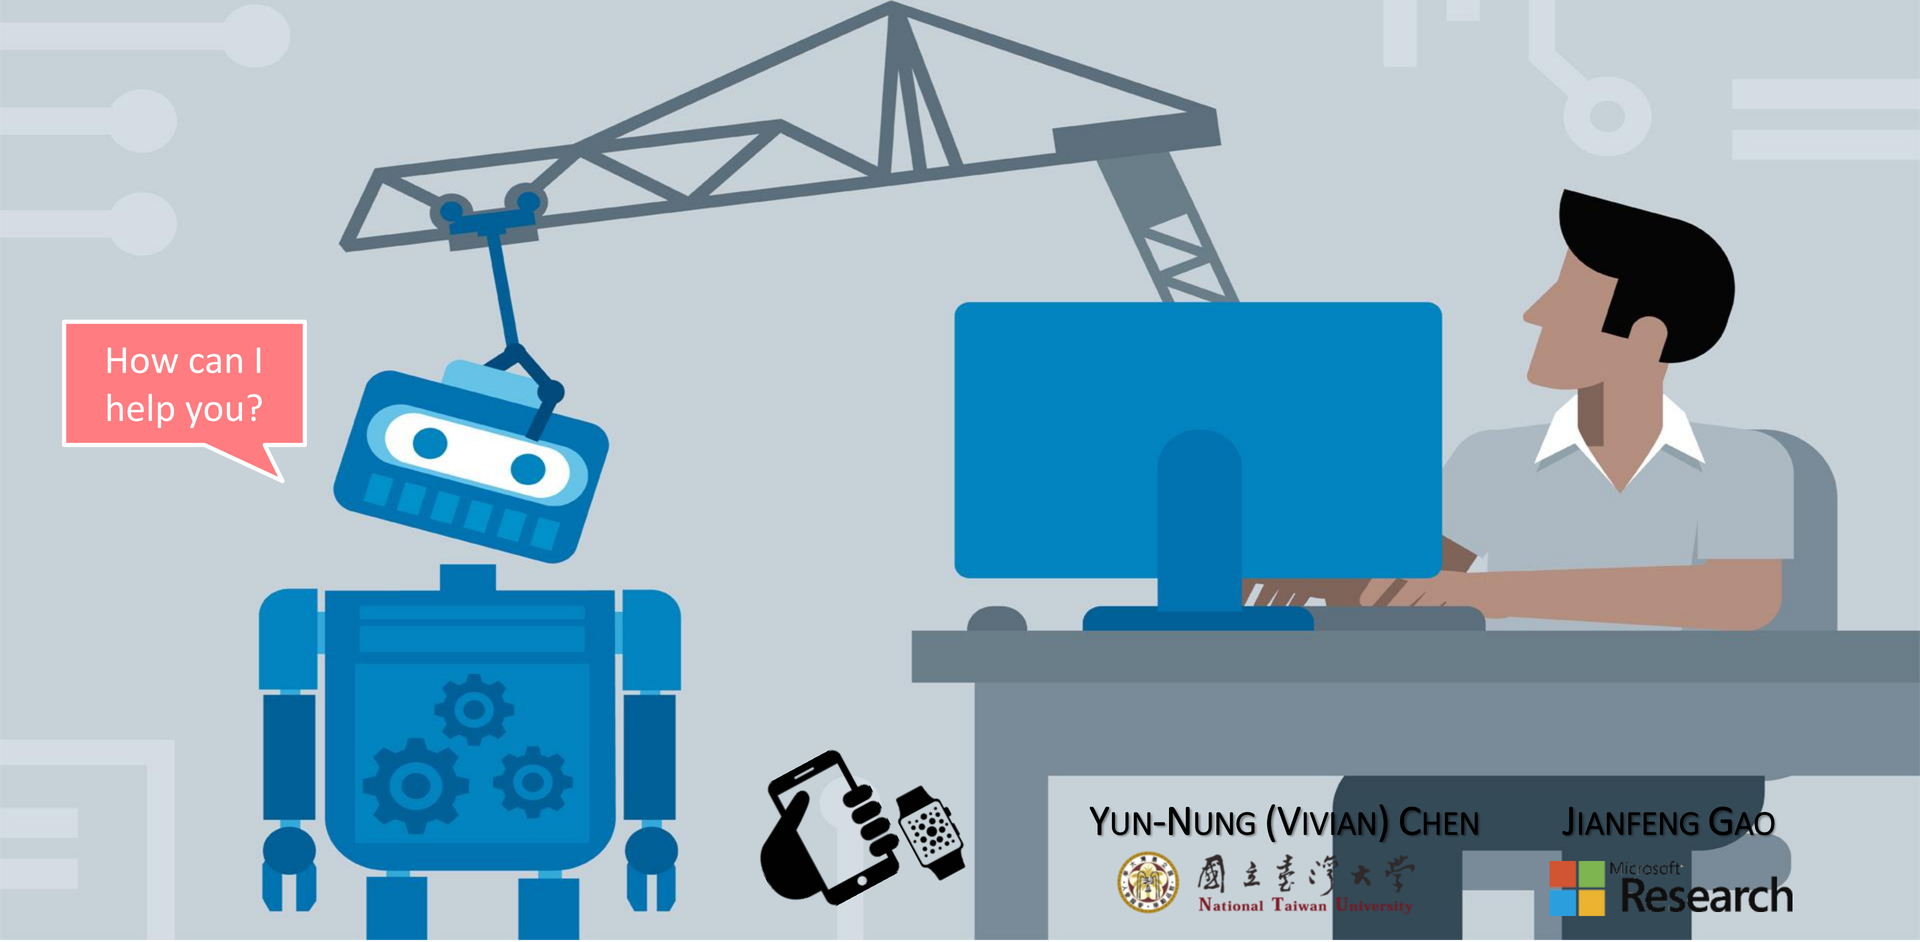

YUN-NUNG (VIVIAN) CHEN
國立臺灣大學 National Taiwan University

JIANFENG GAO
Microsoft Research

Open-Domain Neural Dialogue Systems

opendialogue.miulab.tw

Outline

- **PART I.** Introduction & Background Knowledge
- **PART II.** Task-Oriented Dialogue Systems

Break

- **PART III.** Social Chat Bots
- **PART IV.** Evaluation
- **PART V.** Recent Trends and Challenges

Introduction & Background Knowledge

Introduction

Outline

- **PART I. Introduction & Background Knowledge**
  - ***Dialogue System Introduction***
  - Neural Network Basics
  - Reinforcement Learning
- PART II. Task-Oriented Dialogue Systems
- PART III. Social Chat Bots
- PART IV. Evaluation
- PART V. Recent Trends and Challenges

Brief History of Dialogue Systems

Material: http://opendialogue.miulab.tw

Multi-modal systems
e.g., Microsoft MiPad, Pocket PC

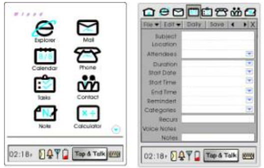

TV Voice Search
e.g., Bing on Xbox

Virtual Personal Assistants

Task-specific argument extraction
(e.g., Nuance, SpeechWorks)
User: "I want to fly from Boston to New York next week."

2017

Early 2000s

IBM WATSON

Early 1990s

Intent Determination
(Nuance's Emily™, AT&T HMIHY)
User: "Uh…we want to move…we want to change our phone line from this house to another house"

DARPA
CALO Project

Keyword Spotting
(e.g., AT&T)
System: "Please say collect, calling card, person, third number, or operator"

Why We Need?

"I am smart"	Turing Test ("I" talk like a human)
"I have a question"	Information consumption
"I need to get this done"	Task completion
"What should I do?"	Decision support

Why We Need?

Material: http://opendialogue.miulab.tw

"I am smart"	Turing Test ("I" talk like a human)
"I have a question"	Information consumption
"I need to get this done"	Task completion
"What should I do?"	Decision support

- *What is the employee review schedule?*
- *Which room is the dialogue tutorial in?*
- *When is the IJCNLP 2017 conference?*
- *What does NLP stand for?*

Why We Need?

"I am smart"	Turing Test ("I" talk like a human)
"I have a question"	Information consumption
"I need to get this done"	Task completion
"What should I do?"	Decision support

- *Book me the flight from Seattle to Taipei*
- *Reserve a table at Din Tai Fung for 5 people, 7PM tonight*
- *Schedule a meeting with Bill at 10:00 tomorrow.*

Why We Need?

"I am smart" — Turing Test ("I" talk like a human)

"I have a question" — Information consumption

"I need to get this done" — Task completion

"What should I do?" — Decision support

- *Is this product worth to buy?*

Why We Need?

"I am smart"

"I have a question"

"I need to get this done"

"What should I do?"

Turing Test ("I" talk like a human)

Information consumption

Task completion

Decision support

Task-Oriented Dialogues

Language Empowering Intelligent Assistant

Material: http://opendialogue.miulab.tw

Apple Siri (2011)

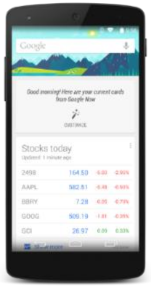

Google Now (2012)
Google Assistant (2016)

Microsoft Cortana (2014)

Amazon Alexa/Echo (2014)

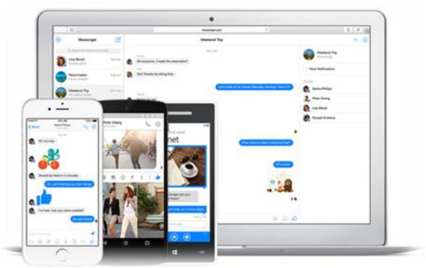

Facebook M & Bot (2015)

Google Home (2016)

Apple HomePod (2017)

Intelligent Assistants

Google Now Siri Cortana echo

Material: http://opendialogue.miulab.tw

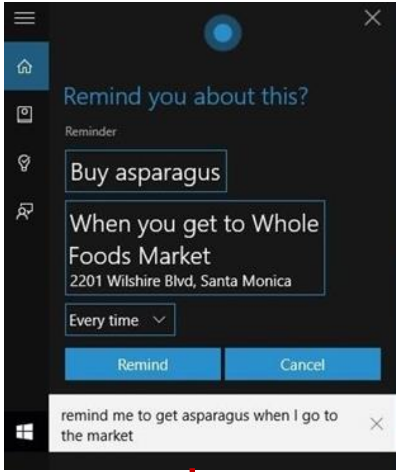

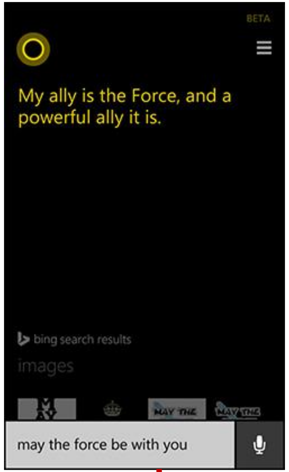

Task-Oriented

Engaging (social bots)

Why Natural Language?

Material: http://opendialogue.miulab.tw

- Global Digital Statistics (2017 January)

Total Population
7.48B

Internet Users
3.77B

Active Social Media Users
2.79B

Unique Mobile Users
4.92B

Active Mobile Social Users
2.55B

The more **natural** and **convenient** input of devices evolves towards speech.

Spoken Dialogue System (SDS)

Material: http://opendialogue.miulab.tw

- Spoken dialogue systems are intelligent agents that are able to help users finish tasks more efficiently via spoken interactions.
- Spoken dialogue systems are being incorporated into various devices (smart-phones, smart TVs, in-car navigating system, etc).

JARVIS – Iron Man's Personal Assistant

Baymax – Personal Healthcare Companion

Good dialogue systems assist users to access information conveniently and finish tasks efficiently.

App → Bot

Material: http://opendialogue.miulab.tw

- A bot is responsible for a “single” domain, similar to an app

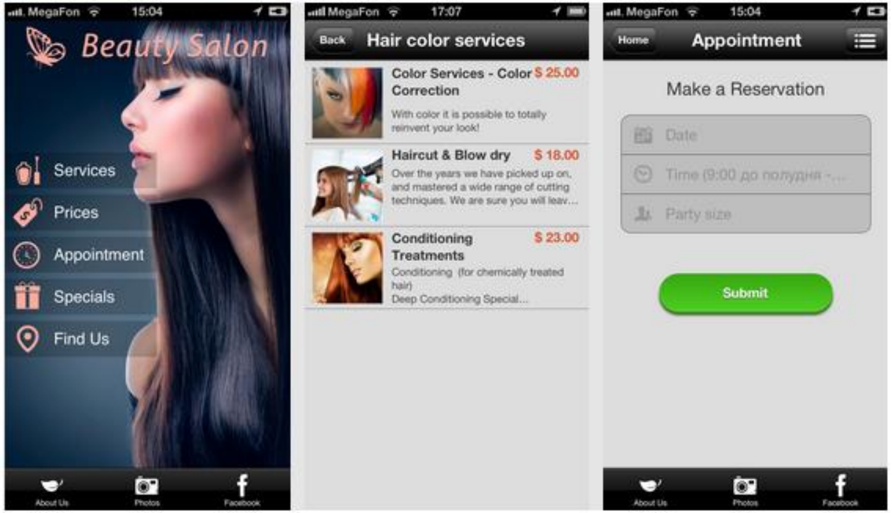

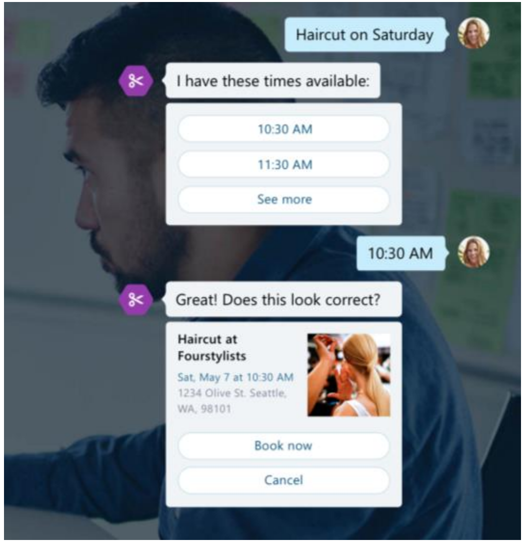

Users can initiate dialogues instead of following the GUI design

GUI v.s. CUI (Conversational UI)

https://github.com/enginebai/Movie-lol-android

Material: http://opendialogue.miulab.tw

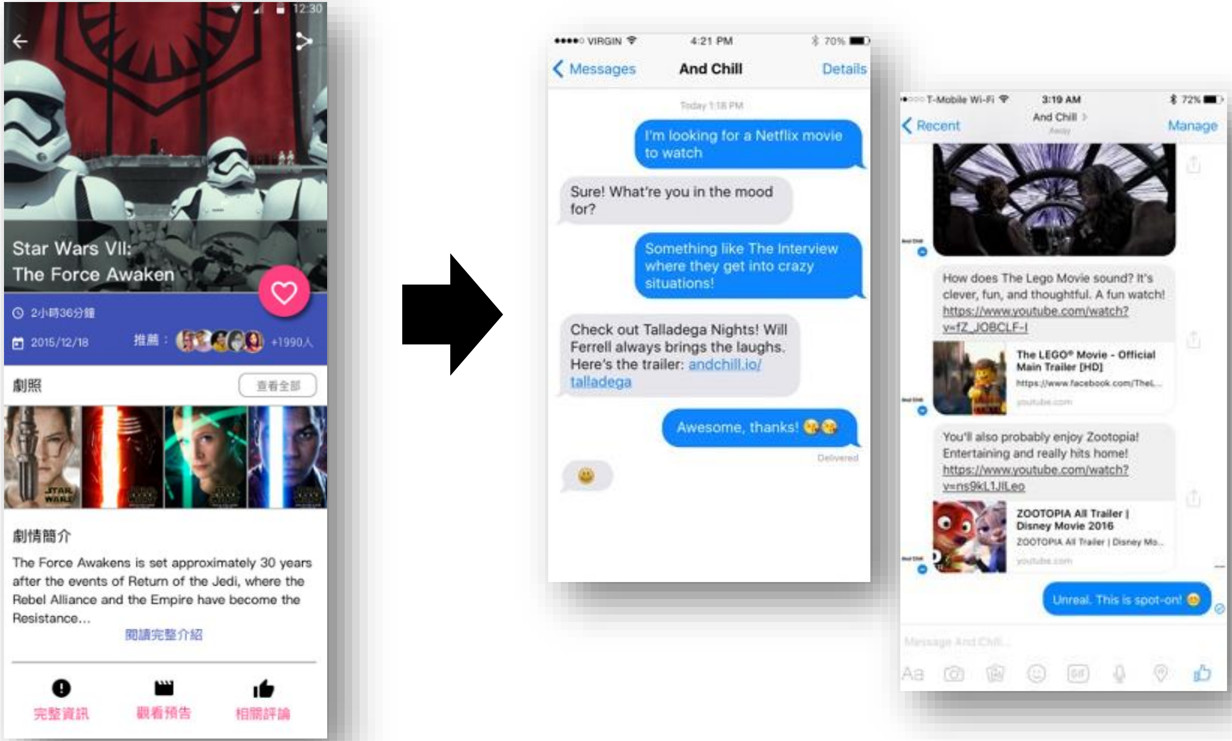

GUI v.s. CUI (Conversational UI)

	Website/APP's GUI	Msg's CUI
Situation	Navigation, no specific goal	Searching, with specific goal
Information Quantity	More	Less
Information Precision	Low	High
Display	Structured	Non-structured
Interface	Graphics	Language
Manipulation	Click	mainly use texts or speech as input
Learning	Need time to learn and adapt	No need to learn
Entrance	App download	Incorporated in any msg-based interface
Flexibility	Low, like machine manipulation	High, like converse with a human

Two Branches of Dialogue Systems

Task-Oriented Bot

- Personal assistant, helps users achieve a certain task
- Combination of rules and statistical components
  - POMDP for spoken dialog systems (Williams and Young, 2007)
  - End-to-end trainable task-oriented dialogue system (Wen et al., 2016)
  - End-to-end reinforcement learning dialogue system (Li et al., 2017; Zhao and Eskenazi, 2016)

Chit-Chat Bot

- No specific goal, focus on natural responses
- Using variants of seq2seq model
  - A neural conversation model (Vinyals and Le, 2015)
  - Reinforcement learning for dialogue generation (Li et al., 2016)
  - Conversational contextual cues for response ranking (Al-Rfou et al., 2016)

Task-Oriented Dialogue System (Young, 2000)

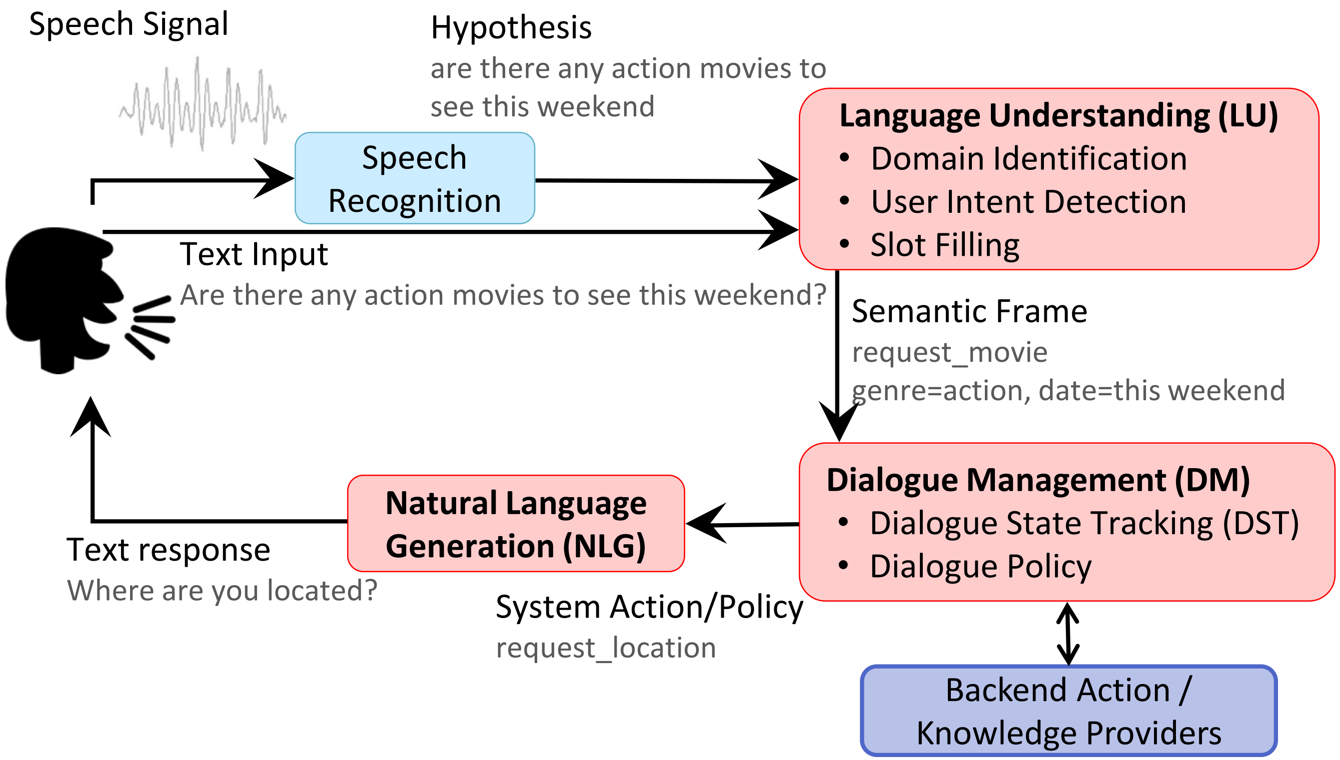

Outline

- **PART I. Introduction & Background Knowledge**
  - Dialogue System Introduction
  - ***Neural Network Basics***
  - Reinforcement Learning
- PART II. Task-Oriented Dialogue Systems
- PART III. Social Chat Bots
- PART IV. Evaluation
- PART V. Recent Trends and Challenges

Machine Learning ≈ Looking for a Function

- Speech Recognition

$f($ 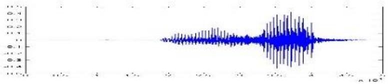 $) =$ "你好 (Hello) "

- Image Recognition

$f($  $) =$ cat

- Go Playing

$f($  $) =$ 5-5 (next move)

- Chat Bot

$f($ "Where is IJCNLP?" $) =$ "The location is…"

Given a large amount of data, the machine learns what the function ***f*** should be.

Machine Learning

Material: http://opendialogue.miulab.tw

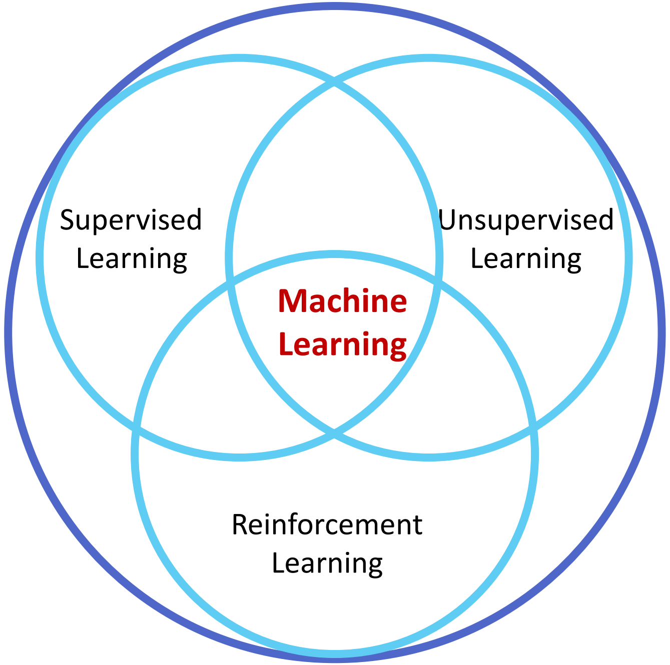

Deep learning is a type of machine learning approaches, called “neural networks”.

