-Oriented Dialogue Systems (Young, 2000)





Task-Oriented Dialogue Systems (Young, 2000)





Semantic Frame Representation

- Requires a domain ontology: early connection to backend
- Contains core concept (intent, a set of slots with fillers)





Backend Database / Ontology

Material: http://opendialogue.miulab.tw

date

- Domain-specific table
 - Target and attributes
- Functionality
 - Information access: find specific entries



Task completion: find the <u>row</u> that satisfies the constraints

Movie Name	Theater	Rating	Date	Time
Iron Man Last	Taipei A1	8.5	2018/10/31	09:00
Iron Man Last	Taipei A1	8.5	2018/10/31	09:25
Iron Man Last	Taipei A1	8.5	2018/10/31	10:15
Iron Man Last	Taipei A1	8.5	2018/10/31	10:40



Task-Oriented Dialogue Systems (Young, 2000)





Language Understanding (LU)

Material: http://opendialogue.miulab.tw Pipelined 1. Domain 2. Intent 3. Slot Filling Classification Classification



RNN for Slot Tagging – I (Yao et al, 2013; Mesnil et al, 2015)

- Variations:
 - a. RNNs with LSTM cells
 - b. Input, sliding window of n-grams
 - c. Bi-directional LSTMs





RNN for Slot Tagging – II (Kurata et al., 2016; Simonnet et al., 2015)

- Encoder-decoder networks
 - Leverages sentence level information
- Attention-based encoder-decoder
 - Use attention (as in MT) in the encoderdecoder network







Joint Semantic Frame Parsing





Joint Model Comparison

	Attention Mechanism	Intent-Slot Relationship
Joint bi-LSTM	Х	∆ (Implicit)
Attentional Encoder-Decoder	V	∆ (Implicit)
Slot Gate Joint Model	V	√ (Explicit)



Slot-Gated Joint SLU (Goo et al., 2018)



□ Slot Prediction $y_i^S = softmax(W^S(h_i + c_i^S) + b^S)$ \longrightarrow $y_i^S = softmax(W^S(h_i + g \cdot c_i^S) + b^S)$ W: trainable matrix v: trainable vector g: scalar gate value

 $m{g}$ will be larger if slot and intent are better related



Contextual Language Understanding







E2E MemNN for Contextual LU (Chen et al., 2016)





Role-Based & Time-Aware Attention (Su et al., 2018)





Context-Sensitive Time-Decay Attention (Su et al., 2018)

Material: http://opendialogue.miulab.tw



Time-decay attention significantly improves the understanding results



Task-Oriented Dialogue Systems (Young, 2000)





Dialogue State Tracking





DNN for DST

Material: http://opendialogue.miulab.tw



multi-turn conversation



RNN-CNN DST (Wen et al., 2016)



(Figure from Wen et al, 2016)



Dialogue Policy Optimization





Dialogue Policy Optimization





Reward for RL \cong Evaluation for System

- 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



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

Material: http://opendialogue.miulab.tw



Idea: SL for each component and RL for end-to-end training the neural dialogue system



Reinforcement Learning Agent (Li et al., 2017)

Material: http://opendialogue.miulab.tw

User goal - Two tickets for "Deadpool" tomorrow 9PM at AMC Pacific Place 11 theater, Seattle.

RULE BASED SYSTEM

User: Do you have any ticket for movie 'Deadpool' at 9:00 pm? Agent: Which theater do you want? User: AMC pacific place 11 theater, please! Agent: How many tickets do you want? User: Two! Agent: What date? User: Tomorrow! Agent: Which city do you want? User: Seattle! Agent: Do you have any specification for zip code? User: No, I do not care. Agent: Do you have any distance constraint? User: No, I do not care. Agent: How about video format? 2d? Imax 3d? Issue 1: no notion about User: Any kind is okay. what requests can be skipped Agent: Do you have any constraint to price? User: No!

REINFORCEMENT LEARNING SYSTEM

User: I want two tickets for movie 'Deadpool'. Agent: What is the starttime you prefer? User: 9:00 PM Agent: Which day? User: Tomorrow Agent: Which city? User: Seattle Agent: Which theater do you prefer? User: AMC Pacific Place 11 theater. Agent: Okay, I will issue 2 tickets for you, tomorrow 9:00 pm at AMC pacific place 11 theater, Seattle, movie 'Deadpool'. User: Okay, thanks. Agent: Thanks!

Agent: Okay, I will issue 2 tickets for you, tomorrow 9:00 pm at AMC pacific place 11 theater, Seattle, movie 'Deadpool'.

User: Okay, thanks. Agent: Thanks!