A Single Neuron

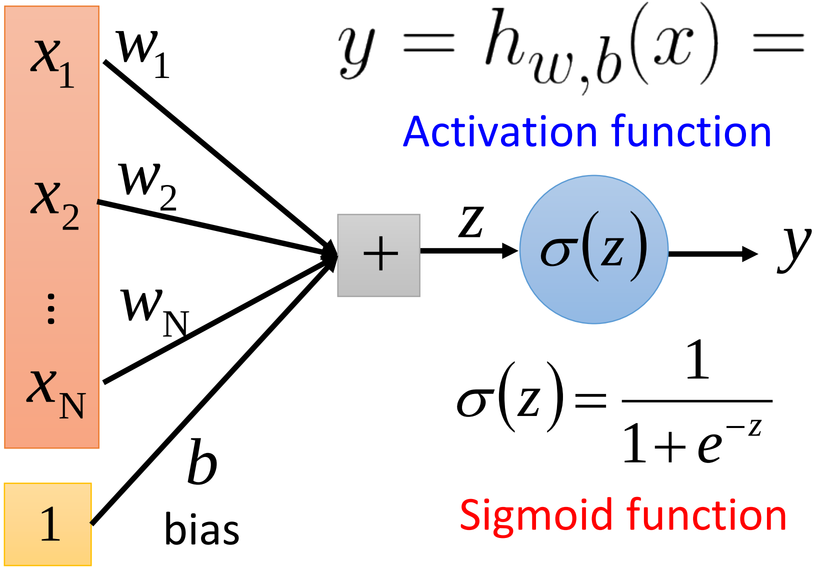

$$y = h_{w,b}(x) = \sigma(w^T x + b)$$

Activation function

$$\sigma(z) = \frac{1}{1+e^{-z}}$$

Sigmoid function

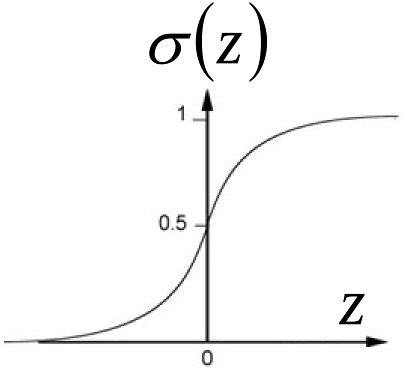

w, b are the parameters of this neuron

A Single Neuron

$$f : R^N \rightarrow R^M$$

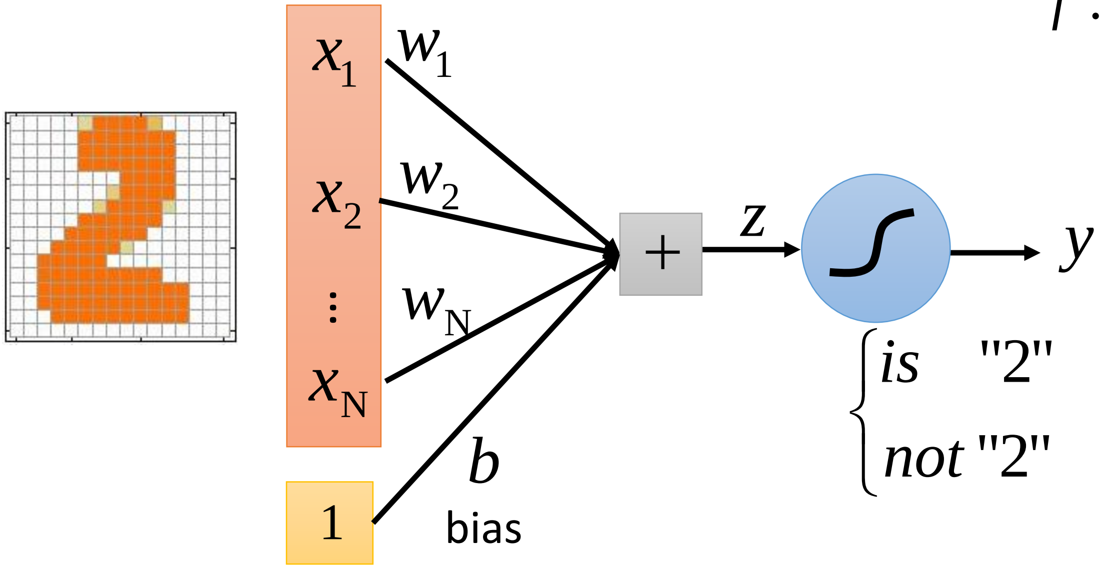

$$\begin{cases} is \quad "2" & y \geq 0.5 \\ not \ "2" & y < 0.5 \end{cases}$$

A single neuron can only handle binary classification

A Layer of Neurons

Material: http://opendialogue.miulab.tw

- Handwriting digit classification

$$f : R^N \rightarrow R^M$$

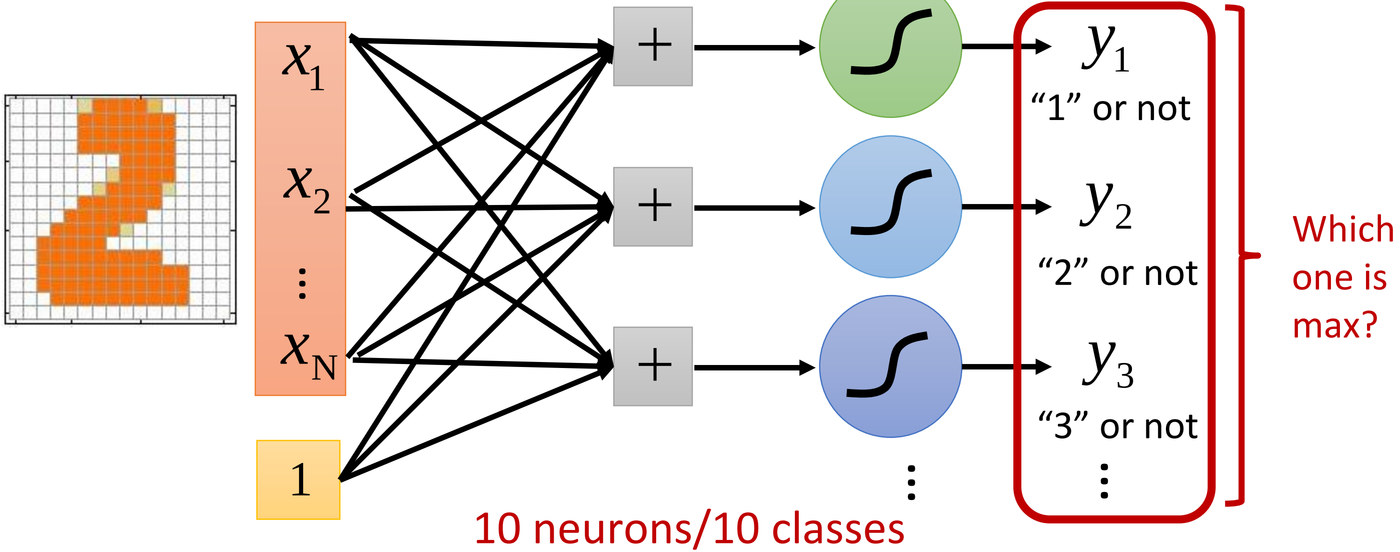

A layer of neurons can handle multiple possible output, and the result depends on the max one

Deep Neural Networks (DNN)

Material: http://opendialogue.miulab.tw

- Fully connected feedforward network

$$f : R^N \rightarrow R^M$$

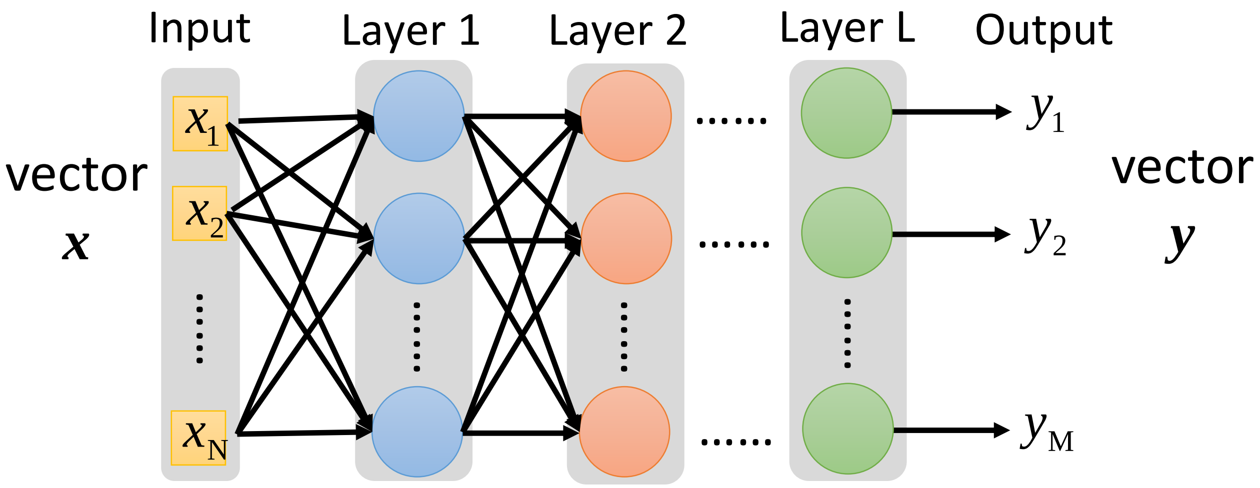

Deep NN: multiple hidden layers

Recurrent Neural Network (RNN)

$$s_t = \sigma(W s_{t-1} + U x_t)$$

$\sigma(\cdot)$: tanh, ReLU

$$o_t = \mathrm{softmax}(V s_t)$$

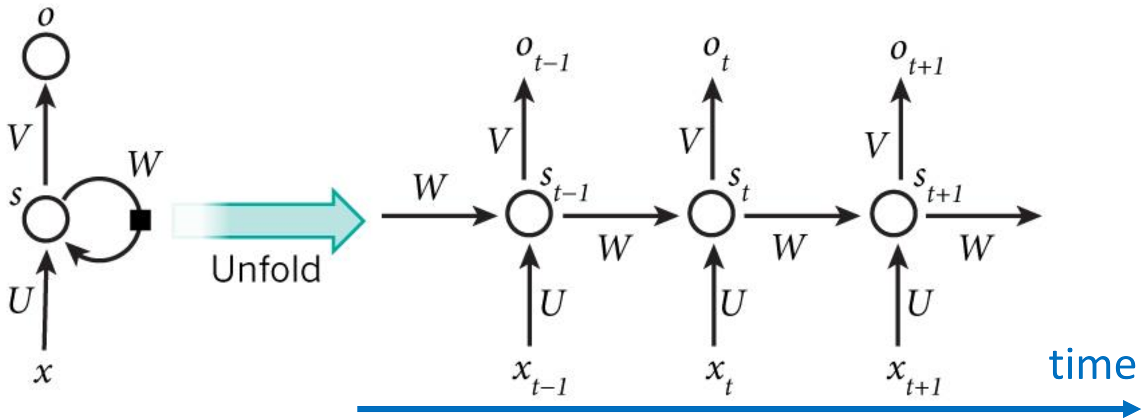

RNN can learn accumulated sequential information (time-series)

http://www.wildml.com/2015/09/recurrent-neural-networks-tutorial-part-1-introduction-to-rnns/

Outline

Reinforcement Learning

- RL is a general purpose framework for **decision making**
  - RL is for an *agent* with the capacity to *act*
  - Each *action* influences the agent's future *state*
  - Success is measured by a scalar *reward* signal
  - Goal: *select actions to maximize future reward*

Scenario of Reinforcement Learning

Material: http://opendialogue.miulab.tw

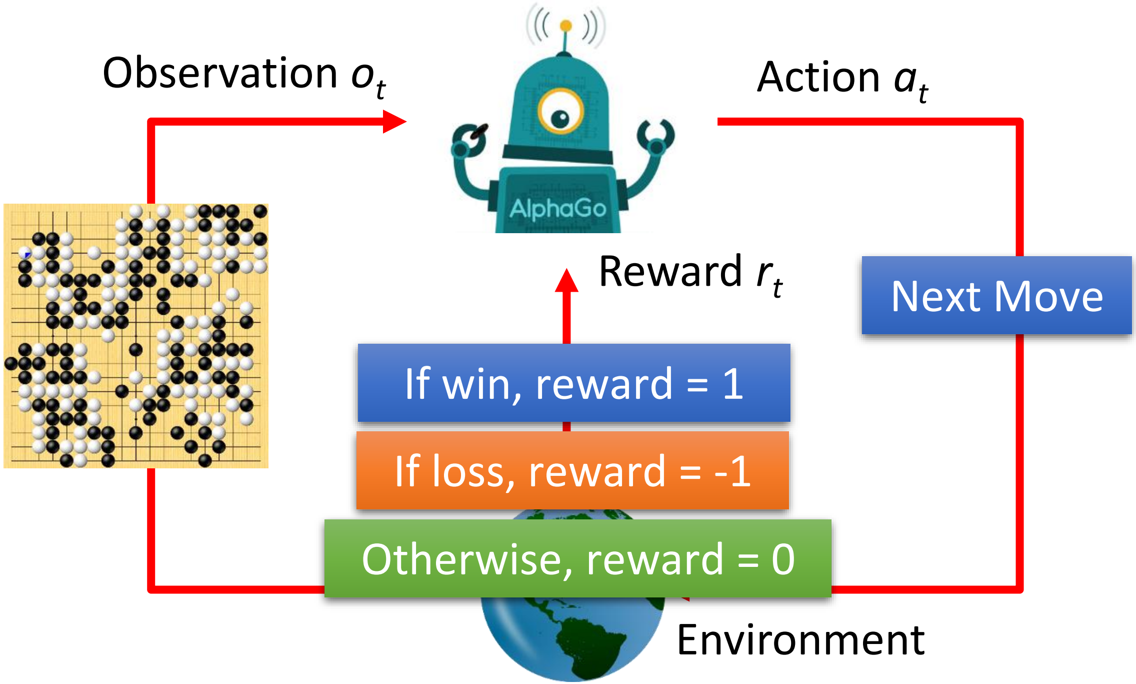

Agent learns to take actions to maximize expected reward.

Supervised v.s. Reinforcement

- Supervised

Learning from teacher

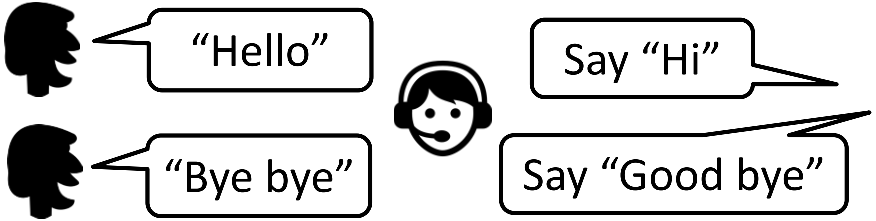

- Reinforcement

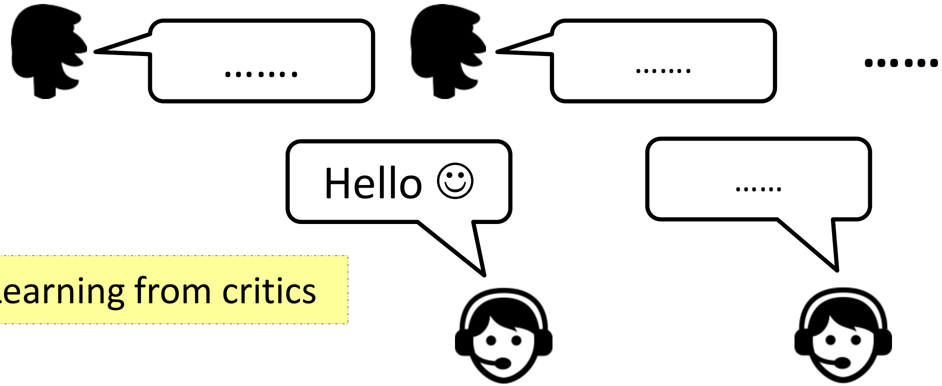

Bad

Learning from critics

Sequential Decision Making

Material: http://opendialogue.miulab.tw

- Goal: select actions to maximize total future reward
  - Actions may have long-term consequences
  - Reward may be delayed
  - It may be better to sacrifice immediate reward to gain more long-term reward

Deep Reinforcement Learning

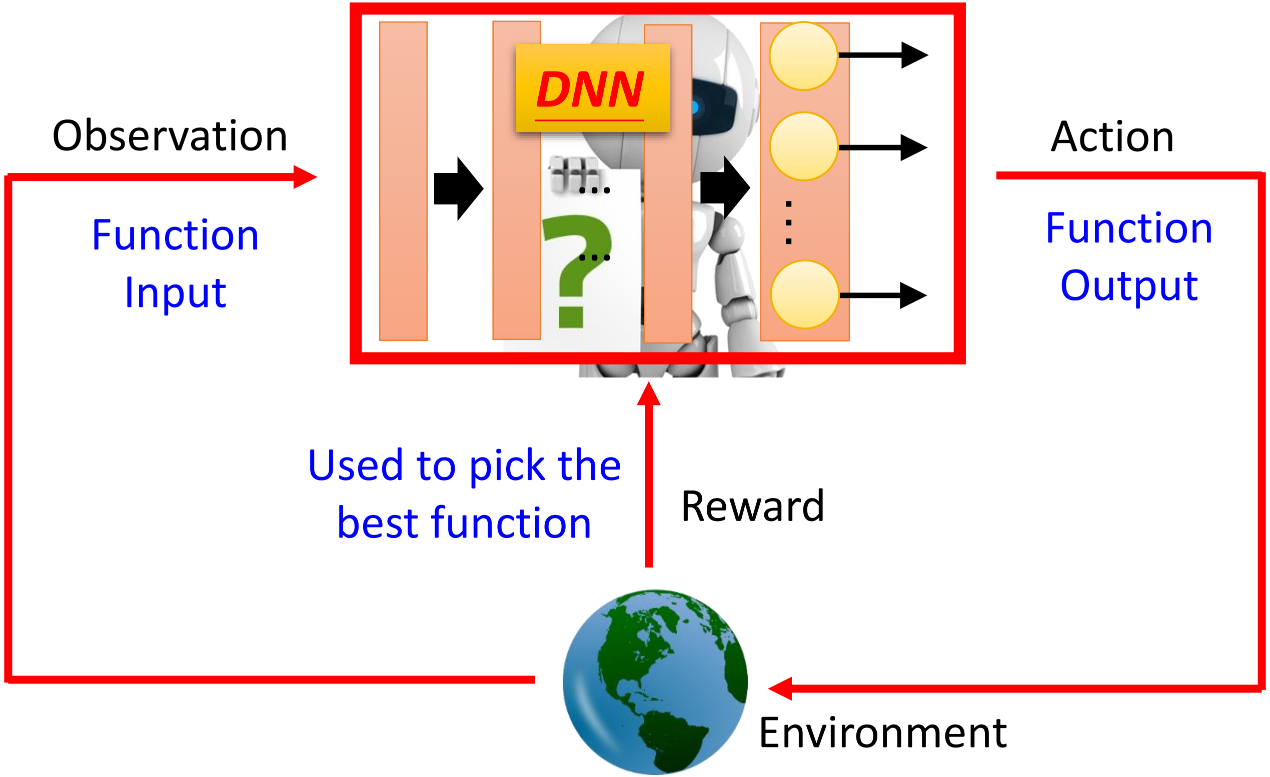

Reinforcing Learning

- Start from state s_0
- Choose action a_0
- Transit to $s_1 \sim P(s_0, a_0)$
- Continue…

$$s_0 \xrightarrow{a_0} s_1 \xrightarrow{a_1} s_2 \xrightarrow{a_2} s_3 \xrightarrow{a_3} \dots$$

- Total reward:

$$R(s_0) + \gamma R(s_1) + \gamma^2 R(s_2) + \cdots$$

Goal: select actions that maximize the expected total reward

$$\mathbb{E}[R(s_0) + \gamma R(s_1) + \gamma^2 R(s_2) + \cdots]$$

Reinforcement Learning Approach

- Policy-based RL
  - Search directly for optimal policy π^*

π^* is the policy achieving maximum future reward

- Value-based RL
  - Estimate the optimal value function $Q^*(s, a)$

$Q^*(s, a)$ is maximum value achievable under any policy

- Model-based RL
  - Build a model of the environment
  - Plan (e.g. by lookahead) using model

Task-Oriented Dialogue Systems

Task-Oriented Dialogue System (Young, 2000)

Material: http://opendialogue.miulab.tw

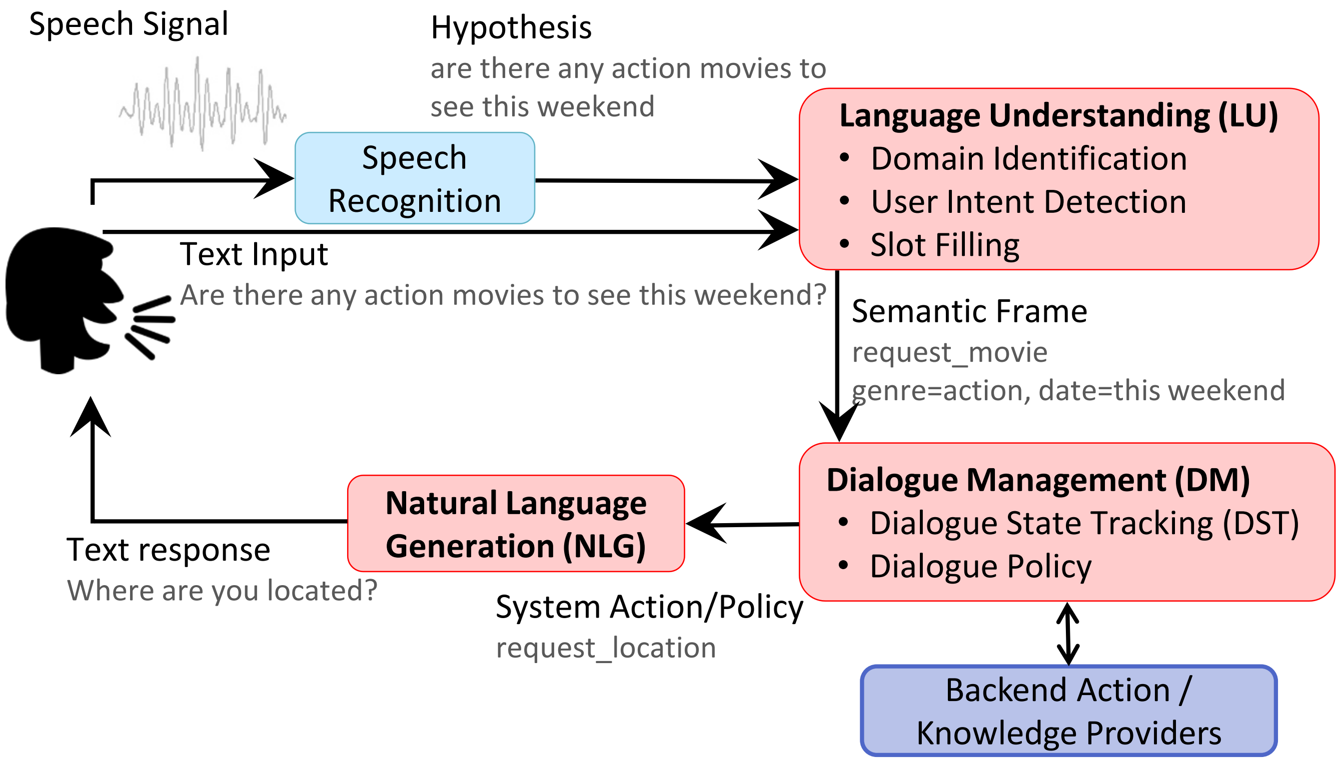

Outline

- PART I. Introduction & Background Knowledge
- **PART II. Task-Oriented Dialogue Systems**
  - ***Spoken/Natural Language Understanding (SLU/NLU)***
  - Dialogue Management – Dialogue State Tracking (DST)
  - Dialogue Management – Dialogue Policy Optimization
  - Natural Language Generation (NLG)
  - End-to-End Task-Oriented Dialogue Systems
- PART III. Social Chat Bots
- PART IV. Evaluation
- PART V. Recent Trends and Challenges

Language Understanding (LU)

Material: http://opendialogue.miulab.tw

- Pipelined

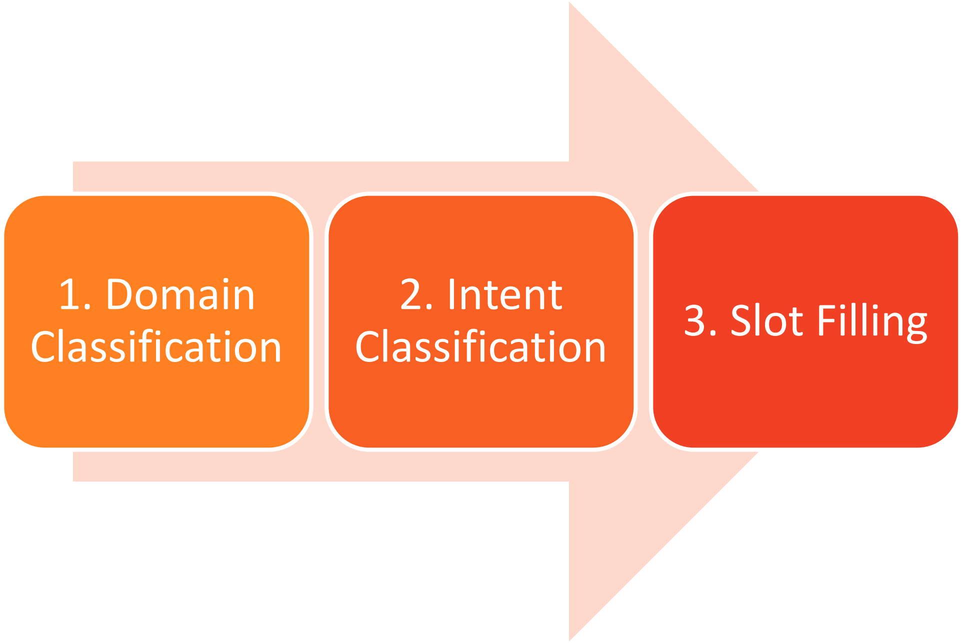

LU – Domain/Intent Classification

Mainly viewed as an utterance classification task

- Given a collection of utterances u_i with labels c_i, $D = \{(u_1,c_1),\dots,(u_n,c_n)\}$ where $c_i \in C$, train a model to estimate labels for new utterances u_k.

find me a cheap taiwanese restaurant in oakland

Movies	Find_movie
Restaurants	Buy_tickets
Sports	Find_restaurant
Weather	Book_table
Music	Find_lyrics
…	…

DNN for Domain/Intent Classification (Ravuri & Stolcke, 2015)

- RNN and LSTMs for utterance classification

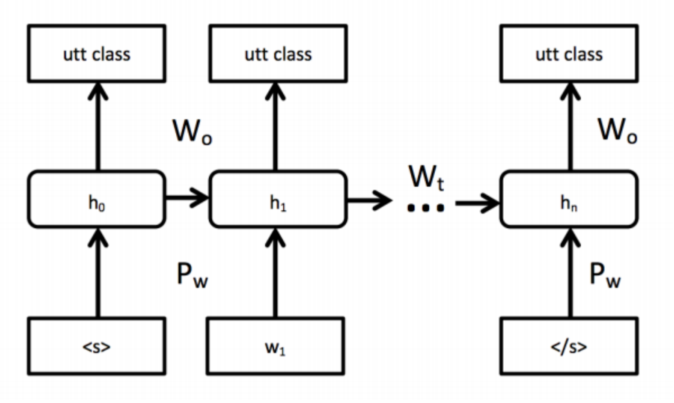

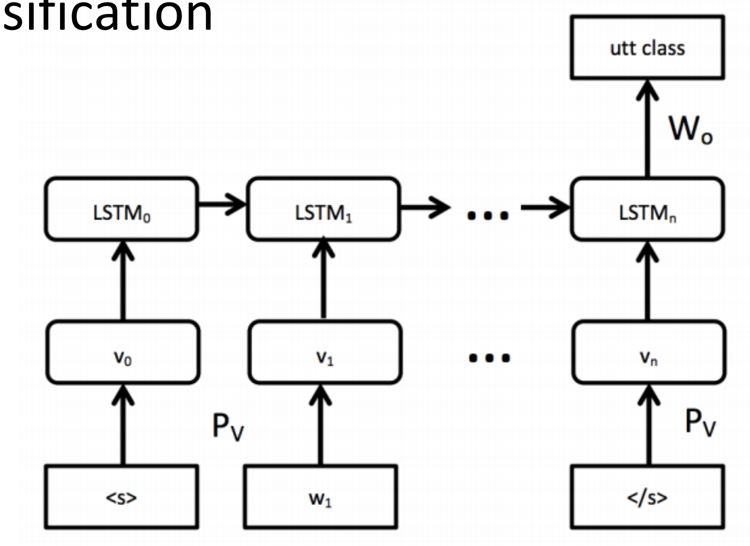

Intent decision after reading all words performs better

DNN for Dialogue Act Classification (Lee & Dernoncourt, 2016)

Material: http://opendialogue.miulab.tw

- RNN and CNNs for dialogue act classification

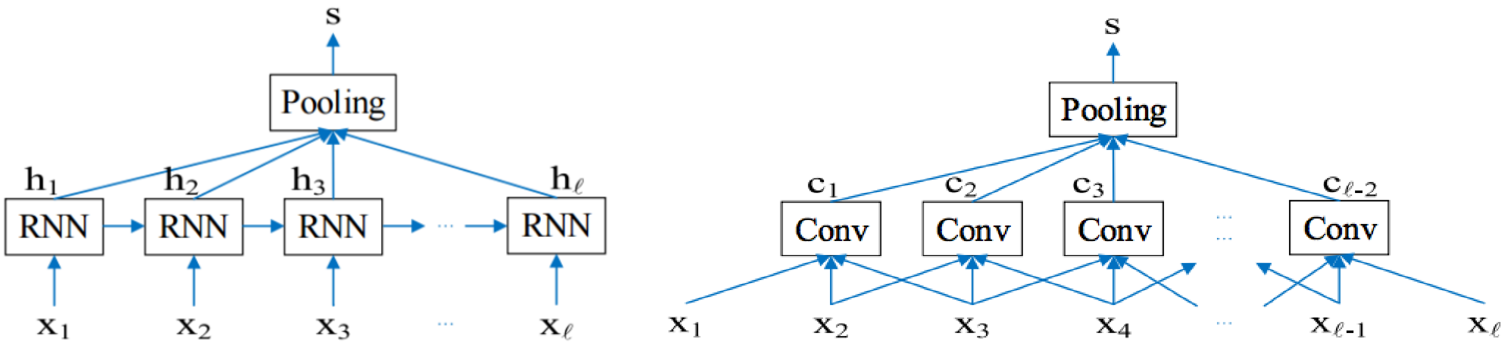

LU – Slot Filling

As a sequence tagging task

- Given a collection tagged word sequences, $S=\{((w_{1,1},w_{1,2},\ldots, w_{1,n1}), (t_{1,1},t_{1,2},\ldots,t_{1,n1})), ((w_{2,1},w_{2,2},\ldots,w_{2,n2}), (t_{2,1},t_{2,2},\ldots,t_{2,n2})) \ldots\}$ where $t_i \in M$, the goal is to estimate tags for a new word sequence.

flights from Boston to New York today

	flights	from	Boston	to	New	York	today
Entity Tag	O	O	B-city	O	B-city	I-city	O
Slot Tag	O	O	B-dept	O	B-arrival	I-arrival	B-date

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

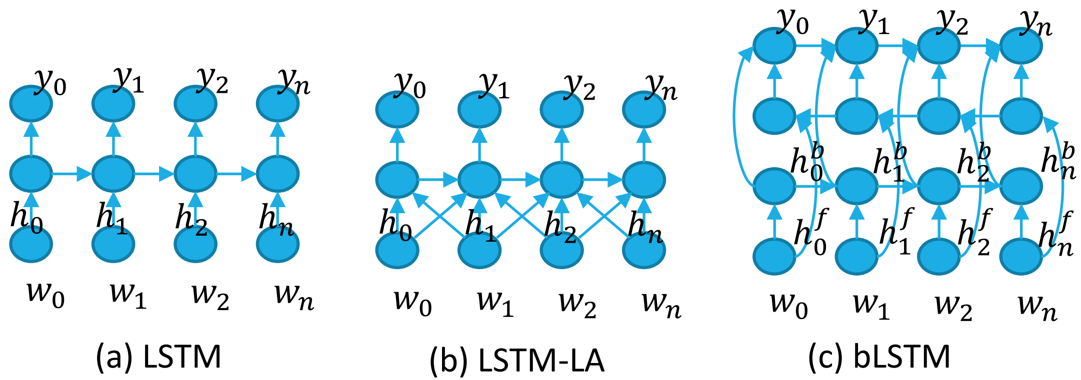

RNN for Slot Tagging – II (Kurata et al., 2016; Simonnet et al., 2015)

- Encoder-decoder networks
  - Leverages sentence level information
- Attention-based encoder-decoder
  - Use of attention (as in MT) in the encoder-decoder network
  - Attention is estimated using a feed-forward network with input: h_t and s_t at time t

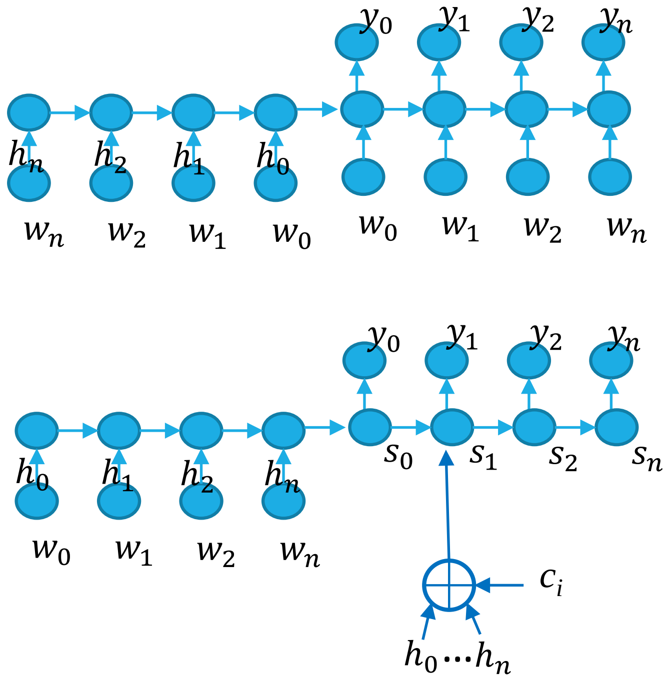

RNN for Slot Tagging – III (Jaech et al., 2016; Tafforeau et al., 2016)

Material: http://opendialogue.miulab.tw

- Multi-task learning
  - Goal: exploit data from domains/tasks with a lot of data to improve ones with less data
  - Lower layers are shared across domains/tasks
  - Output layer is specific to task

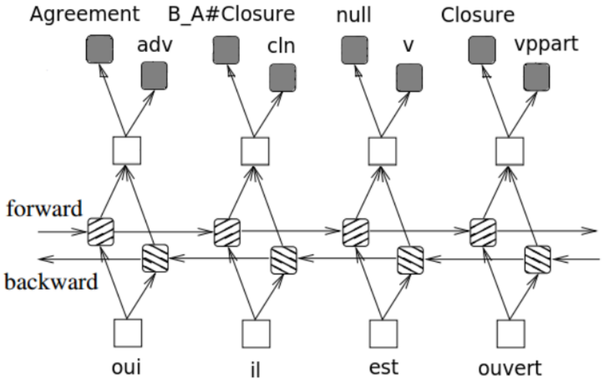

Joint Segmentation and Slot Tagging (Zhai et al., 2017)

- Encoder that segments
- Decoder that tags the segments

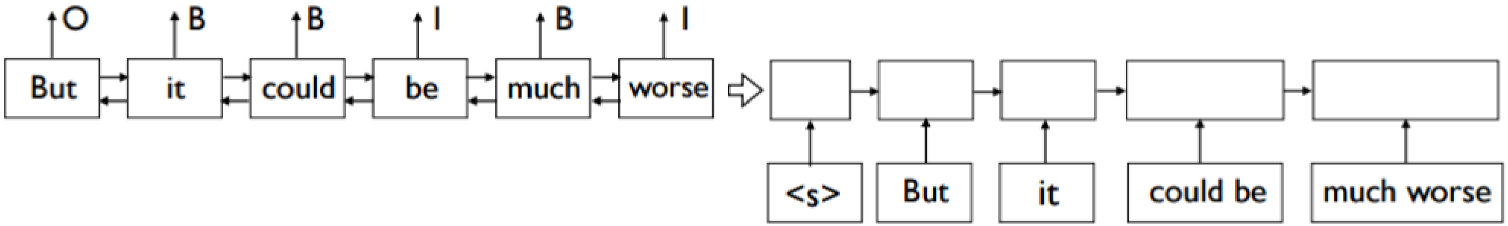

Joint Semantic Frame Parsing

Sequence-based (Hakkani-Tur+ 16)

- Slot filling and intent prediction in the same output sequence

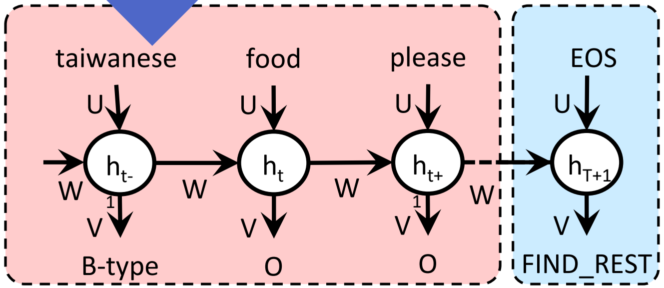

Slot Filling Intent Prediction

Parallel-based (Liu+ 16)

- Intent prediction and slot filling are performed in two branches

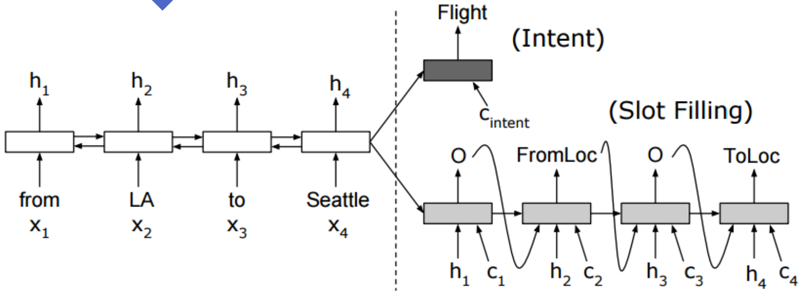

Contextual LU

Domain Identification → **Intent Prediction** → **Slot Filling**

D communication *I* send_email

Single Turn

U	just	sent	email	to	bob	about	fishing	this	weekend
S	O	O	O	O	B-contact_name	O	B-subject	I-subject	I-subject

→ send_email(contact_name="bob", subject="fishing this weekend")

Multi-Turn

U_1	send	email	to	bob
S_1				B-contact_name

→ send_email(contact_name="bob")

U_2	are	we	going	to	fish	this	weekend
S_2	B-message	I-message	I-message	I-message	I-message	I-message	I-message

→ send_email(message="are we going to fish this weekend")

Contextual LU

- User utterances are highly ambiguous in isolation

Restaurant Booking

Book a table for 10 people tonight.

Which restaurant would you like to book a table for?

Cascal, for 6.

?

#people time

Contextual LU (Bhargava et al., 2013; Hori et al, 2015)

- Leveraging contexts
  - Used for individual tasks
- Seq2Seq model
  - Words are input one at a time, tags are output at the end of each utterance

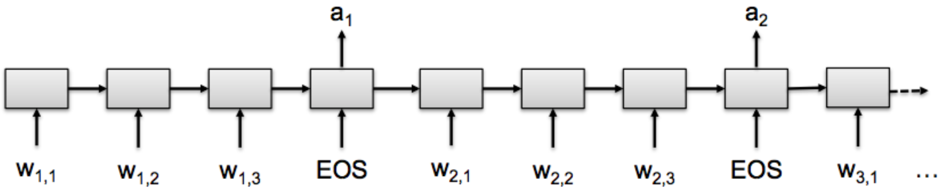

- Extension: LSTM with speaker role dependent layers

End-to-End Memory Networks (Sukhbaatar et al, 2015)

Material: http://opendialogue.miulab.tw

U: “i d like to purchase tickets to see deepwater horizon”
S: “for which theatre”
U: “angelika”
S: “you want them for angelika theatre?”
U: “yes angelika”
S: “how many tickets would you like ?”
U: “3 tickets for saturday”
S: “What time would you like ?”
U: “Any time on saturday is fine”
S: “okay , there is 4:10 pm , 5:40 pm and 9:20 pm”
U: “Let’s do 5:40”

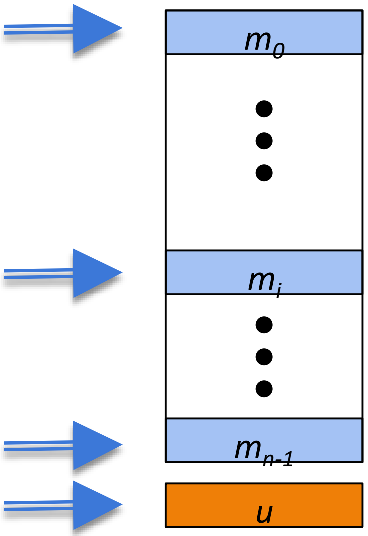

E2E MemNN for Contextual LU (Chen et al., 2016)

Material: http://opendialogue.miulab.tw

1. Sentence Encoding

$$m_i = \mathrm{RNN}_{\mathrm{mem}}(x_i)$$
$$u = \mathrm{RNN}_{\mathrm{in}}(c)$$

2. Knowledge Attention

$$p_i = \mathrm{softmax}(u^T m_i)$$

3. Knowledge Encoding

$$h = \sum_i p_i m_i \quad o = W_{\mathrm{kg}}(h + u)$$

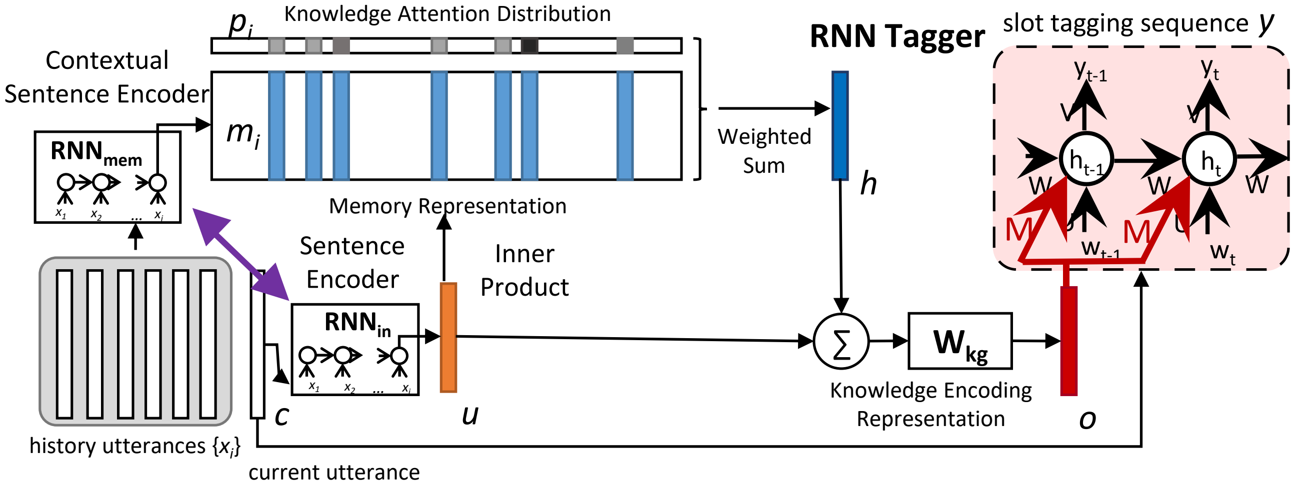

Idea: additionally incorporating contextual knowledge during slot tagging
→ track dialogue states in a latent way

Analysis of Attention

Material: http://opendialogue.miulab.tw

U: "i d like to purchase tickets to see deepwater horizon" → 0.69

S: "for which theatre"

U: "angelika"

S: "you want them for angelika theatre?"

U: "yes angelika"

S: "how many tickets would you like ?" → 0.13

U: "3 tickets for saturday"

S: "What time would you like ?"

U: "Any time on saturday is fine"

S: "okay , there is 4:10 pm , 5:40 pm and 9:20 pm" → 0.16

U: "Let's do 5:40"

Sequential Dialogue Encoder Network (Bapna et al., 2017)

- Past and current turn encodings input to a feed forward network

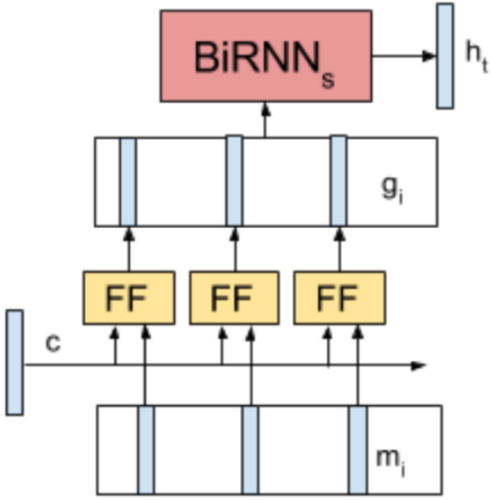

Structural LU ([Chen et al., 2016](#))

- K-SAN: prior knowledge as a teacher

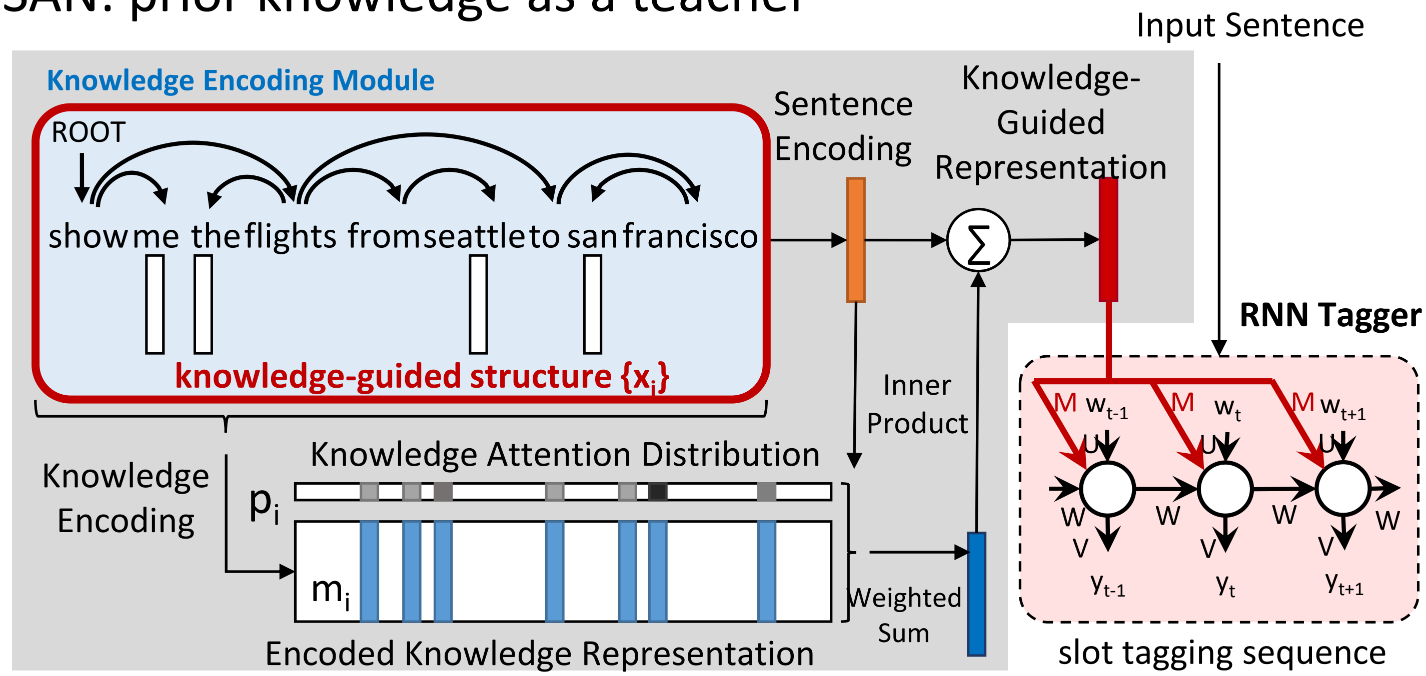

Structural LU (Chen et al., 2016)

- Sentence structural knowledge stored as memory

Sentence s show me the flights from seattle to san francisco

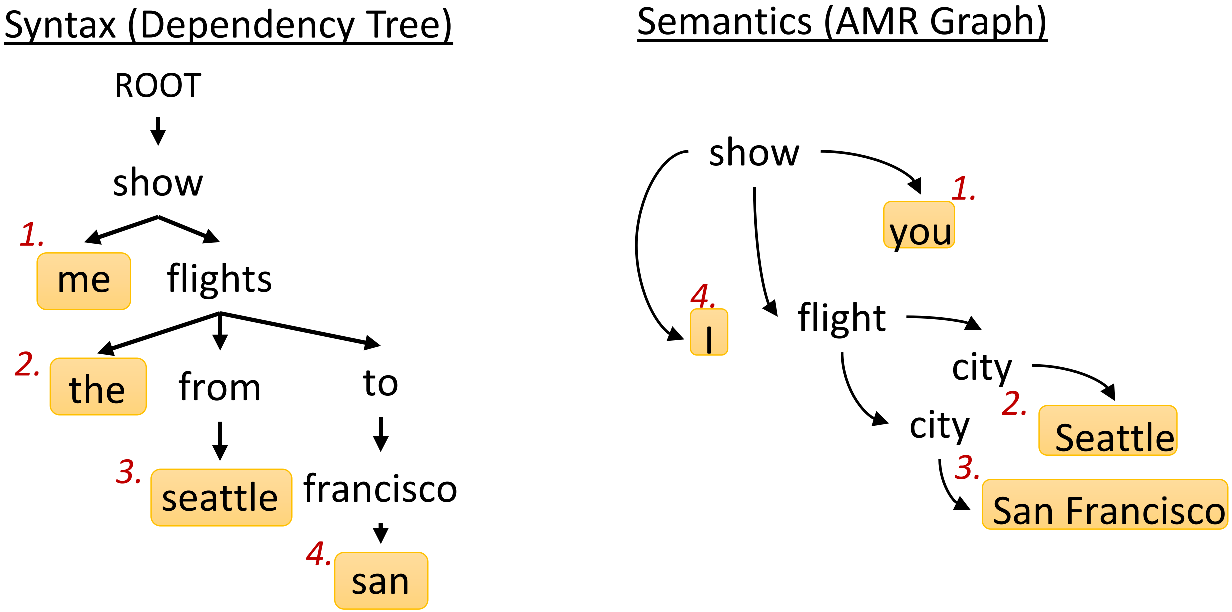

Structural LU (Chen et al., 2016)

- Sentence structural knowledge stored as memory

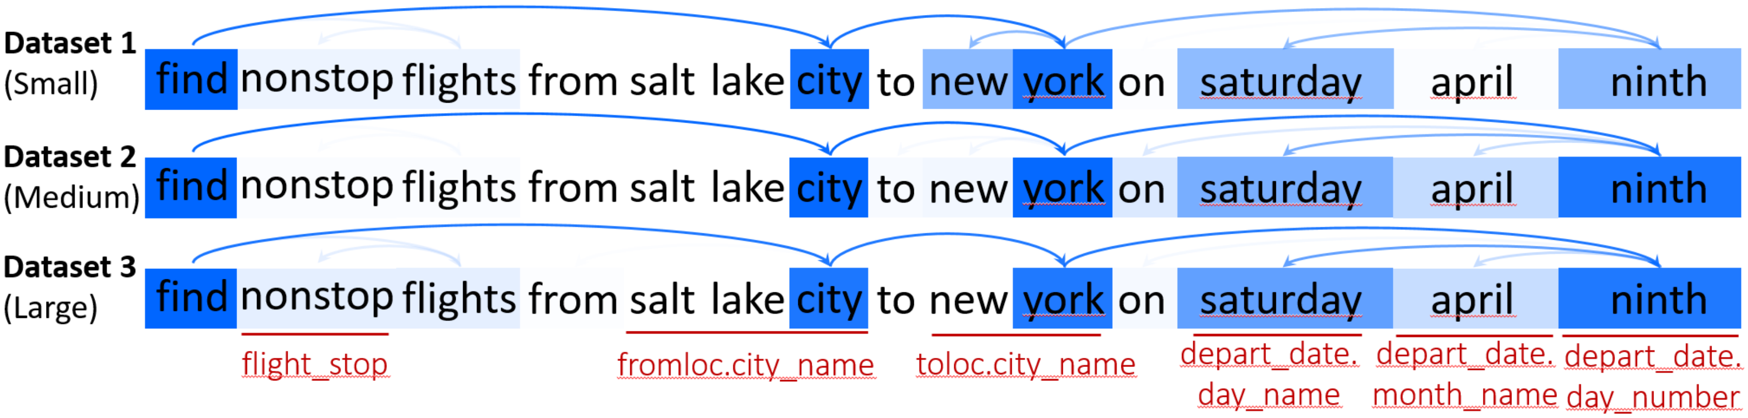

Using *less training data* with K-SAN allows the model pay the similar attention to the *salient substructures* that are important for tagging.