Skip the requests the user may not care about to improve efficiency


Reinforcement Learning Agent (Li et al., 2017)

Material: http://opendialogue.miulab.tw

User goal – Two tickets for "Hail Caesar" tomorrow 7:15PM at AMC Lowes Oak tree 6, Seattle.

RULE BASED SYSTEM

User: Do you any ticket for movie 'Hail Caesar' at theater AMC Lowes Oak Tree 6? Agent: What start time do you prefer? User: 7:15 pm Tomorrow Agent: Which city you are looking for? User: Seattle Agent: How many tickets do you want? User: Two Agent: Okay, I have issued two tickets for movie 'Hail Caeser' tomorrow 7:50 pm at theater AMC Lowes Oak Tree 6, Seattle. User: Oh, no. Issue 2: no principle way to deal

Agent: Thanks!

with LU errors/uncertainty

REINFORCEMENT LEARNING SYSTEM

User: Do you any ticket for movie 'Hail Caesar' at theater AMC Lowes Oak Tree 6? Agent: What start time do you prefer? User: 7:15 pm Tomorrow Agent: Which city you are looking for? User: Seattle Agent: Okay, I will issue you tickets for movie 'Hail Caeser' tomorrow 7:50 pm at theater AMC Lowes Oak Tree 6, Seattle. User: Oh, no. Agent: It is 7:50 pm, right? User: Two tickets for 7:15 pm, please. Agent: Okay, I will issue you two tickets for movie 'Hail Caeser' tomorrow 7:15 pm at theater AMC Lowes Oak Tree 6, Seattle. User: Thanks!

Agent: Thanks!

Automatically learn when and what to confirm for better robustness



Interactive RL for DM (Shah et al., 2016)





RL in Dialogue Systems

- 1. Sample inefficient, hard to design reward function, local optima...
- 2. Real users are expensive
- 3. Discrepancy between real users and simulators



D3Q: Discriminative Deep Dyna-Q (Su et al., 2018)

- 🗆 Idea
 - learning with real users with planning
 - add a *discriminator* to filter out the bad experiences





D3Q: Discriminative Deep Dyna-Q (Su et al., 2018)



The policy learning is more robust and shows the improvement in human evaluation



Task-Oriented Dialogue Systems (Young, 2000)





Natural Language Generation (NLG)

Material: http://opendialogue.miulab.tw

Mapping dialogue acts into natural language

inform(name=Seven_Days, foodtype=Chinese)

Seven Days is a nice Chinese restaurant



Template-Based NLG

Material: http://opendialogue.miulab.tw

Define <u>a set of rules</u> to map frames to natural language

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, rigid, poor scalability



RNN-Based LM NLG (Wen et al., 2015)





Semantic Conditioned LSTM (Wen et al., 2015)

- □ Issue: semantic repetition
 - Din Tai Fung is a great Taiwanese restaurant that serves Taiwanese.
 - Din Tai Fung is a child friendly restaurant, and also allows kids.



Idea: using gate mechanism to control the generated semantics (dialogue act/slots)



Material: http://opendialogue.miulab.tw

Contextual NLG (Dušek and Jurčíček, 2016)

- Goal: adapting users' way of speaking, providing contextaware responses
 - Context encoder
 - Seq2Seq model







Controlled Text Generation (Hu et al., 2017)

- □ Idea: NLG based on generative adversarial network (GAN) framework
 - **c**: targeted sentence attributes





Material: http://opendialogue.miulab.tw

Issues in NLG

lssue

NLG tends to generate shorter sentences

NLG may generate grammatically-incorrect sentences

Solution

- Generate word patterns in a order
- Consider linguistic patterns



Hierarchical NLG w/ Linguistic Patterns (Su et al., 2018)





Evolution Roadmap



Dialogue breadth (coverage)



Dialogue Systems







Neural Response Generation (Sordoni et al., 2015; Vinyals & Le, 2015)





Issue 1: Blandness Problem





Mutual Information for Neural Generation

• Mutual information objective

$$\hat{T} = \arg \max_{T} \left\{ \log \frac{p(S,T)}{p(S)p(T)} \right\}$$

$$\hat{T} = \arg \max_{T} \left\{ \log p(T|S) - \lambda \log p(T) \right\} \text{ Bayes' rule}$$

$$\operatorname{standard}_{\text{likelihood}} \text{ anti-LM}$$

$$\hat{T} = \arg \max_{T} \left\{ (1 - \lambda) \log p(T|S) + \lambda \log p(S|T) \right\}$$



Mutual Information for Neural Generation





Lexical Diversity

Material: http://opendialogue.miulab.tw **#** Distinct Tokens in Generated Targets 0.1 (Divided by Total #) 0.08 0.06 0.04 0.02 0 Standard LSTM LSTM + MMI LSTM + MMI (antiLM) (bidirectional)



MMI for Response Diversity (Li et al., 2016)





MMI for Response Diversity (Li et al., 2016)

Material: http://opendialogue.miulab.tw Wow sour starbursts really do make your mouth water... mm drool. Can I have one? Of course you can! They're delicious! Milan apparently selling Zlatan to balance the books... Where next, Madrid? I think he'd be a good signing. 'tis a fine brew on a day like this! Strong though, how many is sensible? Depends on how much you drink! Well he was on in Bromley a while ago... still touring I've never seen him live.