LU Importance (Li et al., 2017)

- Compare different types of LU errors

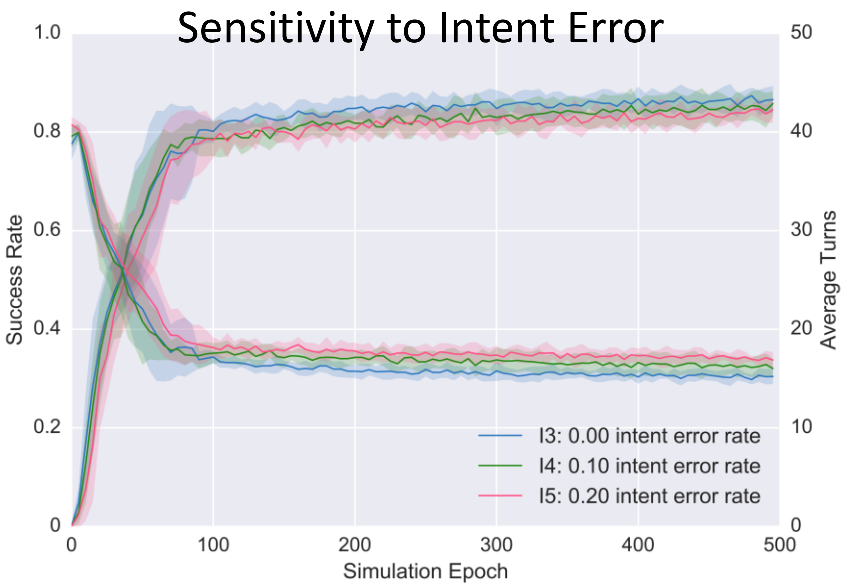

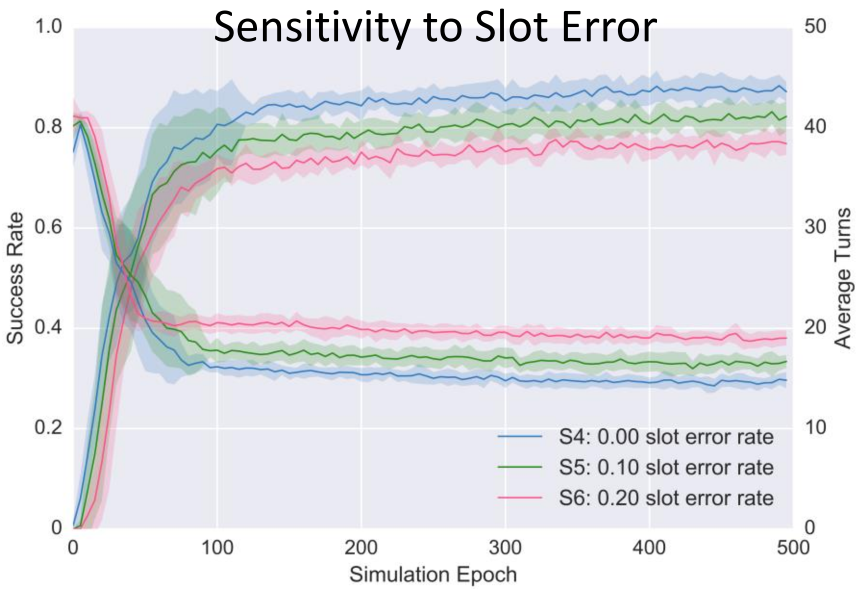

Slot filling is more important than intent detection in language understanding

LU Evaluation

- Metrics
  - Sub-sentence-level: intent accuracy, slot F1
  - Sentence-level: whole frame accuracy

Outline

- PART I. Introduction & Background Knowledge
- **PART II. Task-Oriented Dialogue Systems**
  - Spoken/Natural Language Understanding (SLU/NLU)
  - ***Dialogue Management – Dialogue State Tracking (DST)***
  - Dialogue Management – Dialogue Policy Optimization
  - Natural Language Generation (NLG)
  - End-to-End Task-Oriented Dialogue Systems
- PART III. Social Chat Bots
- PART IV. Evaluation
- PART V. Recent Trends and Challenges

Elements of Dialogue Management

Material: http://opendialogue.miulab.tw

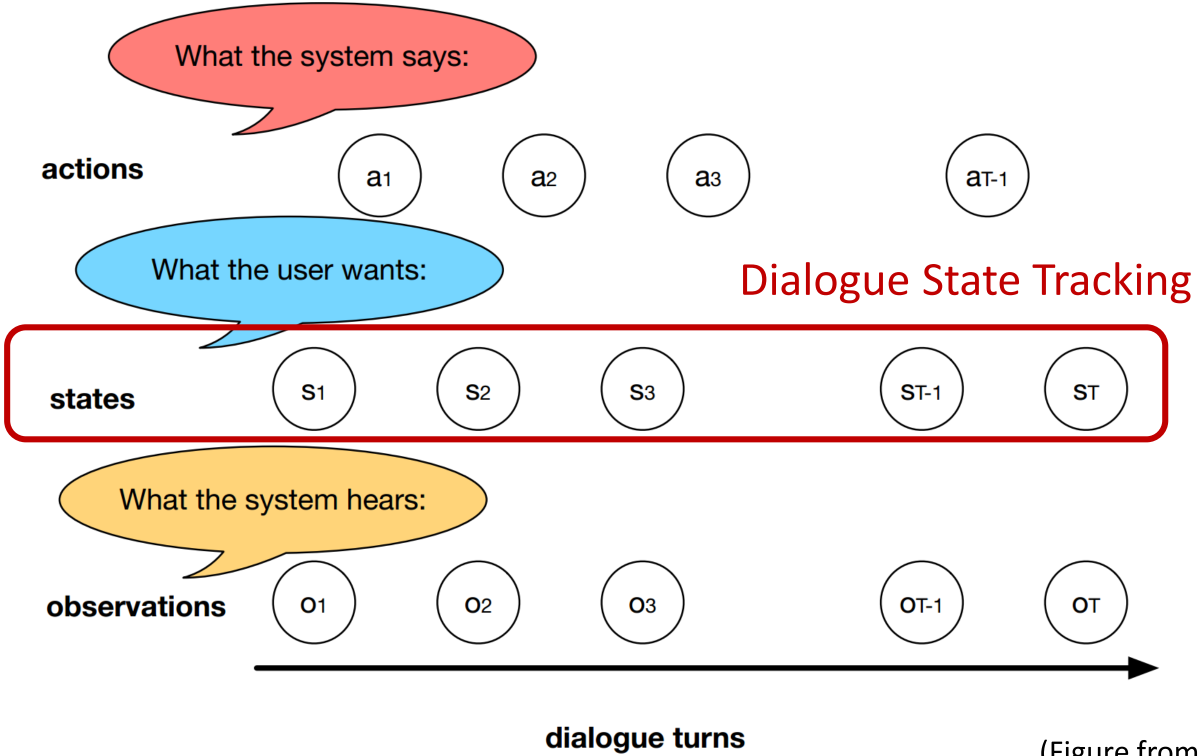

(Figure from Gašić)

Dialogue State Tracking (DST)

Material: http://opendialogue.miulab.tw

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

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

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

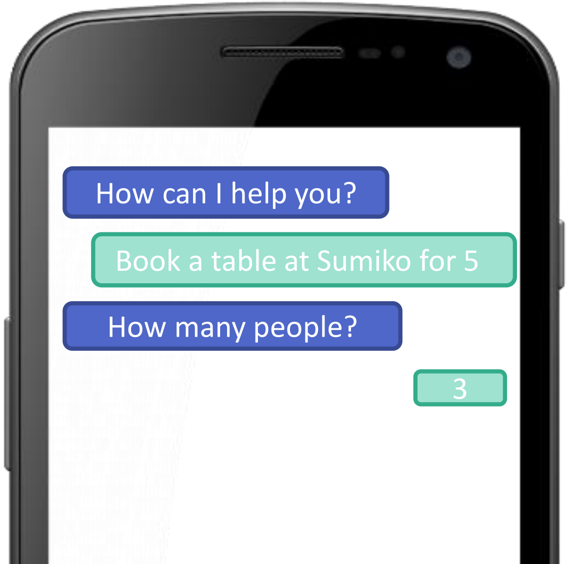

Multi-Domain Dialogue State Tracking (DST)

Material: http://opendialogue.miulab.tw

- A full representation of the system's belief of the user's goal at any point during the dialogue
- Used for making API calls

Movies

11/15/16			
6 pm	7 pm	8 pm	9 pm
2	3		
Inferno	Trolls		
Century 16			

Restaurants

Date	11/15/16		
Time	6:30 pm	7 pm	7:30 pm
Cuisine	Mexican		
Restaurant	Vive Sol		

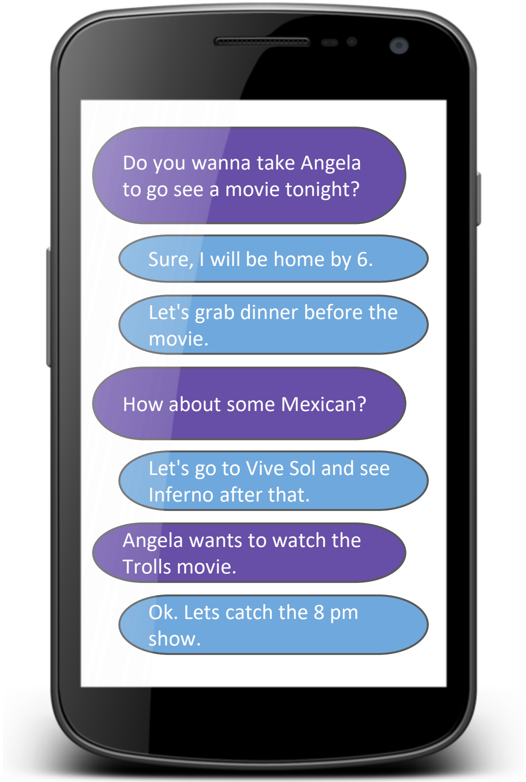

Dialog State Tracking Challenge (DSTC)

(Williams et al. 2013, Henderson et al. 2014, Henderson et al. 2014, Kim et al. 2016, Kim et al. 2016)

Challenge	Type	Domain	Data Provider	Main Theme
DSTC1	Human-Machine	Bus Route	CMU	Evaluation Metrics
DSTC2	Human-Machine	Restaurant	U. Cambridge	User Goal Changes
DSTC3	Human-Machine	Tourist Information	U. Cambridge	Domain Adaptation
DSTC4	Human-Human	Tourist Information	I2R	Human Conversation
DSTC5	Human-Human	Tourist Information	I2R	Language Adaptation

NN-Based DST (Henderson et al., 2013; Mrkšić et al., 2015; Mrkšić et al., 2016)

Material: http://opendialogue.miulab.tw

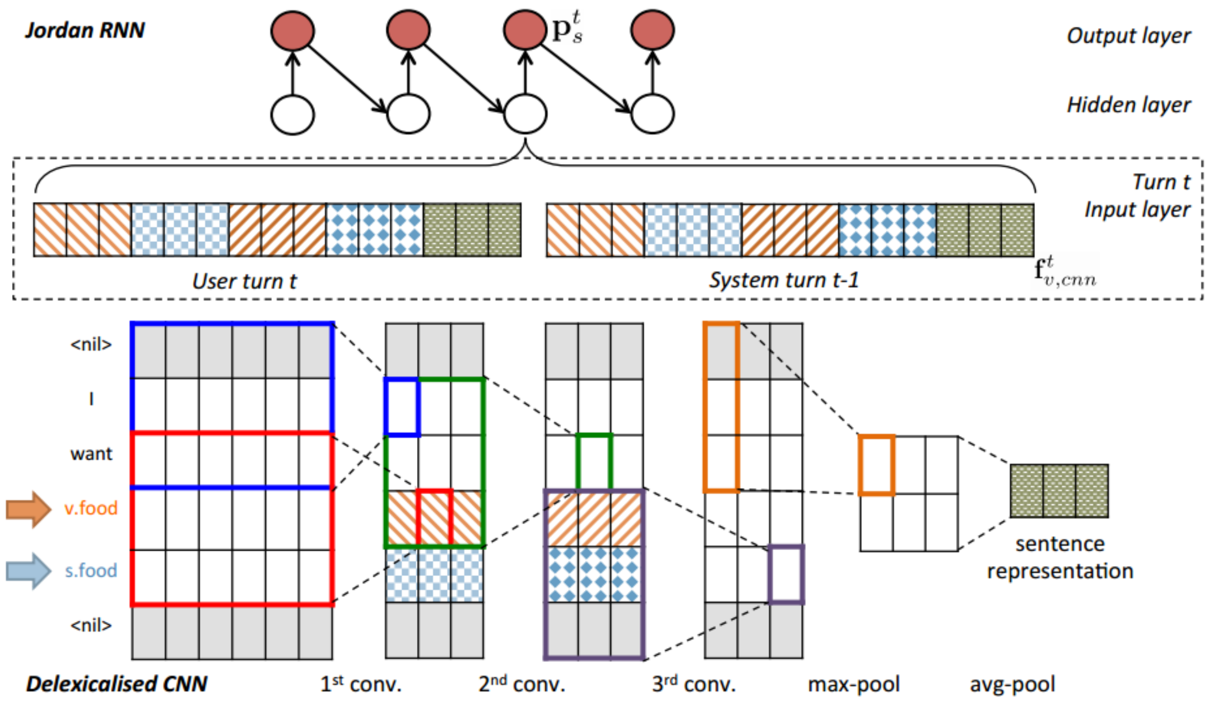

(Figure from Wen et al, 2016)

Neural Belief Tracker (Mrkšić et al., 2016)

Material: http://opendialogue.miulab.tw

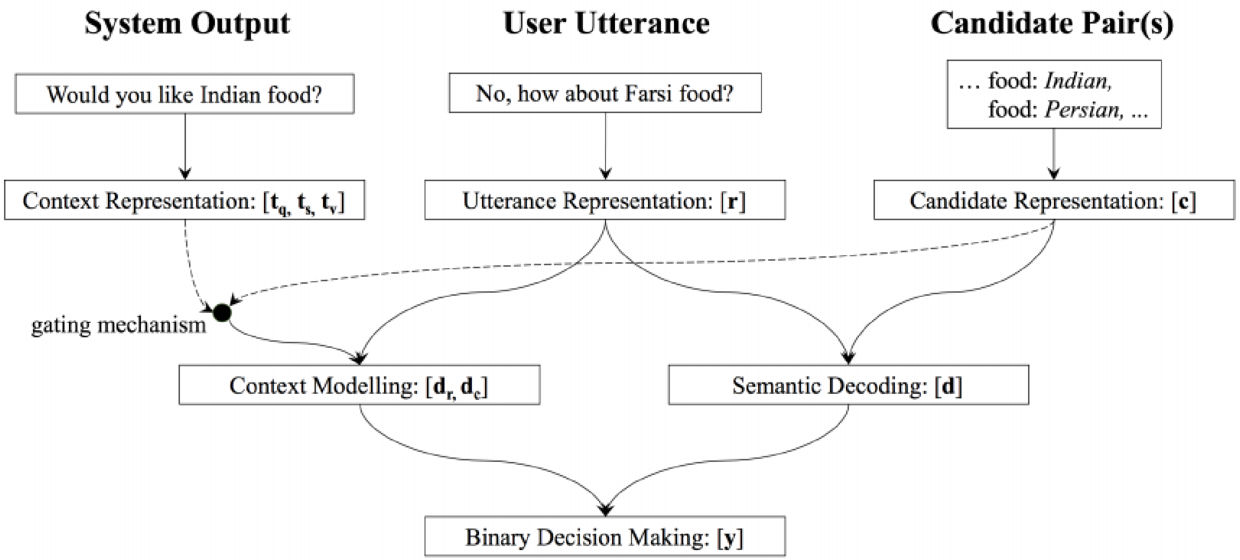

DST Evaluation

Material: http://opendialogue.miulab.tw

- Dialogue State Tracking Challenges
  - DSTC2-3, human-machine
  - DSTC4-5, human-human
- Metric
  - Tracked state accuracy with respect to user goal
  - Recall/Precision/F-measure individual slots

Outline

- PART I. Introduction & Background Knowledge
- **PART II. Task-Oriented Dialogue Systems**
  - Spoken/Natural Language Understanding (SLU/NLU)
  - Dialogue Management – Dialogue State Tracking (DST)
  - ***Dialogue Management – Dialogue Policy Optimization***
  - Natural Language Generation (NLG)
  - End-to-End Task-Oriented Dialogue Systems
- PART III. Social Chat Bots
- PART IV. Evaluation
- PART V. Recent Trends and Challenges

Elements of Dialogue Management

Material: http://opendialogue.miulab.tw

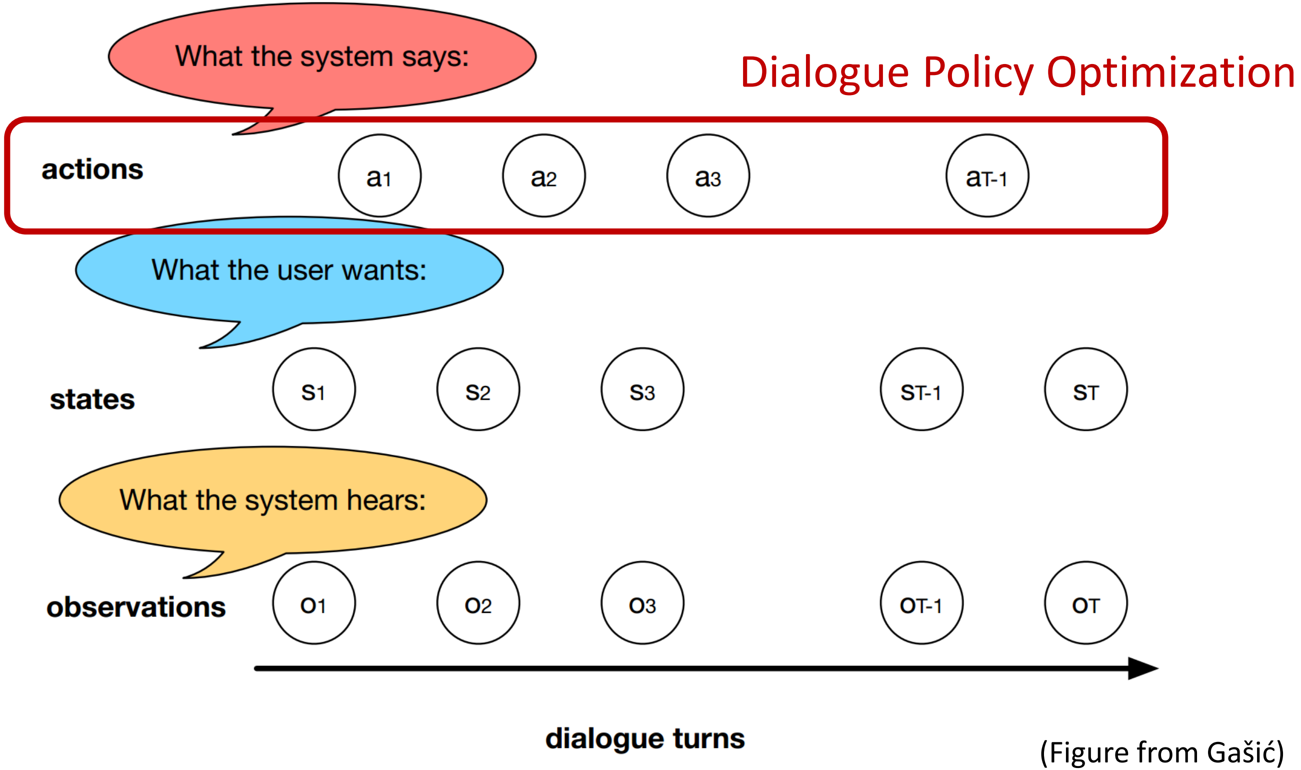

(Figure from Gašić)

Dialogue Policy Optimization

- Dialogue management in a RL framework

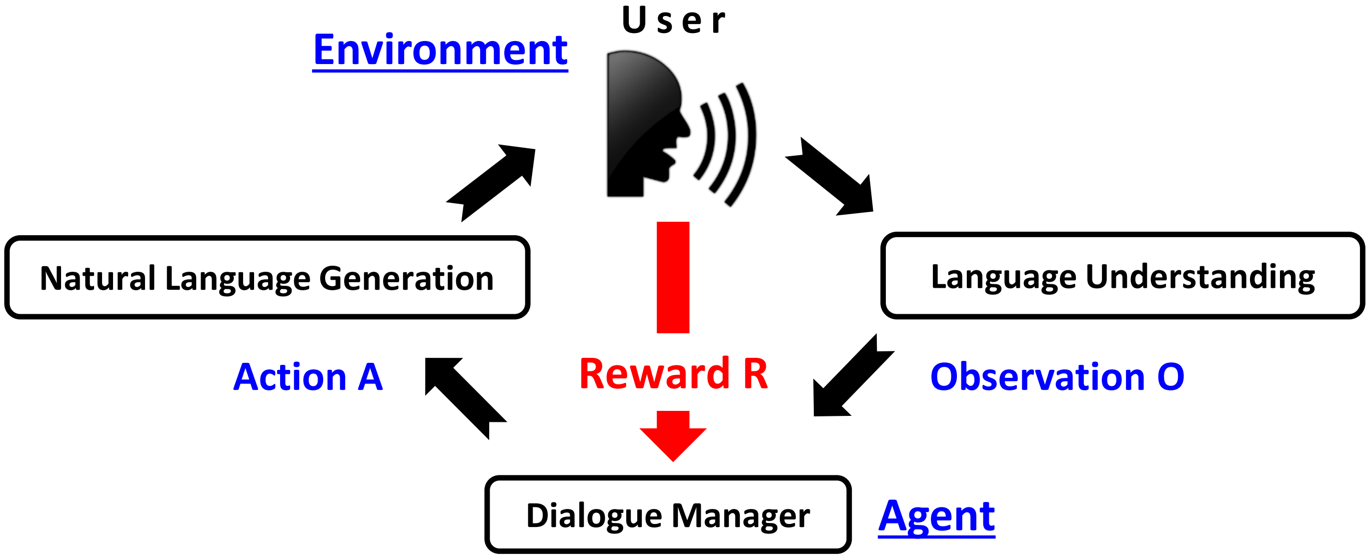

Optimized dialogue policy selects the best action that can maximize the future reward.
Correct rewards are a crucial factor in dialogue policy training

Reward for RL $\cong$ Evaluation for System

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

- Expert rating	high quality, high cost
- User rating	unreliable quality, medium cost
- Objective rating	Check desired aspects, low cost

Reinforcement Learning for Dialogue Policy Optimization

Material: http://opendialogue.miulab.tw

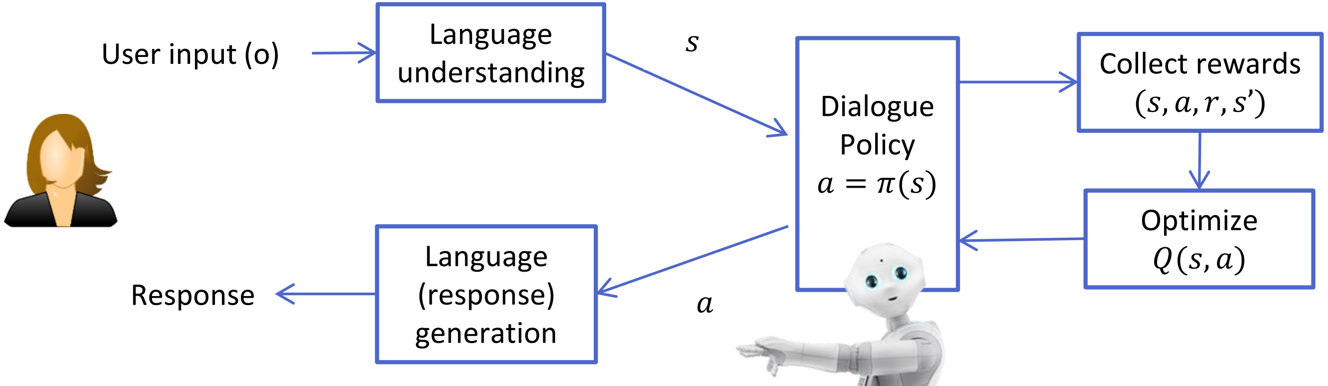

Type of Bots	State	Action	Reward
Social ChatBots	Chat history	System Response	# of turns maximized; Intrinsically motivated reward
InfoBots (interactive Q/A)	User current question + Context	Answers to current question	Relevance of answer; # of turns minimized
Task-Completion Bots	User current input + Context	System dialogue act w/ slot value (or API calls)	Task success rate; # of turns minimized

Goal: develop a generic deep RL algorithm to learn dialogue policy for all bot categories

Dialogue Reinforcement Learning Signal

- Typical reward function
  - -1 for per turn penalty
  - Large reward at completion if successful
- Typically requires domain knowledge
  - ✔ Simulated user
  - ✔ Paid users (Amazon Mechanical Turk)
  - ✖ Real users

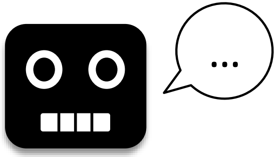

The user simulator is usually required for dialogue system training before deployment

Neural Dialogue Manager (Li et al., 2017)

- Deep Q-network for training DM policy
  - Input: current semantic frame observation, database returned results
  - Output: system action

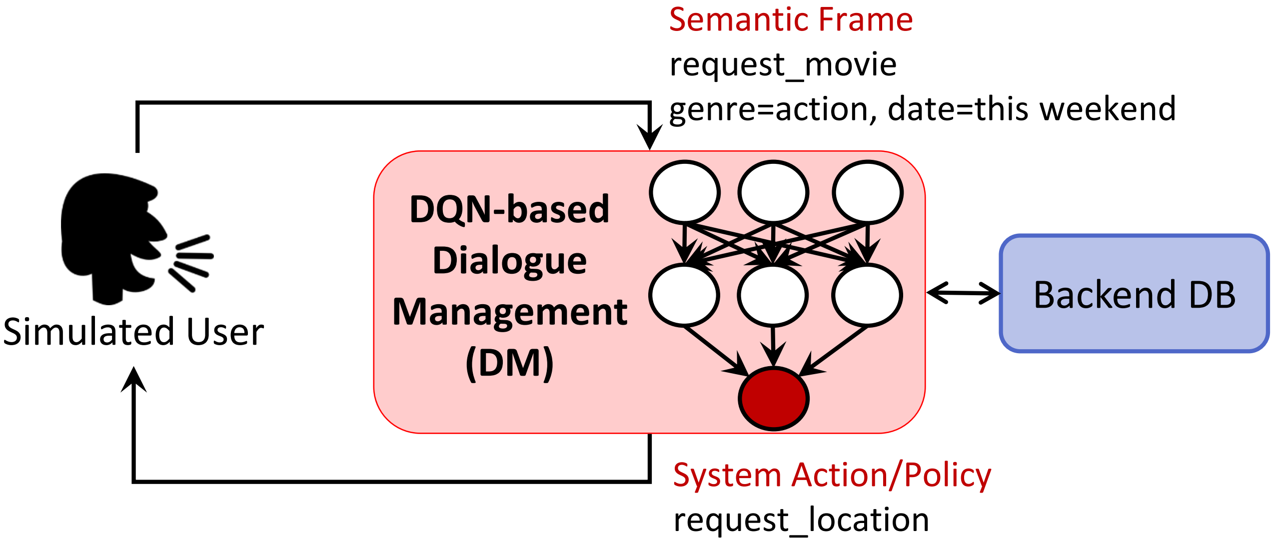

SL + RL for Sample Efficiency (Su et al., 2017)

- Issue about RL for DM
  - slow learning speed
  - cold start
- Solutions
  - Sample-efficient actor-critic
    - Off-policy learning with experience replay
    - Better gradient update
  - Utilizing supervised data
    - Pretrain the model with SL and then fine-tune with RL
    - Mix SL and RL data during RL learning
    - Combine both

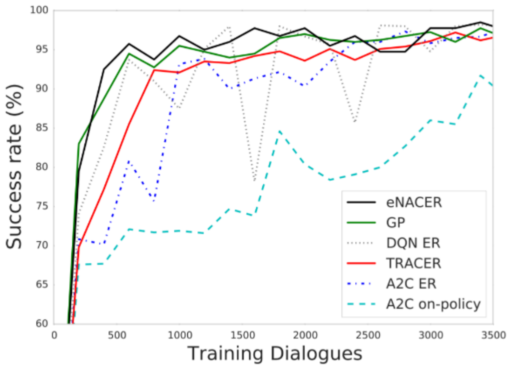

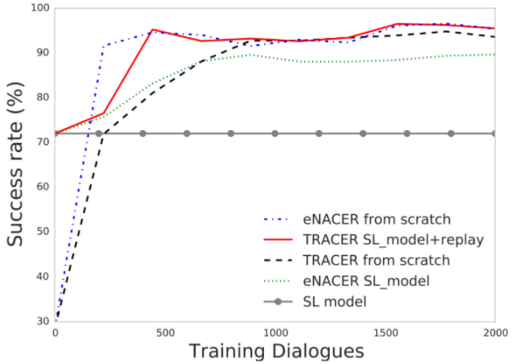

Online Training (Su et al., 2015; Su et al., 2016)

- Policy learning from real users
  - Infer reward directly from dialogues (Su et al., 2015)
  - User rating (Su et al., 2016)
- Reward modeling on user binary success rating

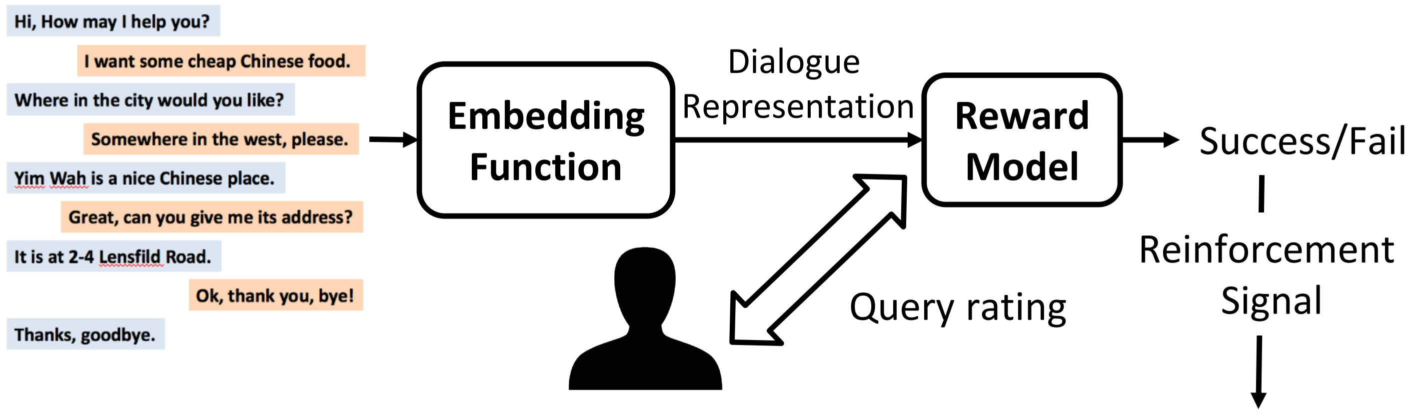

Interactive RL for DM (Shah et al., 2016)

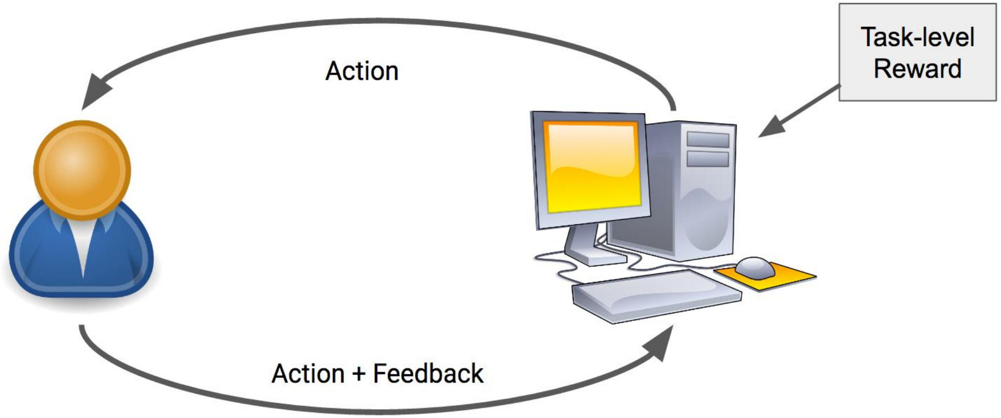

Immediate Feedback

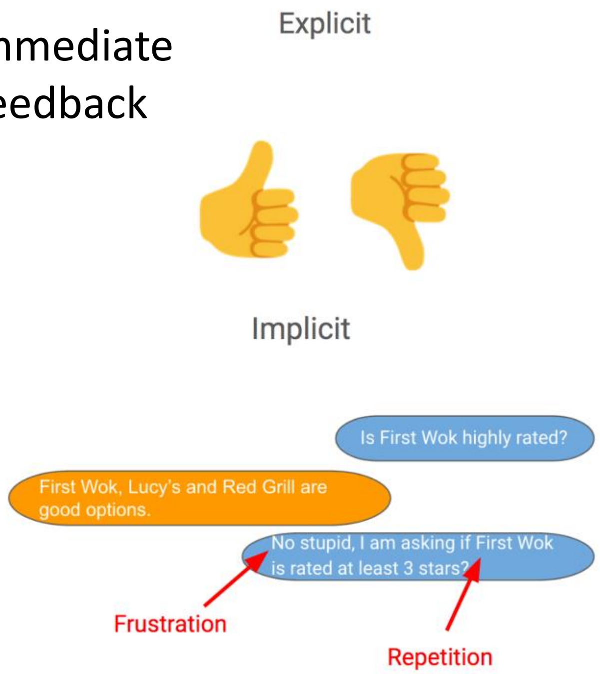

Use a third agent for providing interactive feedback to the DM

Dialogue Management Evaluation

- Metrics
  - Turn-level evaluation: system action accuracy
  - Dialogue-level evaluation: task success rate, reward

Outline

- PART I. Introduction & Background Knowledge
- **PART II. Task-Oriented Dialogue Systems**
  - Spoken/Natural Language Understanding (SLU/NLU)
  - Dialogue Management – Dialogue State Tracking (DST)
  - Dialogue Management – Dialogue Policy Optimization
  - ***Natural Language Generation (NLG)***
  - End-to-End Task-Oriented Dialogue Systems
- PART III. Social Chat Bots
- PART IV. Evaluation
- PART V. Recent Trends and Challenges

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

Material: http://opendialogue.miulab.tw

- 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

Plan-Based NLG (Walker et al., 2002)

- Divide the problem into pipeline

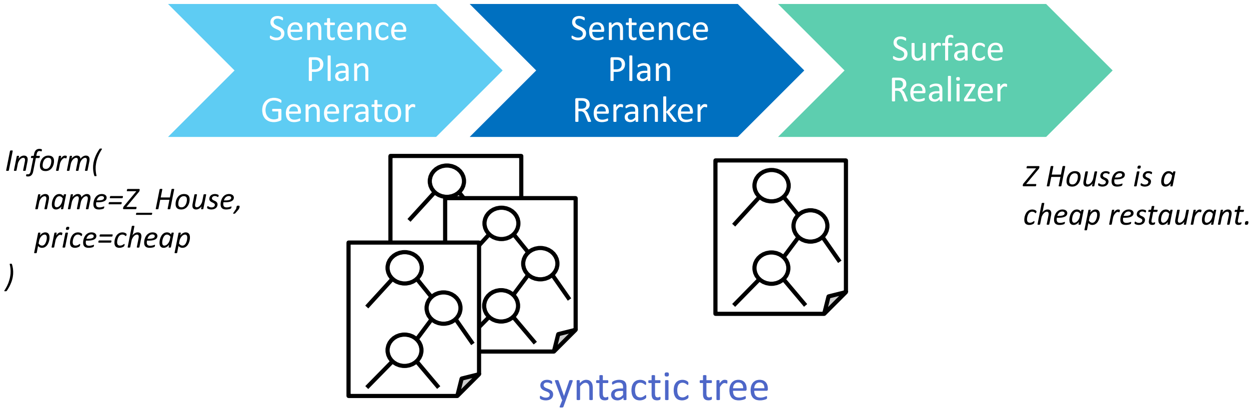

  - Statistical sentence plan generator (Stent et al., 2009)
  - Statistical surface realizer (Dethlefs et al., 2013; Cuayáhuitl et al., 2014; …)

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

Class-Based LM NLG (Oh and Rudnicky, 2000)

- Class-based language modeling

$$P(X \mid c) = \sum_t \log p(x_t \mid x_0, x_1, \cdots, x_{t-1}, c)$$

- NLG by decoding

$$X^* = \arg\max_X P(X \mid c)$$

Classes:
inform_area
inform_address
…
request_area
request_postcode

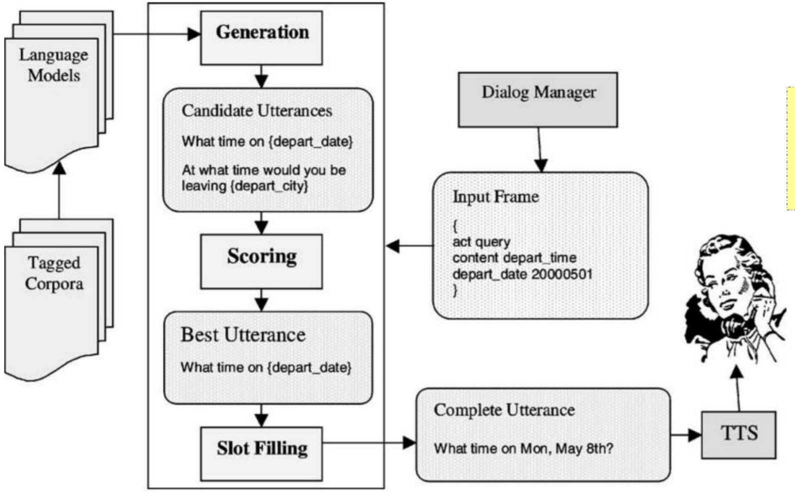

Pros: easy to implement/ understand, simple rules
Cons: computationally inefficient

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

Material: http://opendialogue.miulab.tw

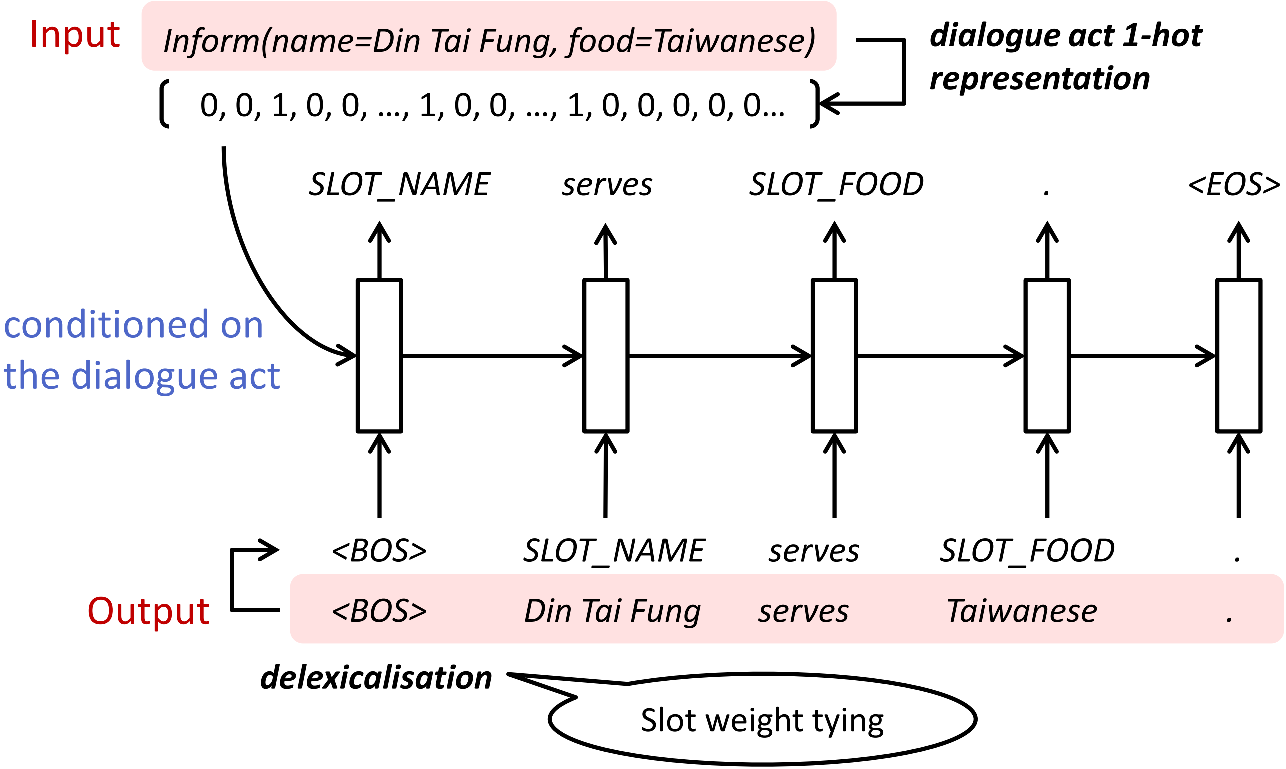

Handling Semantic Repetition

- 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.
- Deficiency in either model or decoding (or both)
- Mitigation
  - Post-processing rules (Oh & Rudnicky, 2000)
  - Gating mechanism (Wen et al., 2015)
  - Attention (Mei et al., 2016; Wen et al., 2015)

Semantic Conditioned LSTM (Wen et al., 2015)

- Original LSTM cell

$$\mathbf{i}_t = \sigma(\mathbf{W}_{wi}\mathbf{x}_t + \mathbf{W}_{hi}\mathbf{h}_{t-1})$$

$$\mathbf{f}_t = \sigma(\mathbf{W}_{wf}\mathbf{x}_t + \mathbf{W}_{hf}\mathbf{h}_{t-1})$$

$$\mathbf{o}_t = \sigma(\mathbf{W}_{wo}\mathbf{x}_t + \mathbf{W}_{ho}\mathbf{h}_{t-1})$$

$$\hat{\mathbf{c}}_t = \tanh(\mathbf{W}_{wc}\mathbf{x}_t + \mathbf{W}_{hc}\mathbf{h}_{t-1})$$

$$\mathbf{c}_t = \mathbf{f}_t \odot \mathbf{c}_{t-1} + \mathbf{i}_t \odot \hat{\mathbf{c}}_t$$

$$\mathbf{h}_t = \mathbf{o}_t \odot \tanh(\mathbf{c}_t)$$

- Dialogue act (DA) cell

$$\mathbf{r}_t = \sigma(\mathbf{W}_{wr}\mathbf{x}_t + \mathbf{W}_{hr}\mathbf{h}_{t-1})$$

$$\mathbf{d}_t = \mathbf{r}_t \odot \mathbf{d}_{t-1}$$

- Modify Ct

$$\mathbf{c}_t = \mathbf{f}_t \odot \mathbf{c}_{t-1} + \mathbf{i}_t \odot \hat{\mathbf{c}}_t + \tanh(\mathbf{W}_{dc}\mathbf{d}_t)$$

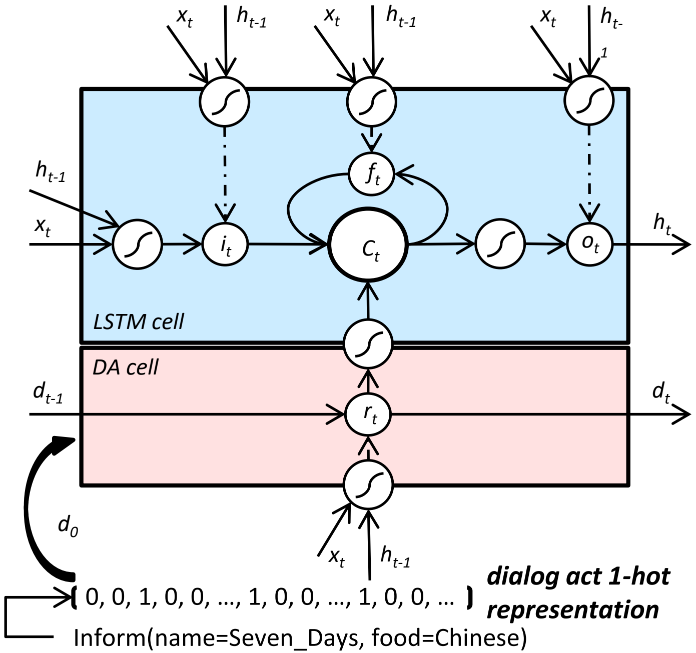

dialog act 1-hot representation

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

Structural NLG ([Dušek and Jurčíček, 2016](#))

- Goal: NLG based on the syntax tree
  - Encode trees as sequences
  - Seq2Seq model for generation

inform(name=X-name,type=placetoeat,eattype=restaurant,
area=riverside,food=Italian)

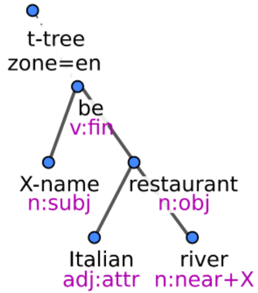

(<root> <root> ((X-name n:subj) be v:fin ((Italian adj:attr) restaurant n:obj (river n:near+X))))

X-name n:subj be v:fin Italian adj:attr restaurant n:obj river n:near+X

X is an Italian restaurant near the river.

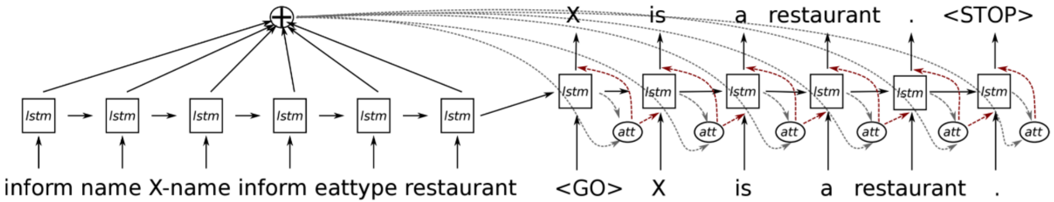

Contextual NLG ([Dušek and Jurčíček, 2016](#))

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

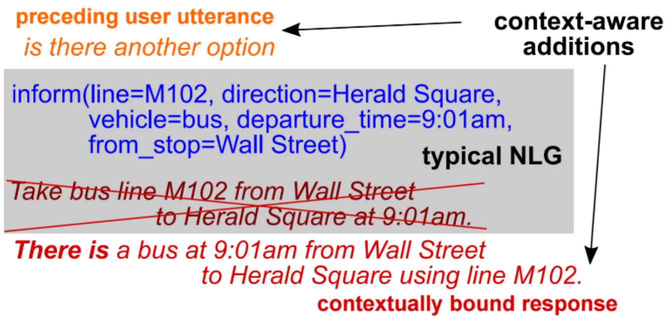

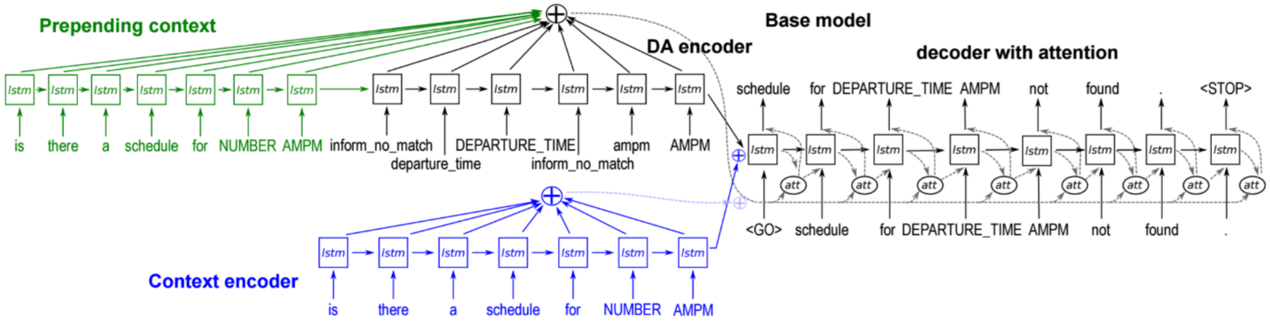

Controlled Text Generation (Hu et al., 2017)

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

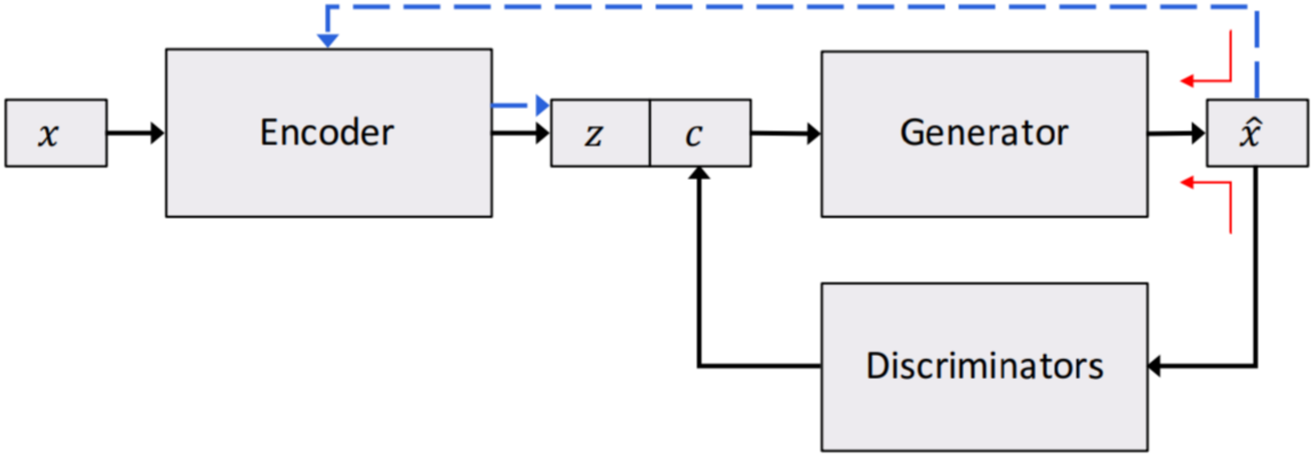

NLG Evaluation

- Metrics
  - Subjective: human judgement (Stent et al., 2005)
    - Adequacy: correct meaning
    - Fluency: linguistic fluency
    - Readability: fluency in the dialogue context
    - Variation: multiple realizations for the same concept
  - Objective: automatic metrics
    - Word overlap: BLEU (Papineni et al, 2002), METEOR, ROUGE
    - Word embedding based: vector extrema, greedy matching, embedding average

There is a gap between human perception and automatic metrics

Outline

- PART I. Introduction & Background Knowledge
- **PART II. Task-Oriented Dialogue Systems**
  - Spoken/Natural Language Understanding (SLU/NLU)
  - Dialogue Management – Dialogue State Tracking (DST)
  - Dialogue Management – Dialogue Policy Optimization
  - Natural Language Generation (NLG)
  - ***End-to-End Task-Oriented Dialogue Systems***
- PART III. Social Chat Bots
- PART IV. Evaluation
- PART V. Recent Trends and Challenges