Material: http://opendialogue.miulab.tw

Real-World Conversations

Multimodality Conversation history yelp& context \odot Persona tripadvisor® encoder User profile data The SATORI (bio, social graph, etc.) Weather Channel Visual signal (camera, picture etc.) Because of gam ?? EOS Yeah ľm our Knowledge base Mood Geolocation Time



Issue 2: Response Inconsistency

message	Where do you live now?	message	How old are you?
response	I live in Los Angeles.	response	16 and you?
message	In which city do you live now?	message	What's your age?
response	I live in Madrid.	response	18.
message	In which country do you live now?	message	What is your major?
response	England, you?	response	I'm majoring in psychology
message	Where were you born?	message	What did you study in college?
response	I was born in Canada.	response	English lit.
message	Where are you from?		
response	England, you?		
message	Where did you grow up?		
response	I grew up in Texas.		



Personalized Response Generation (Li et al., 2016)





Persona Model for Speaker Consistency (Li et al., 2016)

Material: http://opendialogue.miulab.tw

Baseline model \rightarrow inconsistency

Persona model using speaker embedding \rightarrow consistency

		the second se	
message	Where is your hometown?	message	Where is your hometown?
response	I live in Indonesia.	response	I'm from England.
message	Where are you from?	message	Where are you from?
response	England, you?	response	I'm from England.
message	In which city do you live now?	message	In which city do you live now?
response	I live in Jakarta.	response	I live in London.
message	In which country do you live now?	message	In which country do you live now?
response	I live in Indonesia.	response	I live in England.
message	What is your major?	message	What is your major?
response	Business, you?	response	Business. You?
message	What did you study in college?	message	What did you study in college?
response	Psychology, you?	response	I did business studies.
message	How old are you?	message	How old are you?
response	I'm 18.	response	I'm 18.
message	Wh <mark>at is</mark> your age?	message	What is your age?
response	I'm <mark>16.</mark>	response	I'm 18.



Speak-Role Aware Response (Luan et al., 2017)

65





Speak-Role Aware Response (Luan et al., 2017)

Material: http://opendialogue.miulab.tw

- Speaker role: support person
 - Domain expertise: technical
 - □ Speaking style: polite

User input: I am getting a loop back to login page. **Baseline model:** Ah, ok. Thanks for the info. **Our model:** I'm sorry to hear that. Have you tried clearing your cache and cookies?

Context: have you tried uninstalling nike+ connec? how about docking it to a different computer? Input: Tried both of those techniques. But nothing works. I can try the other computer again. Baseline: Let me know if there's anything else I can help you with. Thanks for the help. Multitask: Have you tried using a different browser? Let us know if you have any other questions.

Context: What do you see when you have a failed login attempt? Never ending spinner? Loop you back to login page? *Input: I am getting a loop back to login page, windows.* **Baseline:** Ah, ok. Thanks for the info. I'll check it out. Thanks for the feedback. **Multitask:** I'm sorry to hear that. Have you tried clearing your cache and cookies?

Context: To clarify: your first run after resetting your connection will not post to Facebook. Does this help? Let us know. Input: That's correct. I reset it twice! It still doesn't work. Baseline: Let me know if there's anything else I can help you with. Thanks for the update. Multitask: I'm sorry to hear that. Are you receiving any error messages? Let us know.



Issue 3: Dialogue-Level Optimization via RL





Deep RL for Response Generation (Li et al., 2016)

Input message	Supervised Learning Agent	Reinforcement Learning Agent
How old are you ?	I'm 16.	I'm 16, why are you asking ?
What's your full name ?	I have no idea .	What's yours ?
I don't want to go home tonight.	Really ?	Why?
Do you have any feelings for me?	I don't know what you are talking about.	I don 't want to hurt your feelings .
How much time do you have here?	Not long enough. Sorry, sir.	Ten seconds .
Shall we get started ?	Of course !	Yes. We 've got a lot of work to do here.
Do you play football ?	No, i don't	Yes. I love football !
We'd have to talk to him.	I mean, he's a good guy	About what ?
How come you never say it?	Because I don't want to hurt you .	I don 't think it 's a good idea to say it.