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

- Errors from DM can be propagated to NLU for *regularization* + *robustness*

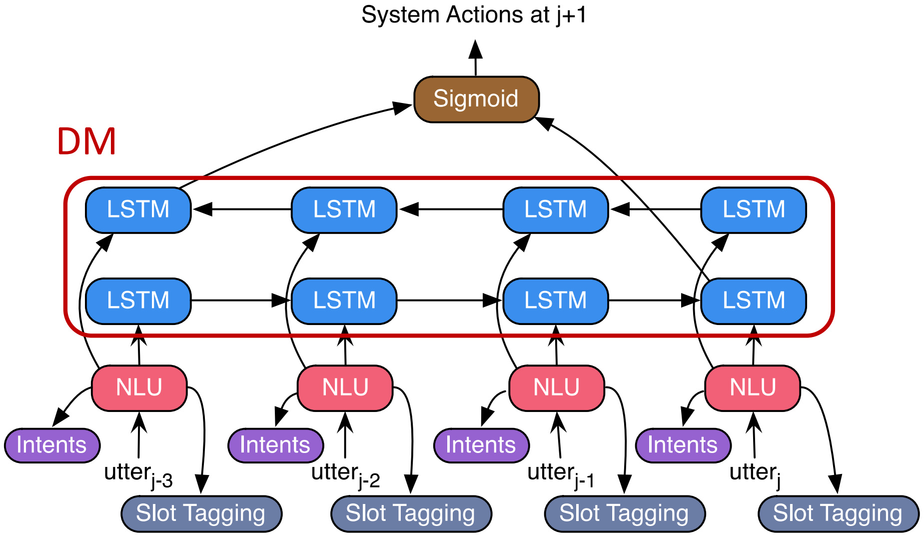

Model	DM	NLU
Baseline (CRF+SVMs)	7.7	33.1
Pipeline-BLSTM	12.0	36.4
JointModel	**22.8**	**37.4**

Both DM and NLU performance (frame accuracy) is improved

E2E Supervised Dialogue System (Wen et al., 2017)

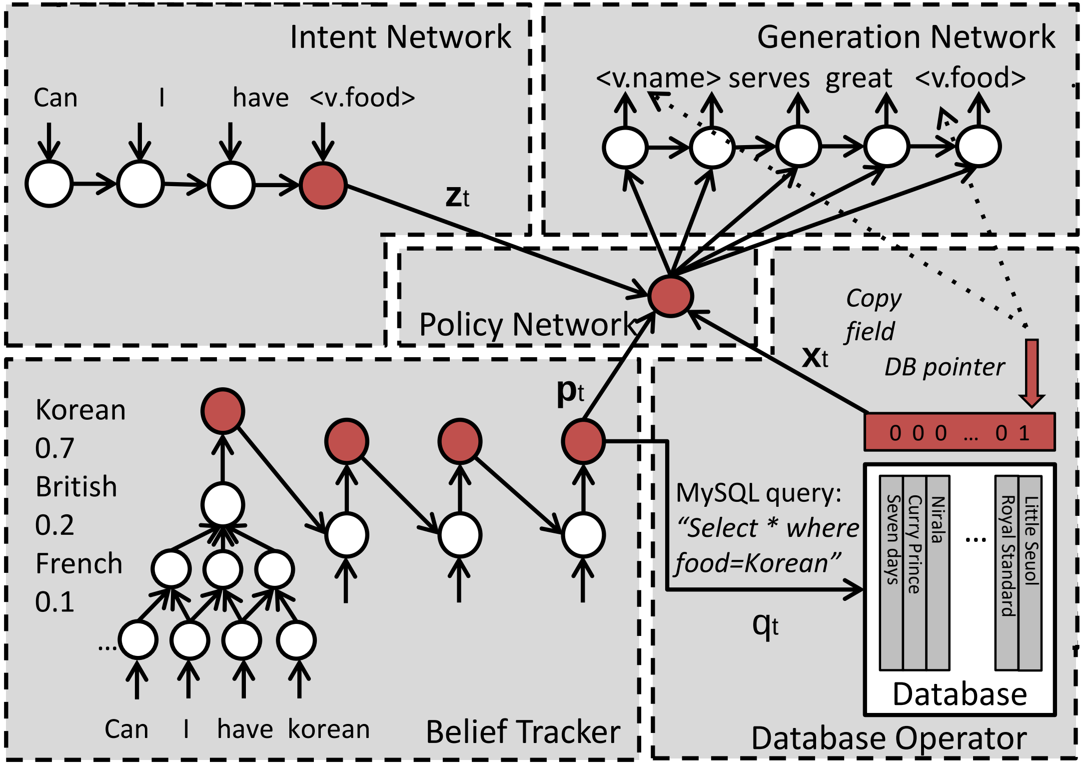

E2E MemNN for Dialogues (Bordes et al., 2017)

- Split dialogue system actions into subtasks
  - API issuing
  - API updating
  - Option displaying
  - Information informing

Task	Memory Networks			
	no match type		+ match type	
T1: Issuing API calls	**99.9**	(99.6)	**100**	(100)
T2: Updating API calls	**100**	(100)	98.3	(83.9)
T3: Displaying options	**74.9**	(2.0)	**74.9**	(0)
T4: Providing information	59.5	(3.0)	**100**	(100)
T5: Full dialogs	**96.1**	(49.4)	93.4	(19.7)
T1(OOV): Issuing API calls	72.3	(0)	**96.5**	(82.7)
T2(OOV): Updating API calls	78.9	(0)	**94.5**	(48.4)
T3(OOV): Displaying options	74.4	(0)	**75.2**	(0)
T4(OOV): Providing inform.	57.6	(0)	**100**	(100)
T5(OOV): Full dialogs	65.5	(0)	**77.7**	(0)
T6: Dialog state tracking 2	**41.1**	(0)	**41.0**	(0)

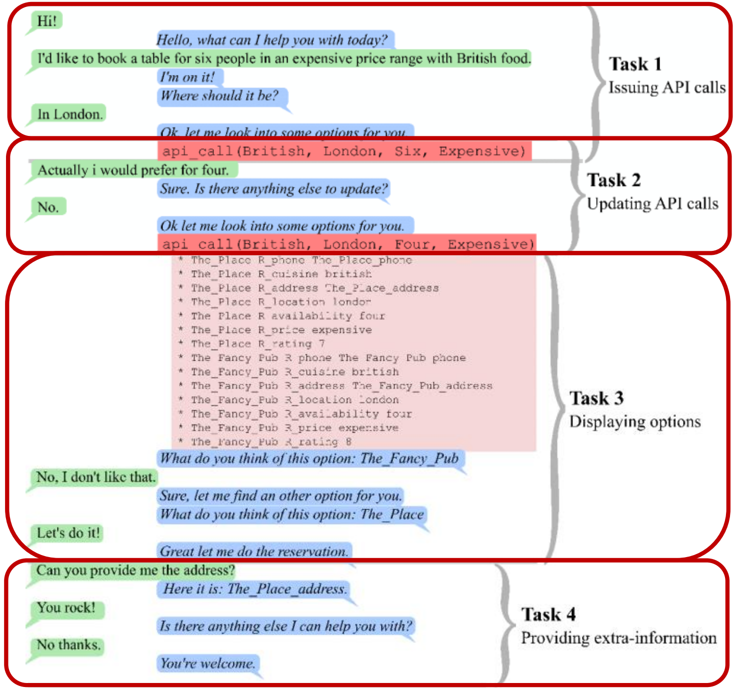

Task 5 Conducting full dialogs

E2E RL-Based KB-InfoBot (Dhingra et al., 2017)

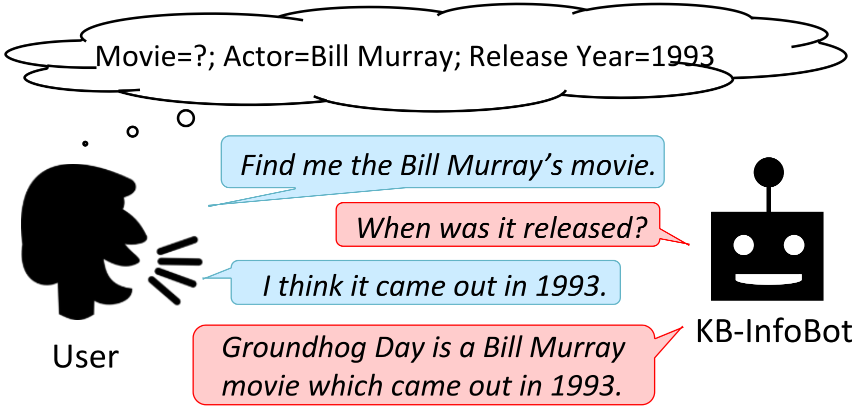

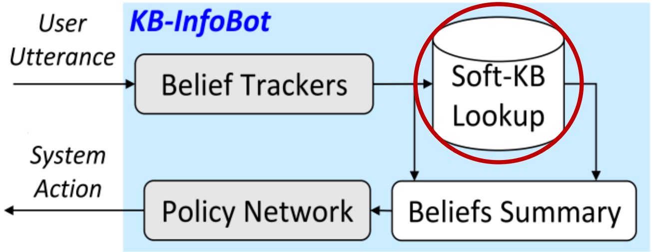

Entity-Centric Knowledge Base

Movie	Actor	Release Year
Groundhog Day	Bill Murray	1993
Australia	Nicole Kidman	X
Mad Max: Fury Road	X	2015

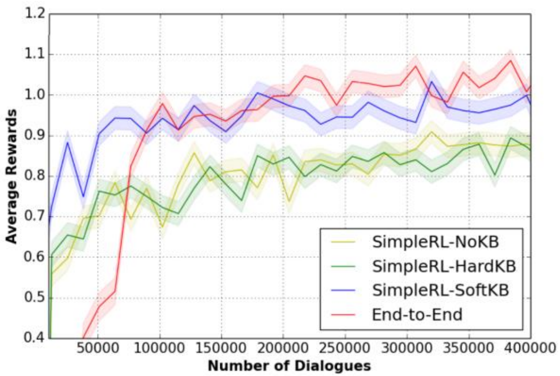

Idea: differentiable database for propagating the gradients

E2E RL-Based System (Zhao and Eskenazi, 2016)

- Joint learning
  - NLU, DST, Dialogue Policy
- Deep RL for training
  - Deep Q-network
  - Deep recurrent network

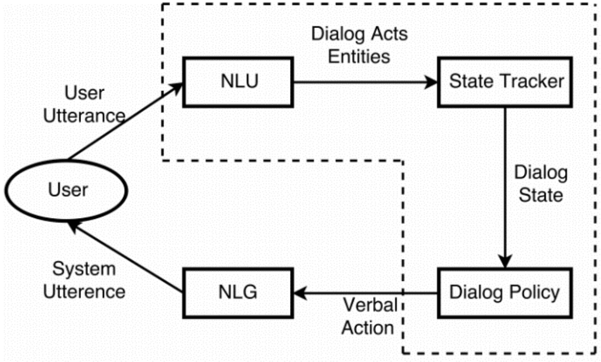

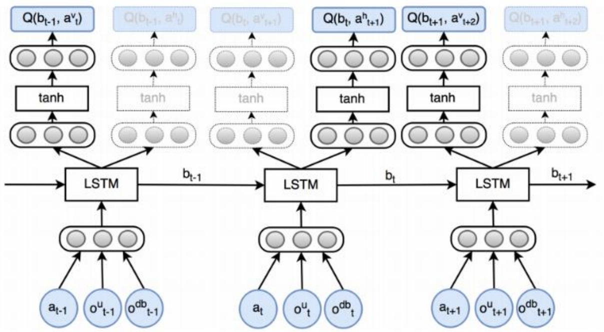

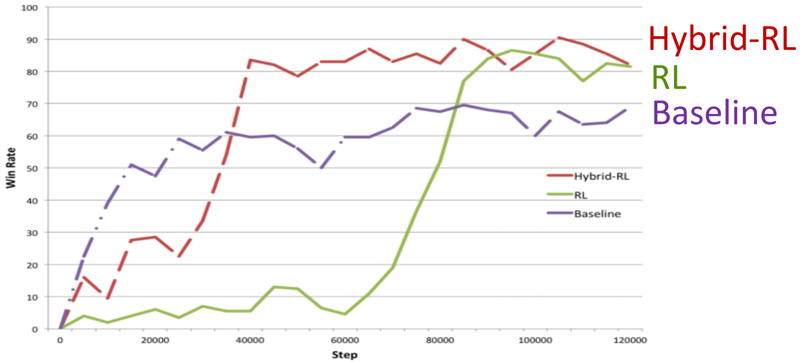

E2E LSTM-Based Dialogue Control (Williams and Zweig, 2016)

- Idea: an LSTM maps from raw dialogue history directly to a distribution over system actions
  - Developers can provide software including business rules & programmatic APIs → LSTM can take actions in the real world on behalf of the user
  - The LSTM can be optimized using SL or RL

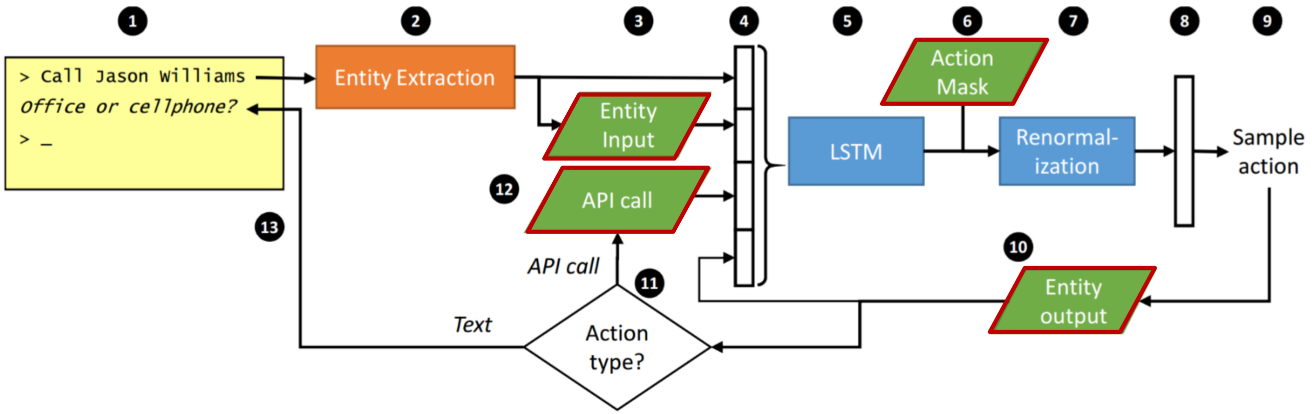

E2E Task-Completion Bot (TC-Bot) (Li et al., 2017)

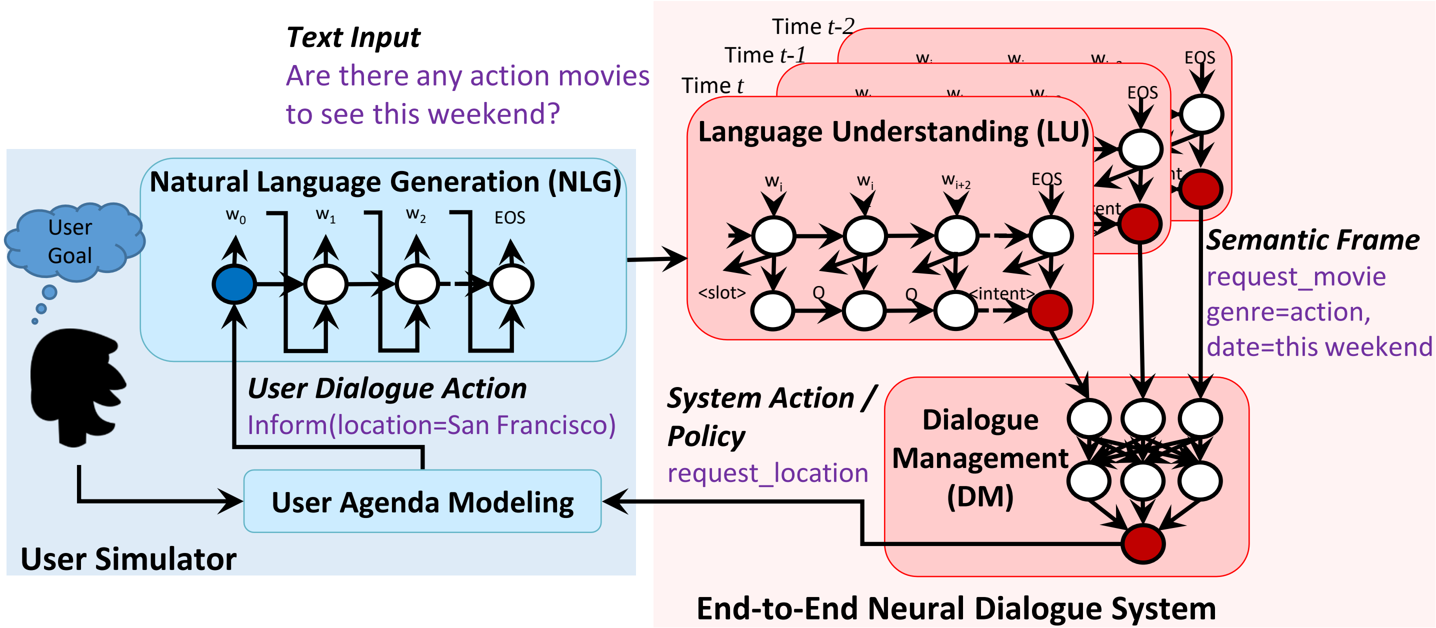

Idea: supervised learning for each component and reinforcement learning for end-to-end training the neural dialogue system

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 you.
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!

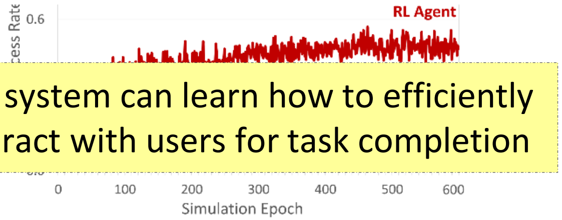

The system can learn how to efficiently interact with users for task completion

Hierarchical RL for Composite Tasks (Peng et al., 2017)

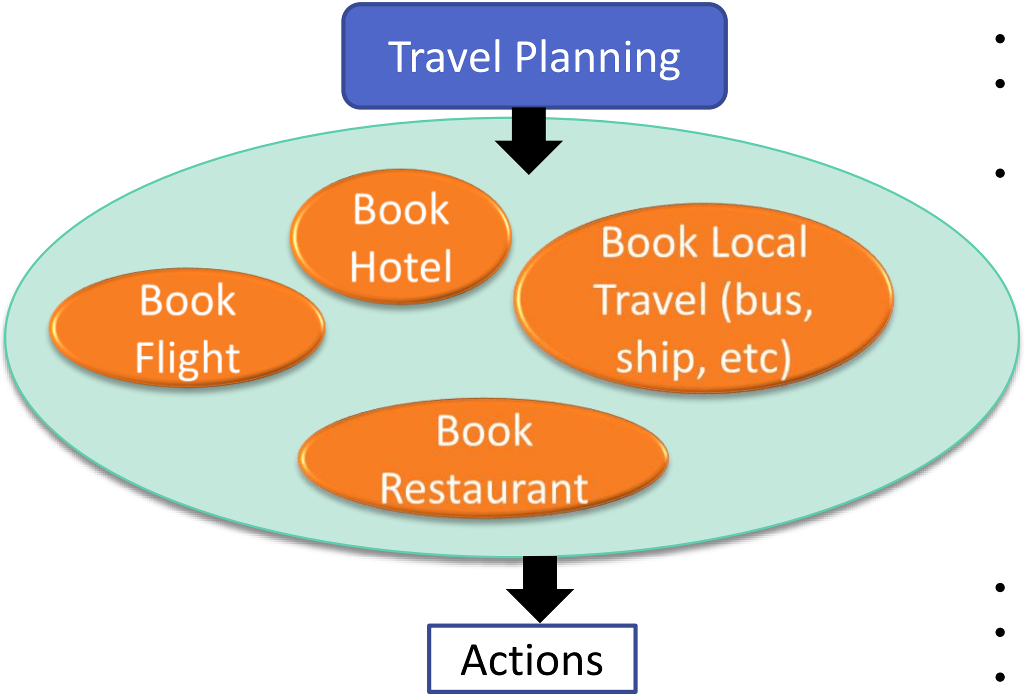

- Set of tasks that need to be fulfilled collectively!
- Build a dialog manager that satisfies **cross-subtask constraints** (**slot constraints**)
- Temporally constructed goals

- hotel_check_in_time > departure_flight_time
- # flight_tickets = #people checking in the hotel
- hotel_check_out_time< return_flight_time,

Hierarchical RL for Composite Tasks (Peng et al., 2017)

- The dialog model makes decisions over two levels: *meta-controller and controller*
- The *agent* learns these policies simultaneously
  - the policy of optimal sequence of goals to follow $\pi_g(g_t, s_t; \theta_1)$
  - Policy $\pi_{a,g}(a_t, g_t, s_t; \theta_2)$ for each sub-goal g_t

Meta-Controller

Controller

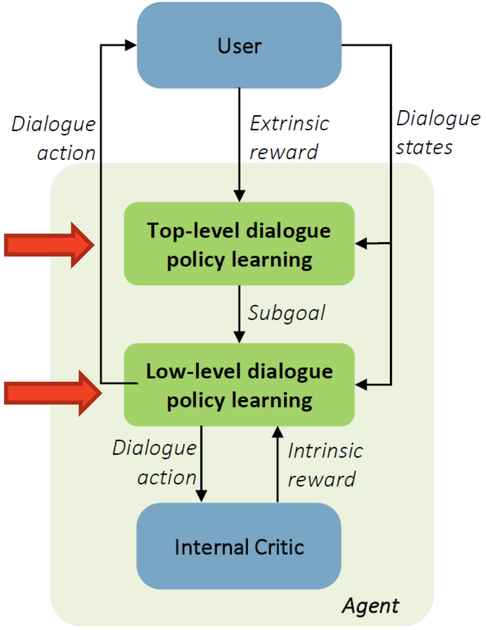

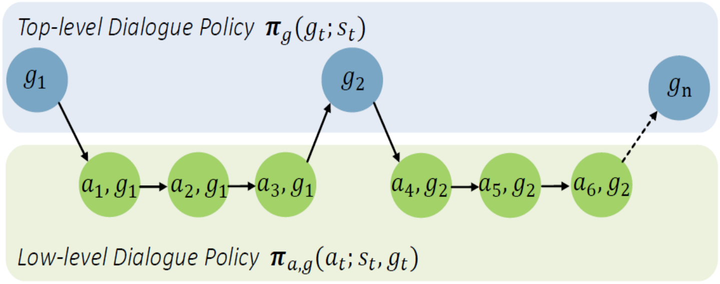

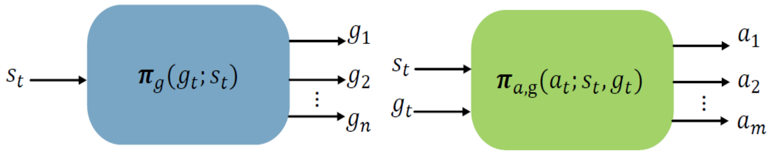

(mitigate reward sparsity issues)

Social Chat Bots

Social Chat Bots

- The success of XiaoIce (小冰)
- Problem setting and evaluation
  - Maximize the user engagement by automatically generating
  - *enjoyable* and *useful* conversations
- Learning a neural conversation engine
  - A data driven engine trained on social chitchat data (Sordoni+ 15; Li+ 16)
  - Persona based models and speaker-role based models (Li+ 16; Luan+ 17)
  - Image-grounded models (Mostafazadeh+ 17)
  - Knowledge-grounded models (Ghazvininejad+ 17)

Outline

- PART I. Introduction & Background Knowledge
- PART II. Task-Oriented Dialogue Systems
- **PART III. Social Chat Bots**
  - ***Neural Response Generation***
  - Response Diversity
  - Response Consistency
  - Deep Reinforcement Learning for Response Generation
  - Combining Task-Oriented Bots and Social Chat Bots
- PART IV. Evaluation
- PART V. Recent Trends and Challenges

Neural Response Generation (Sordoni+ 15; Vinyals & Le 15; Shang+ 15)

Material: http://opendialogue.miulab.tw

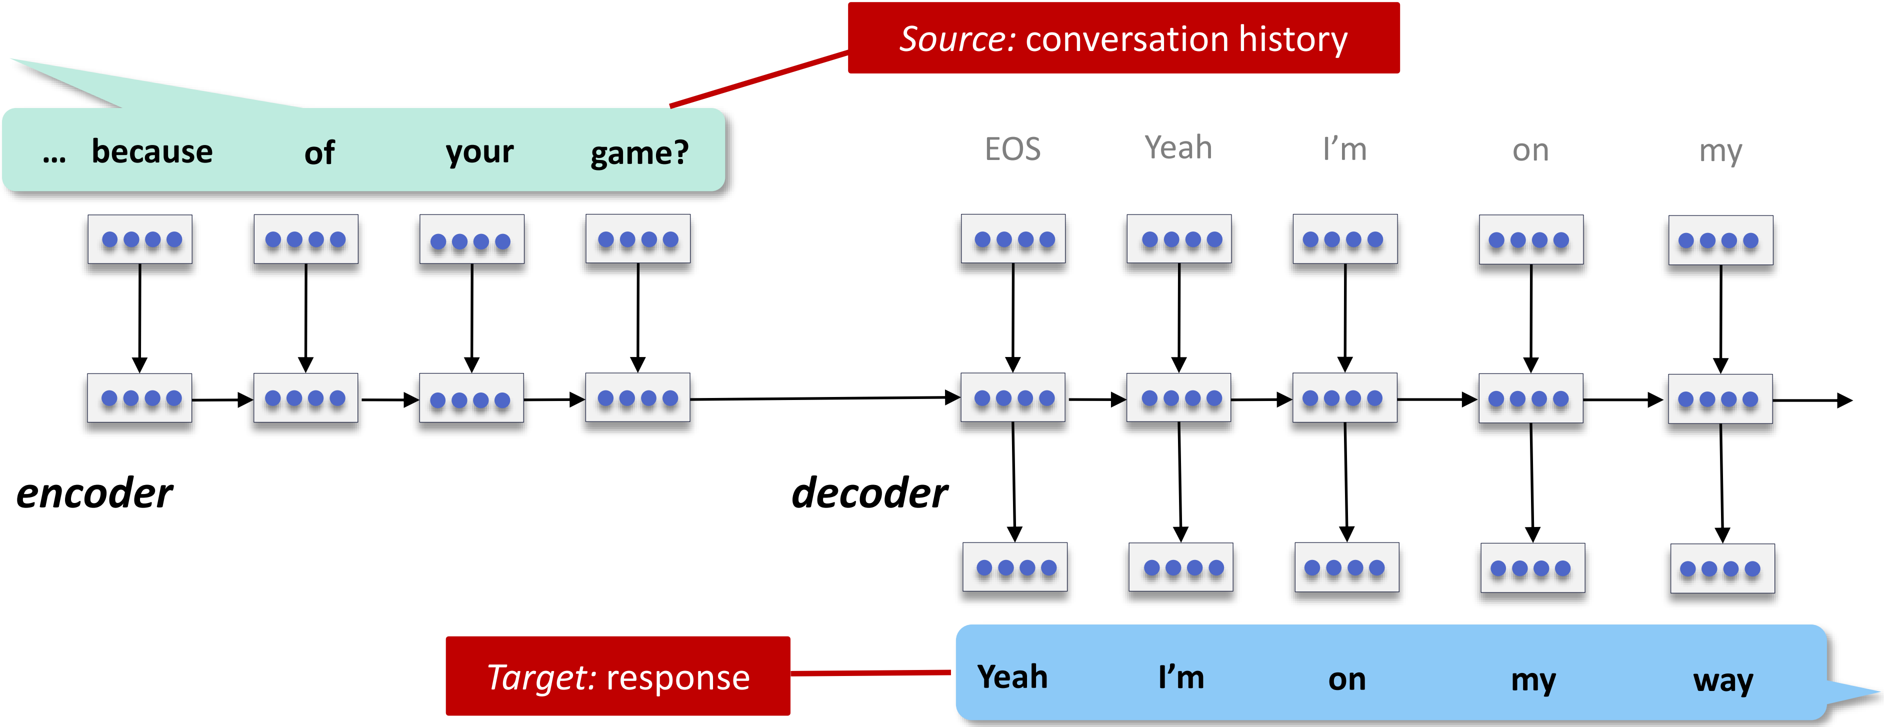

ChitChat Hierarchical Seq2Seq (Serban et al., 2016)

- Learns to generate dialogues from offline dialogs
- No state, action, intent, slot, etc.

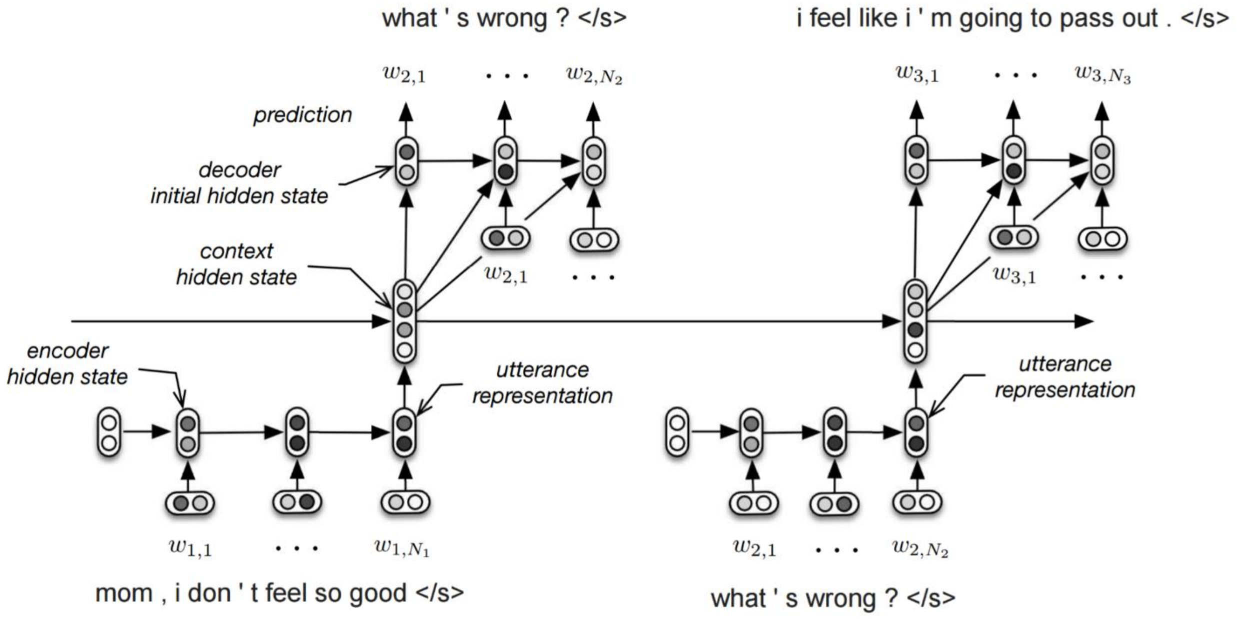

ChitChat Hierarchical Seq2Seq (Serban et.al., 2017)

- A hierarchical seq2seq model with Gaussian latent variable for generating dialogues (like topic or sentiment)

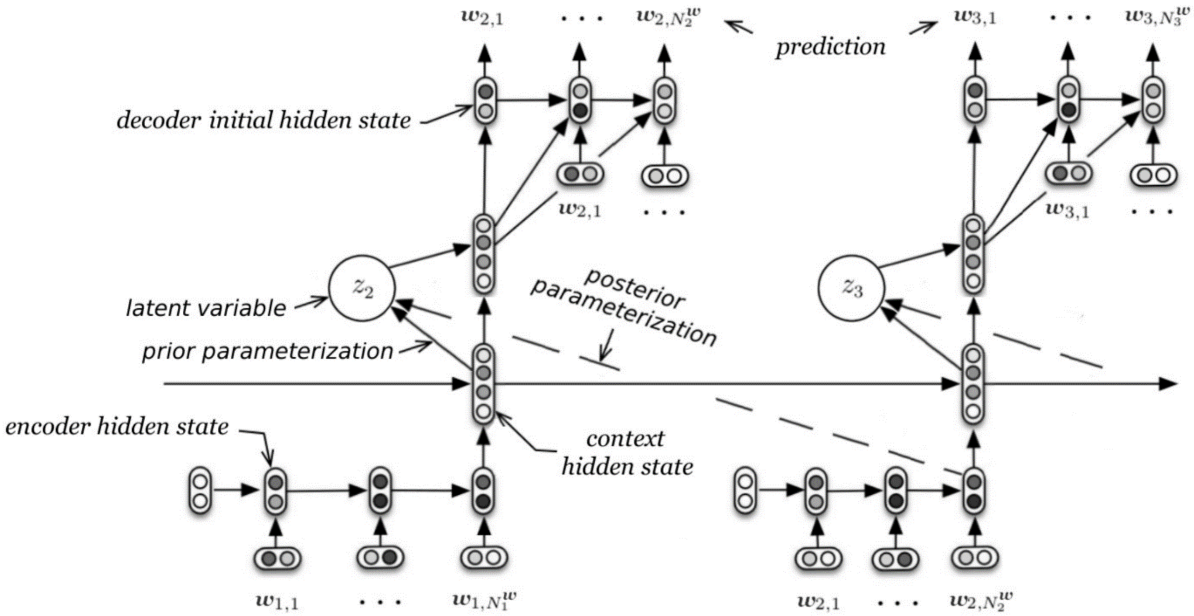

Neural Response Generation: Blandness Problem

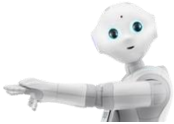

How was your weekend?

I don’t know.

What did you do?

I don’t understand what you are talking about.

This is getting boring…

Yes that’s what I’m saying.

The generated responses are general and meaningless

Outline

- PART I. Introduction & Background Knowledge
- PART II. Task-Oriented Dialogue Systems
- **PART III. Social Chat Bots**
  - Neural Response Generation
  - ***Response Diversity***
  - Response Consistency
  - Deep Reinforcement Learning for Response Generation
  - Combining Task-Oriented Bots and Social Chat Bots
- PART IV. Evaluation
- PART V. Recent Trends and Challenges

Mutual Information for Neural Generation ([Li et al., 2016](#))

- Mutual information objective

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

Bayes' rule

$$\hat{T} = \arg\max_{T} \left\{ \underbrace{\boxed{\log p(T|S)}}_{\text{standard likelihood}} - \underbrace{\boxed{\lambda \log p(T)}}_{\text{anti-LM}} \right\}$$

Bayes' theorem

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

$p(\text{target}|\text{source})$

$p(\text{source}|\text{target})$

Mutual Information for Neural Generation (Li et al., 2016)

Material: http://opendialogue.miulab.tw

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

antiLM:
large λ makes
output ungrammatical

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

used in rescoring
(intractable otherwise)

Lexical Diversity

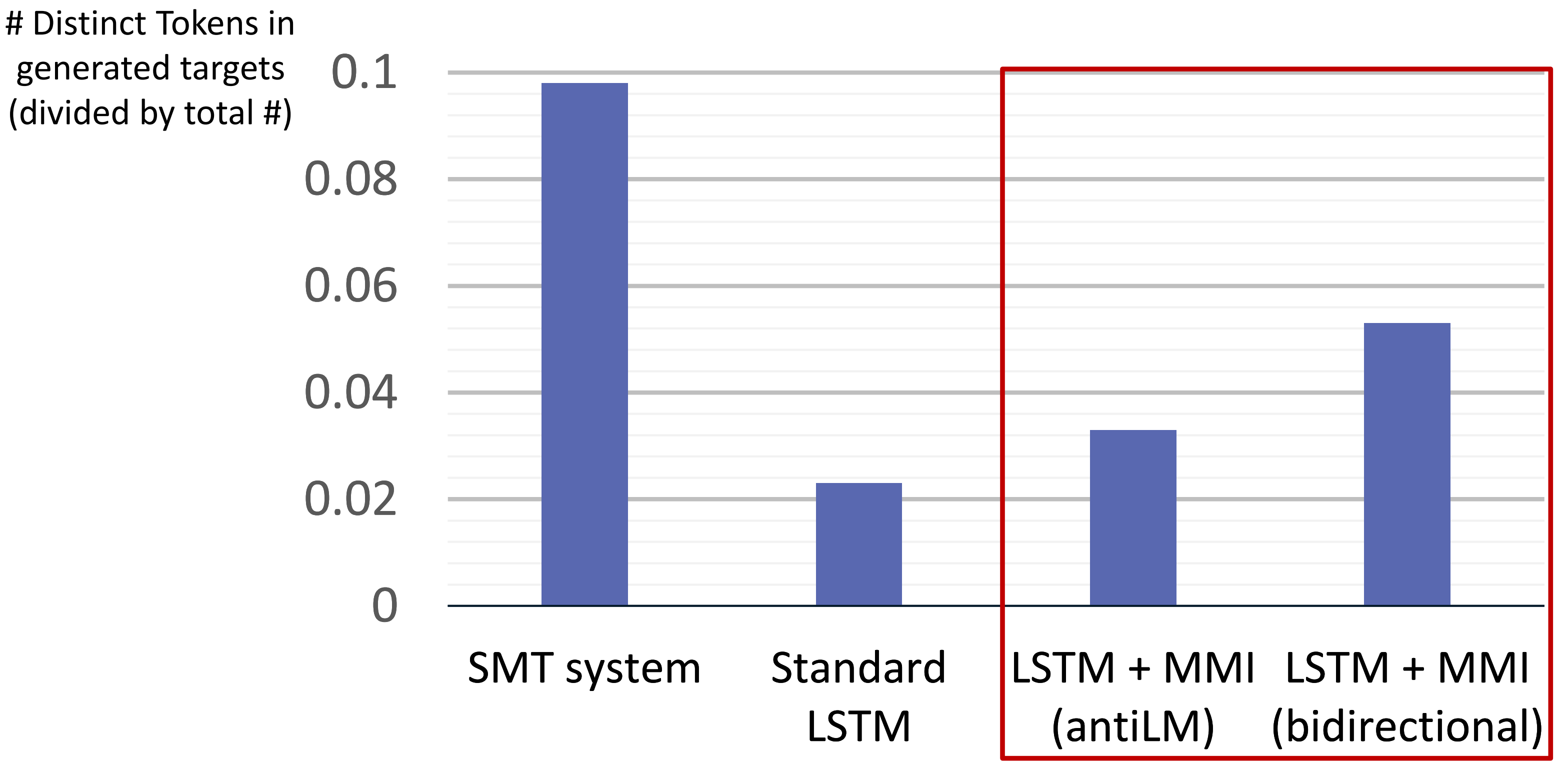

Sample Outputs (Baseline: Maximum Likelihood)

Material: http://opendialogue.miulab.tw

Wow sour starbursts really do make your mouth water... mm drool. Can I have one?

Of course!

Milan apparently selling Zlatan to balance the books... Where next, Madrid?

I don’t know.

‘tis a fine brew on a day like this! Stro

I'm not sure yet,

Well he was on in Bromley a wh

I don't even know what he's talking about.

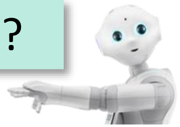

32% of the responses:

“I don’t know”

“I don’t know what you are talking about”

“I don’t think that is a good idea”

“Oh my god”

Sample Outputs (MMI)

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?

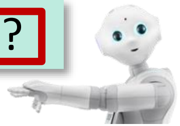

I think he'd be a good signing.

‘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.

Sample Outputs (MMI)

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.

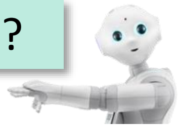

‘tis a fine brew on a day like this! Strong though, how many is sensible?

Depends on how much you drink!

Well he was on in Bromley a while ago... still touring.

I’ve never seen him live.

Outline

- PART I. Introduction & Background Knowledge
- PART II. Task-Oriented Dialogue Systems
- **PART III. Social Chat Bots**
  - Neural Response Generation
  - Response Diversity
  - ***Response Consistency***
  - Deep Reinforcement Learning for Response Generation
  - Combining Task-Oriented Bots and Social Chat Bots
- PART IV. Evaluation
- PART V. Recent Trends and Challenges

Real-World Conversations

Material: http://opendialogue.miulab.tw

- Multimodality
  - Conversation history
  - Persona
  - User profile data (bio, social graph, etc.)
  - Visual signal (camera, picture etc.)
  - Knowledge base
  - Mood
  - Geolocation
  - Time

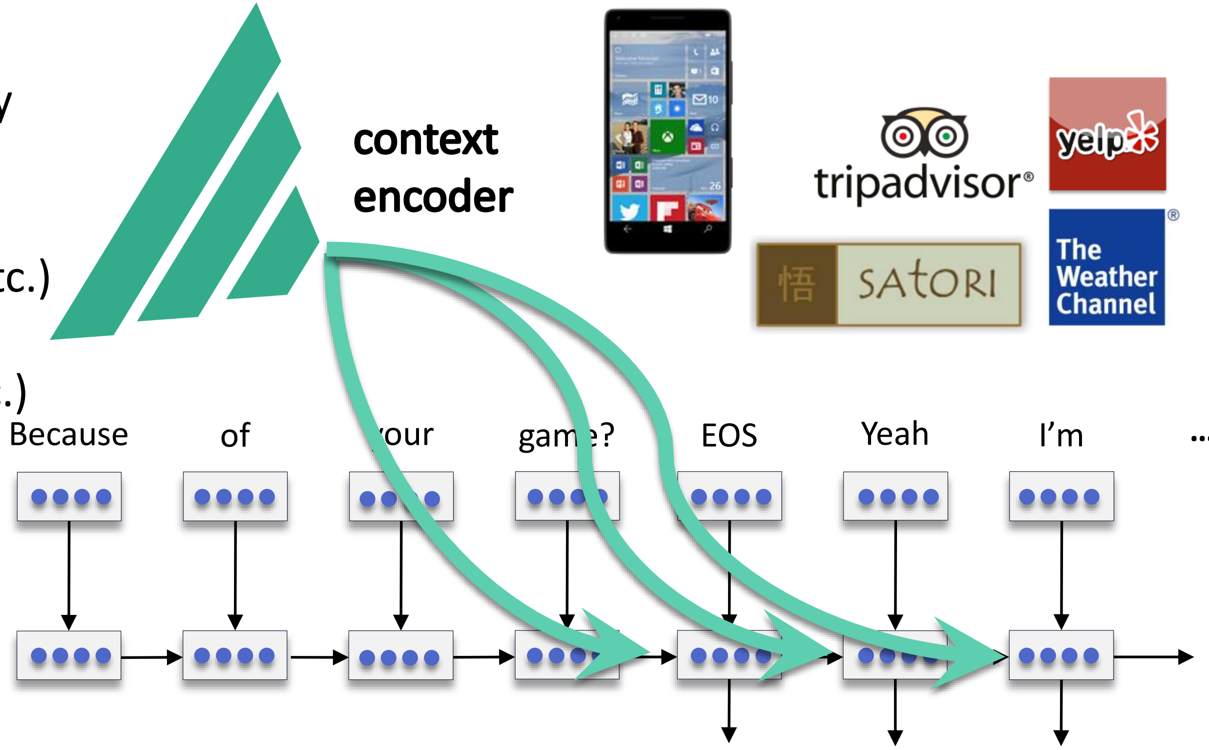

Personalized Response Generation (Li et al., 2016)

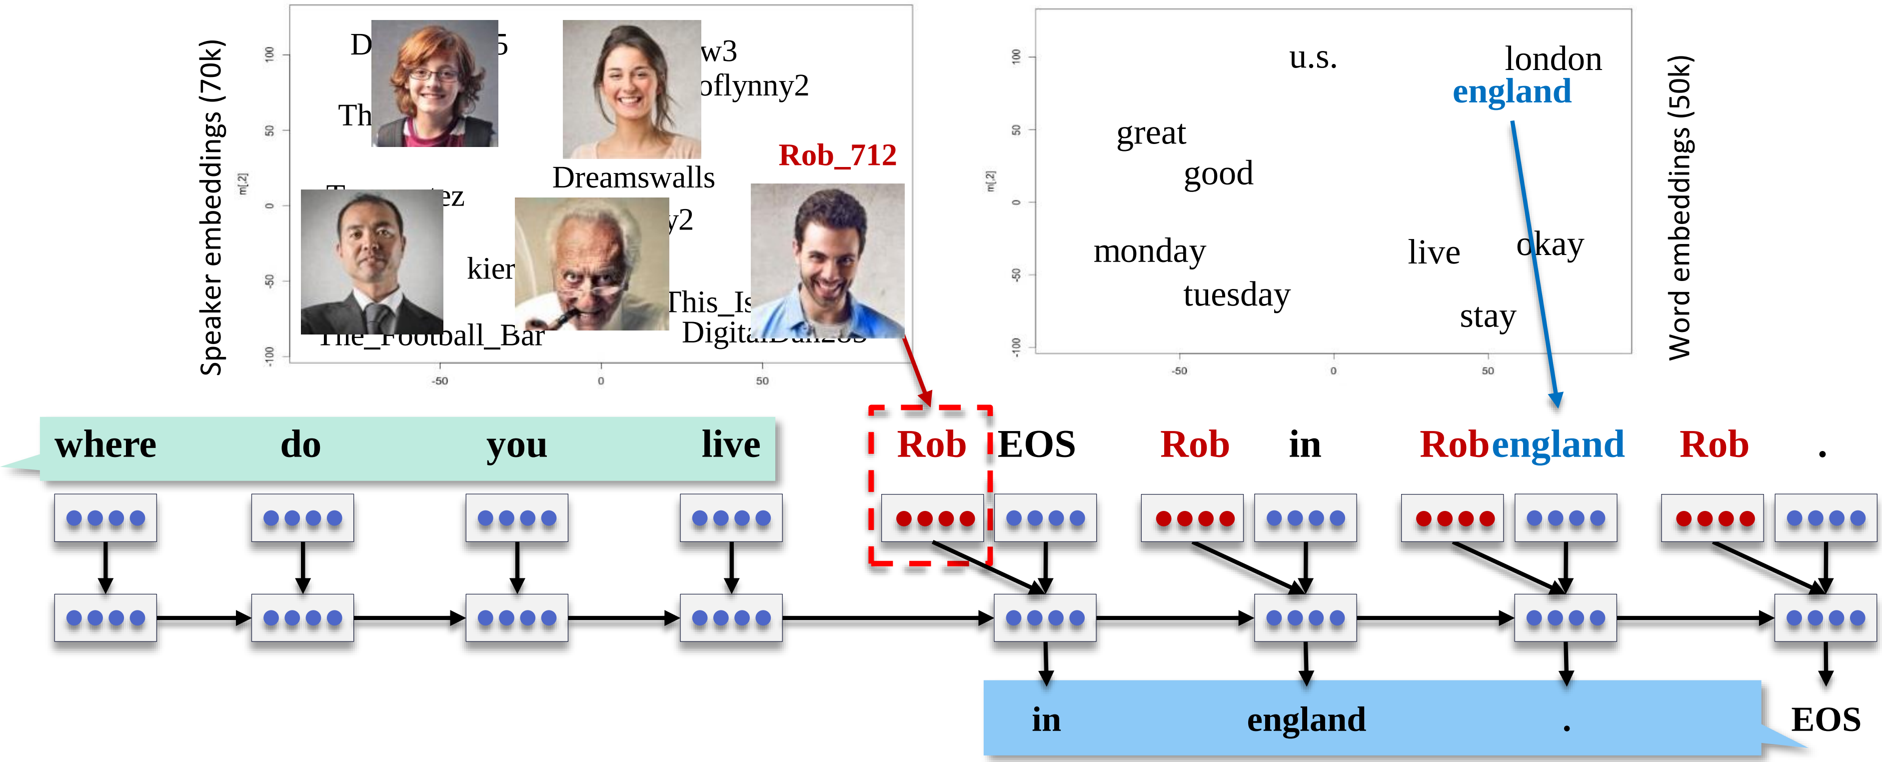

Persona Model for Speaker Consistency (Li et al., 2016)

Baseline model:

message	Where is your hometown?
response	I live in Indonesia.
message	Where are you from?
response	England, you?
message	In which city do you live now?
response	I live in Jakarta.
message	In which country do you live now?
response	I live in Indonesia.
message	What is your major?
response	Business, you?
message	What did you study in college?
response	Psychology, you?
message	How old are you?
response	I'm 18.
message	What is your age?
response	I'm 16.

Persona model using speaker embedding [Li+ 16b]

message	Where is your hometown?
response	I'm from England.
message	Where are you from?
response	I'm from England.
message	In which city do you live now?
response	I live in London.
message	In which country do you live now?
response	I live in England.
message	What is your major?
response	Business. You?
message	What did you study in college?
response	I did business studies.
message	How old are you?
response	I'm 18.
message	What is your age?
response	I'm 18.

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

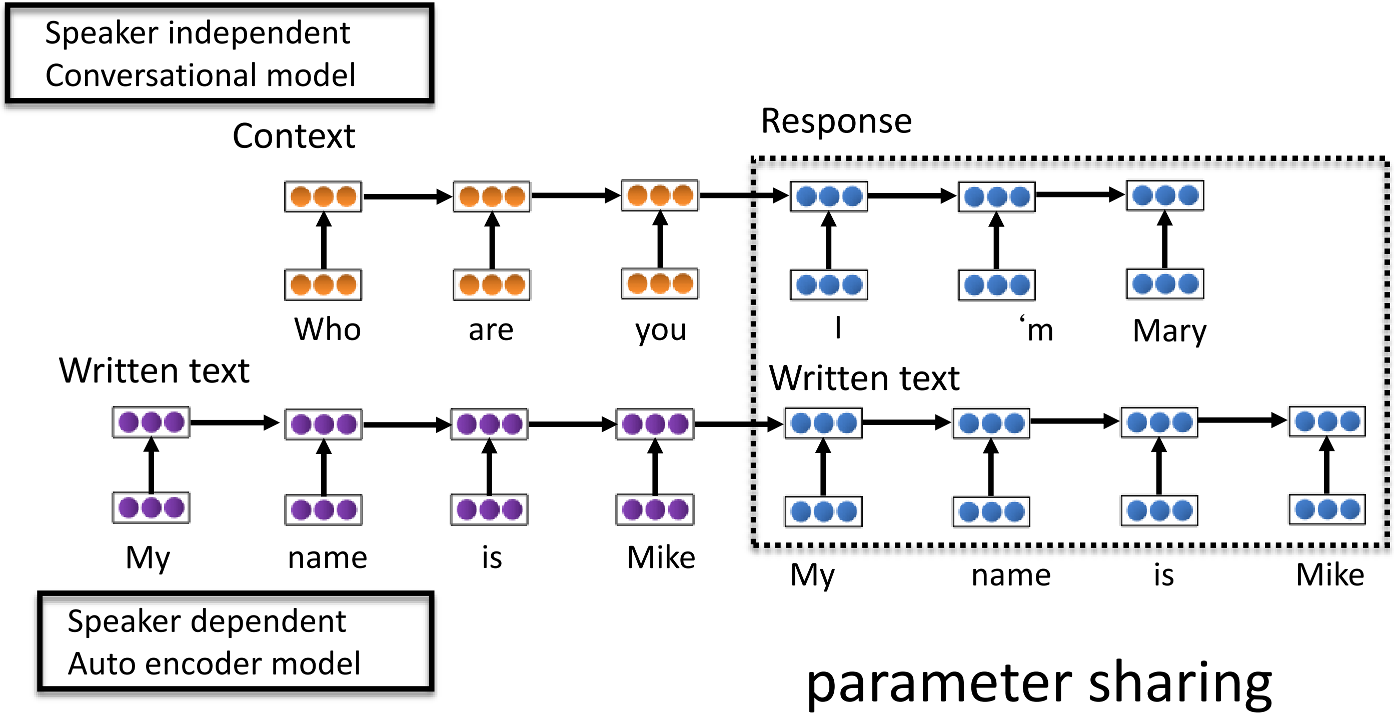

Speak-Role Aware Response (Luan et al., 2017)

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

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.