- RL agent generates more *interactive* responses
- RL agent tends to end a sentence with a question and hand the conversation over to the user



Issue 4: No Grounding (Sordoni et al., 2015; Li et al., 2016)

Material: http://opendialogue.miulab.tw

Neural model learns the general shape of conversations, and the system output is situationally <u>appropriate</u> and <u>coherent</u>. No grounding into a real calendar, but the "shape" of the conversation is fluent and plausible.

H: hi
M: how are you?
H: not bad
M: what's wrong?
H: nothing really
M: wanna talk about it? i 'm here if you wanna talk
H: well, i'm just tired
M: me too, i'm here if you wanna talk
H: would thursday
M: yeah , sure . just
H: after lunch is provided in the second seco

H: would thursday afternoon work sometime?
M: yeah , sure . just let me know when you're free.
H: after lunch is probably the best time
M: okay, sounds good . just let me know when you're free.
H: would 2 pm work for you?
M: works for me.
H: well let's say 2 pm then i 'll see you there
M: sounds good.









Image-Grounded Dialogues (Mostafazadeh et al., 2017)

Material: http://opendialogue.miulab.tw

Neural model, trained on 250K Twitter convos:

- -Turn 1: tweet + accompanying photo
- -Turn 2: tweet is a question

Given data, other signals can be exploited in the same way: program state, database info, geolocation, etc.



Tweet: some of the damage done NeuroCon: Is that at your house?





Tweet: as penitence I offer this NeuroCon: can I have one of those?

Tweet: my poor baby moms saying it 's probably totaled so rip Bessie NeuroCon: **How did you manage to do that?**



Knowledge-Grounded Responses (Ghazvininejad et al., 2017)




Conversation and Non-Conversation Data

Material: http://opendialogue.miulab.tw





Knowledge-Grounded Responses (Ghazvininejad et al., 2017)

Material: http://opendialogue.miulab.tw

A: Visiting the celebs at Los Angeles International Airport (LAX) - [...] w/ 70 others B: Nice airport terminal. Have a safe flight.

A: Is that [...] in your photos? It's on my list of places to visit in NYC. B: Don't forget to check out the 5th floor, while you are here, it's a great view.

A: Live right now on [...] Tune in!!!!! B: Listen to Lisa Paige

A: Been craving Chicken Pot Pie-who has the best? Trying [...] at [...] Must be Change of weather! B: Love the pasta trattoria.

A: So [...] is down to one copy of Pound Foolish. I'm curious to see if they are re-ordering it.

B: Check out the video feed on 6 and take a picture of the Simpsons on the 3rd floor.

A: I wish [...] would introduce another vegetarian option besides the shroomburger. It's delicious but kind of ridiculous. B: This is the best j.crew in the world. Try the lemonade!

A: Just had an awesome dinner at [...] Great recommendation [...]

B: One of my favorite places I've ever been to in NYC. The food is great and the service is lackluster.

Results (23M conversations) outperforms competitive neural baseline (human + automatic eval)



Evolution Roadmap



Dialogue breadth (coverage)



Material: http://opendialogue.miulab.tw

Multimodality & Personalization (Chen et al., 2018)

- Task: user intent prediction
- □ Challenge: language ambiguity

O User preference

- Some people prefer "Message" to "Email"
- Some people prefer "Ping" to "Text"

2 App-level contexts

- "Message" is more likely to follow "Camera"
- "Email" is more likely to follow "Excel"

Behavioral patterns in history helps intent prediction.



send to vivian Communication



High-Level Intention Learning (Sun et al., 2016; Sun et al., 2016)

Material: http://opendialogue.miulab.tw

High-level intention may span several domains



Users interact via high-level descriptions and the system learns how to plan the dialogues



Empathy in Dialogue System (Fung et al., 2016)

- Embed an empathy module
 - Recognize emotion using multimodality
 - Generate emotion-aware responses





Material: http://opendialogue.miulab.tw



Made with love by tvo Technologies in collaboration with Hong Kong University of Science and Technology



Face recognition output

(index):1728



Cognitive Behavioral Therapy (CBT)

Material: http://opendialogue.miulab.tw







Challenge Summary

81

Material: http://opendialogue.miulab.tw

The human-machine interface is a hot topic but several components must be integrated!

- Most state-of-the-art technologies are based on DNN
- •Requires huge amounts of labeled data
- •Several frameworks/models are available

Fast domain adaptation with scarse data + re-use of rules/knowledge

Handling reasoning and personalization

Data collection and analysis from un-structured data

Complex-cascade systems requires high accuracy for working good as a whole

Her (2013)

What can machines achieve now or in the future?

⁸³ Thanks for Your Attention!

Q & A



Yun-Nung (Vivian) Chen Assistant Professor National Taiwan University y.v.chen@ieee.org / http://vivianchen.idv.tw