Outline

- PART I. Introduction & Background Knowledge
- PART II. Task-Oriented Dialogue Systems
- **PART III. Social Chat Bots**
  - Neural Response Generation
  - Response Diversity
  - Response Consistency
  - ***Deep Reinforcement Learning for Response Generation***
  - Combining Task-Oriented Bots and Social Chat Bots
- PART IV. Evaluation
- PART V. Recent Trends and Challenges

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

- RL agent generates more interactive responses
- RL agent tends to end a sentence with a question and hand the conversation over to the user
- Next step: explore intrinsic rewards, large-scale training

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 .

Outline

- PART I. Introduction & Background Knowledge
- PART II. Task-Oriented Dialogue Systems
- **PART III. Social Chat Bots**
  - Neural Response Generation
  - Response Diversity
  - Response Consistency
  - Deep Reinforcement Learning for Response Generation
  - ***Combining Task-Oriented Bots and Social Chat Bots***
- PART IV. Evaluation
- PART V. Recent Trends and Challenges

Chit-Chat v.s. Task-Oriented

Social Chat
Engaging, Human-Like Interaction
(Ungrounded)

Task-Oriented
Task Completion, Decision Support
(Grounded)

The weather is so depressing these days.

I know, I dislike rain too.
What about a day trip to eastern Washington?

Any recommendation?

Try Dry Falls, it’s spectacular!

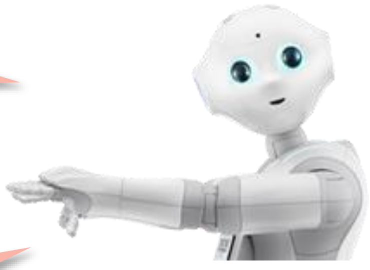

Image-Grounded Dialogues (Mostafazadeh et al., 2017)

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

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

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

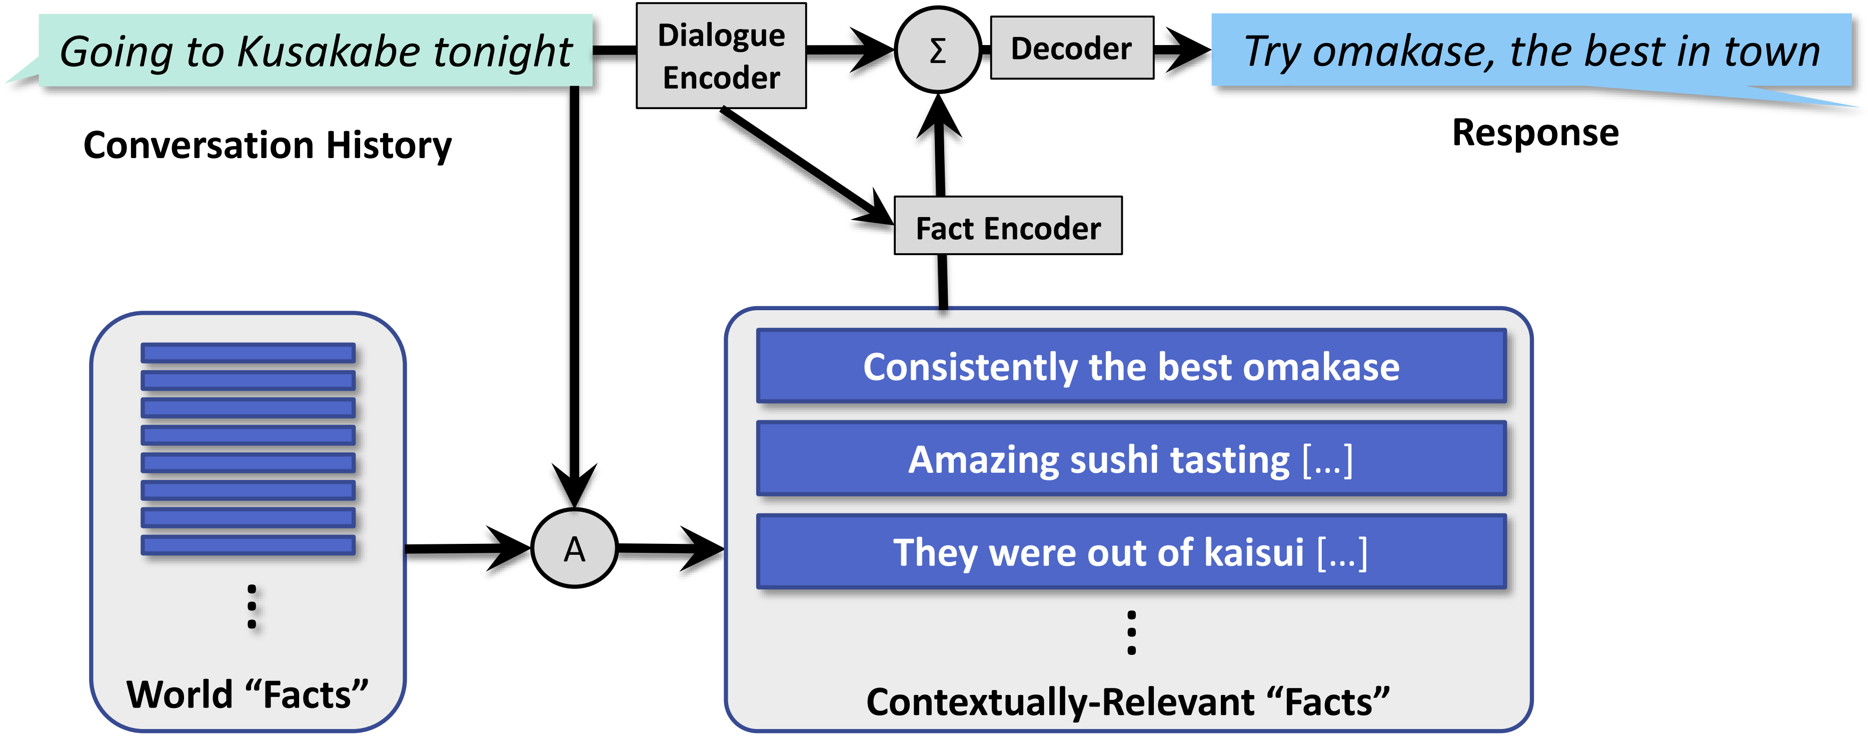

Conversation and Non-Conversation Data

Material: http://opendialogue.miulab.tw

*You know any good **A** restaurant in **B**?*

*Try **C**, one of the best **D** in the city.*

Conversation Data

Knowledge Resource

*You know any good **Japanese** restaurant in **Seattle**?*

*Try **Kisaku**, one of the best **sushi restaurants** in the city.*

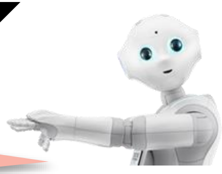

Knowledge-Grounded Responses (Ghazvininejad et al., 2017)

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)

Evaluation

Dialogue System Evaluation

Material: http://opendialogue.miulab.tw

- Dialogue model evaluation
  - Crowd sourcing
  - User simulator
- Response generator evaluation
  - Word overlap metrics
  - Embedding based metrics

Outline

- PART I. Introduction & Background Knowledge
- PART II. Task-Oriented Dialogue Systems
- PART III. Social Chat Bots
- **PART IV. Evaluation**
  - ***Human Evaluation***
  - User Simulation
  - Objective Evaluation
- PART V. Recent Trends and Challenges

Crowdsourcing for Dialogue System Evaluation (Yang et al., 2012)

Q1	**Do you think you understand from the dialog what the user wanted?**
Opt	1) No clue 2) A little bit 3) Somewhat 4) Mostly 5) Entirely
Aim	*elicit the Worker's confidence in his/her ratings.*
Q2	**Do you think the system is successful in providing the information that the user wanted?**
Opt	1) Entirely unsuccessful 2) Mostly unsuccessful 3) Half successful/unsuccessful 4) Mostly successful 5) Entirely successful
Aim	*elicit the Worker's perception of whether the dialog has fulfilled the informational goal of the user.*
Q3	**Does the system work the way you expect it?**
Opt	1) Not at all 2) Barely 3) Somewhat 4) Almost 5) Completely
Aim	*elicit the Worker's impression of whether the dialog flow suits general expectations.*
Q4	**Overall, do you think that this is a good system?**
Opt	1) Very poor 2) Poor 3) Fair 4) Good 5) Very good
Aim	*elicit the Worker's overall impression of the SDS.*
Q5	**What category do you think the dialog belongs to?**
Opt	1) Task is incomplete 2) Out of scope 3) Task is complete
Aim	*elicit the Worker's impression of whether the dialog reflects task completion.*

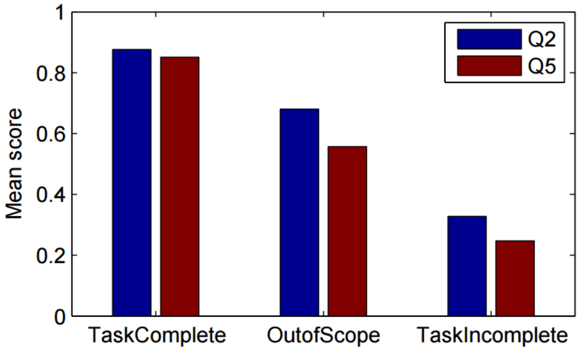

The normalized mean scores of Q2 and Q5 for approved ratings in each category. A higher score maps to a higher level of task success

Outline

- PART I. Introduction & Background Knowledge
- PART II. Task-Oriented Dialogue Systems
- PART III. Social Chat Bots
- **PART IV. Evaluation**
  - Human Evaluation
  - ***User Simulation***
  - Objective Evaluation
- PART V. Recent Trends and Challenges

User Simulation

- Goal: generate natural and reasonable conversations to enable reinforcement learning for exploring the policy space

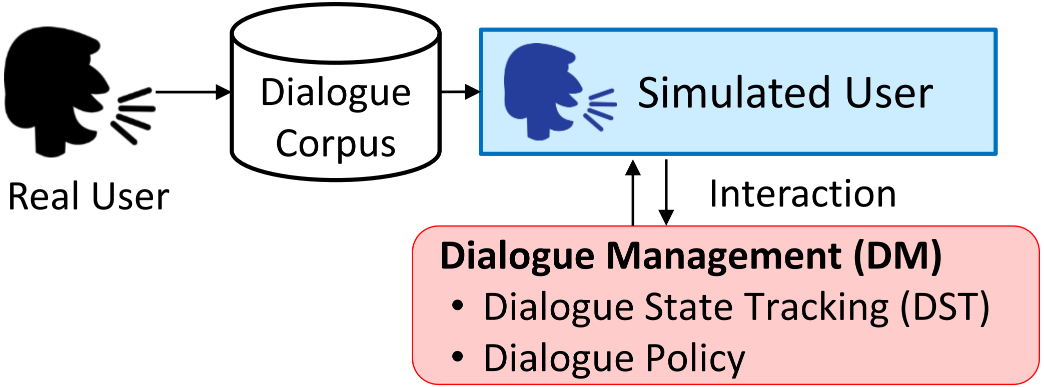

- Approach
  - Rule-based crafted by experts (Li et al., 2016)
  - Learning-based (Schatzmann et al., 2006; El Asri et al., 2016, Crook and Marin, 2017)

Elements of User Simulation

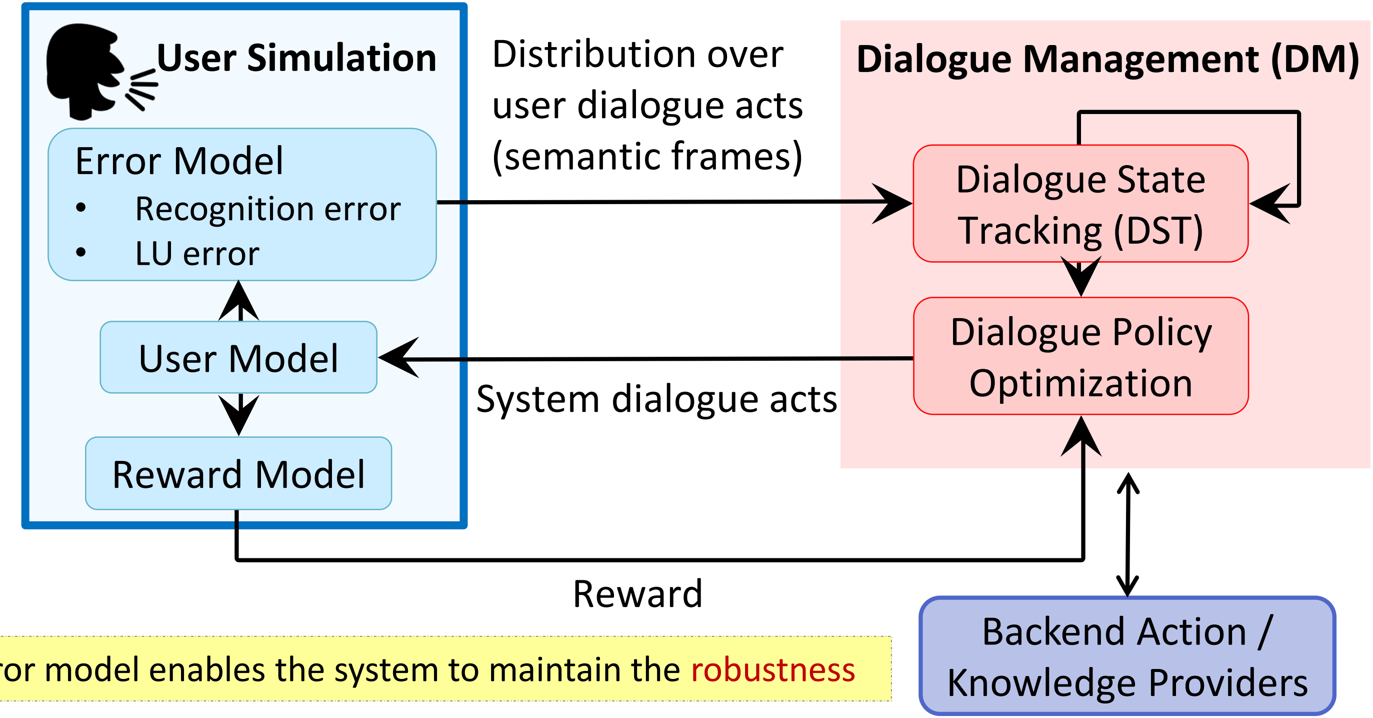

The error model enables the system to maintain the robustness

Rule-Based Simulator for RL Based System (Li et al., 2016)

- rule-based simulator + collected data
- starts with sets of goals, actions, KB, slot types
- publicly available simulation framework
- movie-booking domain: ticket booking and movie seeking
- provide procedures to add and test own agent

```
class AgentDQN(Agent):
    def run_policy(self, representation):
        """ epsilon-greedy policy """

        if random.random() < self.epsilon:
            return random.randint(0, self.num_actions - 1)
        else:
            if self.warm_start == 1:
                if len(self.experience_replay_pool) > self.experience_replay_pool_size:
                    self.warm_start = 2
                return self.rule_policy()
            else:
                return self.dqn.predict(representation, {}, predict_model=True)

    def train(self, batch_size=1, num_batches=100):
        """ Train DQN with experience replay """

        for iter_batch in range(num_batches):
            self.cur_bellman_err = 0
            for iter in range(len(self.experience_replay_pool)/(batch_size)):
                batch = [random.choice(self.experience_replay_pool) for i in xrange(batch_size)]
                batch_struct = self.dqn.singleBatch(batch, {'gamma': self.gamma}, self.clone_dqn)
```

Model-Based User Simulators

- Bi-gram models (Levin et.al. 2000)
- Graph-based models (Scheffler and Young, 2000)
- Data Driven Simulator (Jung et.al., 2009)
- Neural Models (deep encoder-decoder)

Seq2Seq User Simulation (El Asri et al., 2016)

- Seq2Seq trained from dialogue data
  - Input: c_i encodes contextual features, such as the previous system action, consistency between user goal and machine provided values
  - Output: a dialogue act sequence form the user
- Extrinsic evaluation for policy

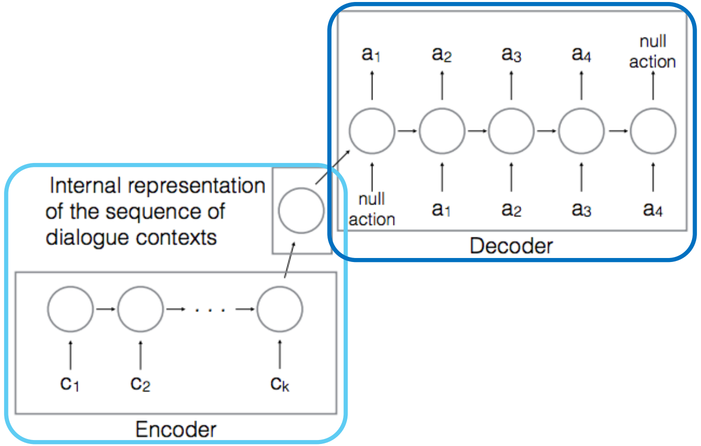

Seq2Seq User Simulation (Crook and Marin, 2017)

Material: http://opendialogue.miulab.tw

- Seq2Seq trained from dialogue data
  - No labeled data
  - Trained on just human to machine conversations

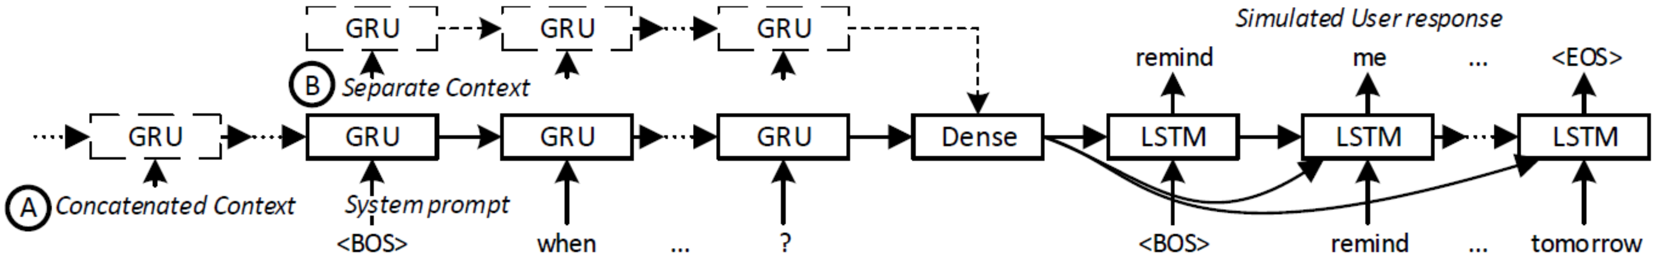

User Simulator for Dialogue Evaluation Measures

Material: http://opendialogue.miulab.tw

Understanding Ability

- whether constrained values specified by users can be understood by the system
- agreement percentage of system/user understandings over the entire dialog (averaging all turns)

Efficiency

- Number of dialogue turns
- Ratio between the dialogue turns (larger is better)

Action Appropriateness

- an explicit confirmation for an uncertain user utterance is an appropriate system action
- providing information based on misunderstood user requirements

Outline

- PART I. Introduction & Background Knowledge
- PART II. Task-Oriented Dialogue Systems
- PART III. Social Chat Bots
- **PART IV. Evaluation**
  - Human Evaluation
  - User Simulation
  - ***Objective Evaluation***
- PART V. Recent Trends and Challenges

How NOT to Evaluate Dialog System (Liu et al., 2017)

- How to evaluate the quality of the generated response ?
  - Specifically investigated for chat-bots
  - Crucial for task-oriented tasks as well

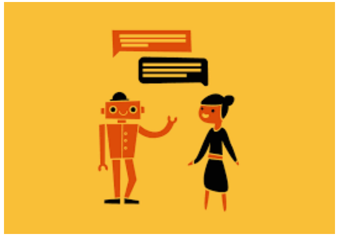

- Metrics:
  - Word overlap metrics, e.g., BLEU, METEOR, ROUGE, etc.
  - Embeddings based metrics, e.g., contextual/meaning representation between target and candidate

Dialogue Response Evaluation (Lowe et al., 2017)

- Problems of existing automatic evaluation
  - can be biased
  - correlate poorly with human judgements of response quality
  - using word overlap may be misleading
- Solution
  - collect a dataset of accurate human scores for variety of dialogue responses (e.g., coherent/un-coherent, relevant/irrelevant, etc.)
  - use this dataset to train an automatic dialogue evaluation model – learn to compare the reference to candidate responses!
  - Use RNN to predict scores by comparing against human scores!

Context of Conversation **Speaker A:** *Hey, what do you want to do tonight?* **Speaker B:** *Why don't we go see a movie?*
Model Response *Nah, let's do something active.*
Reference Response *Yeah, the film about Turing looks great!*

Recent Trends and Challenges

Multimodality

Dialogue Breath

Dialogue Depth

Outline

- PART I. Introduction & Background Knowledge
- PART II. Task-Oriented Dialogue Systems
- PART III. Social Chat Bots
- PART IV. Evaluation
- **PART V. Recent Trends and Challenges**
  - ***Multimodality***
  - Dialogue Breath
  - Dialogue Depth

Brain Signal for Understanding (Sridharan et al., 2012)

- Misunderstanding detection by brain signal
  - Green: listen to the correct answer
  - Red: listen to the wrong answer

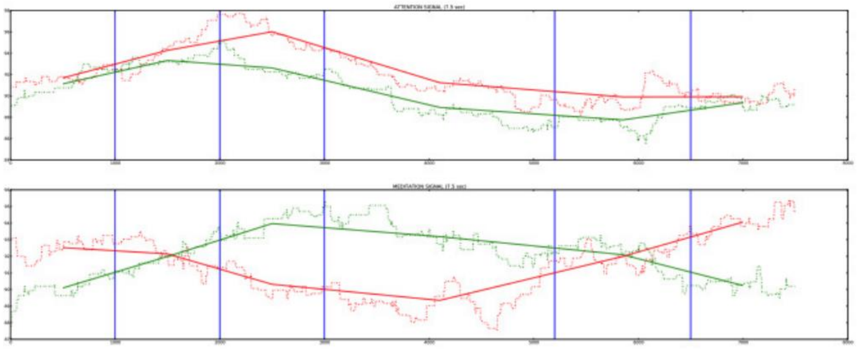

Detecting misunderstanding via brain signal in order to correct the understanding results

Video for Intent Understanding

Material: http://opendialogue.miulab.tw

I want to see a movie on TV!

Intent: turn_on_tv

Proactive (from camera)

May I turn on the TV for you?

Proactively understanding user intent to initiate the dialogues.

App Behavior for Understanding (Chen et al., 2015)

- Task: user intent prediction
- Challenge: language ambiguity

send to vivian

Communication

Email? v.s. **Message?**

① **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”

Considering behavioral patterns in history to model understanding for intent prediction.

Video Highlight Prediction ([Fu et al., 2017](#))

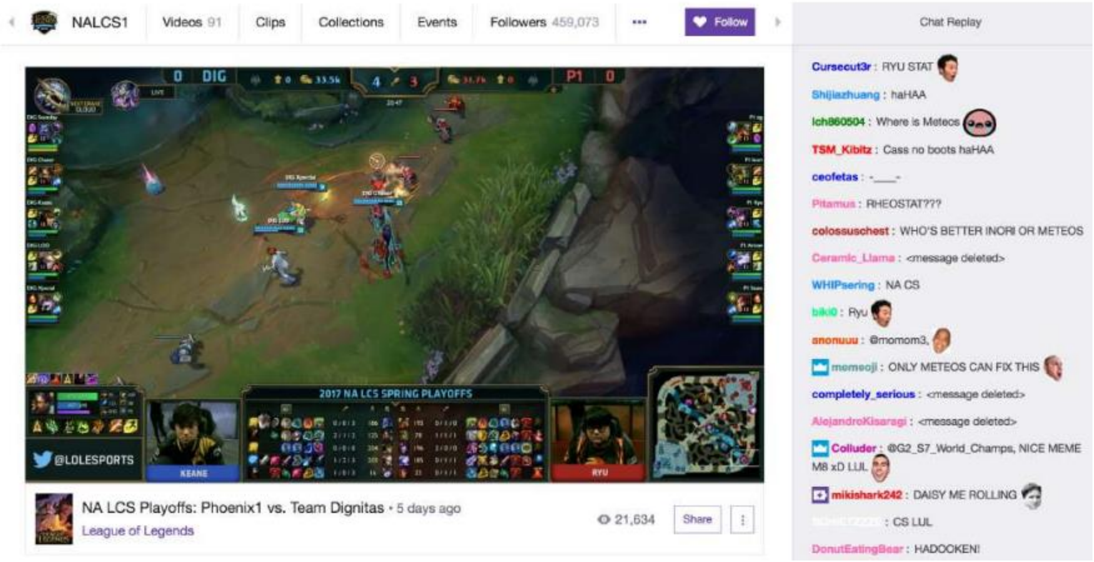

Video Highlight Prediction (Fu et al., 2017)

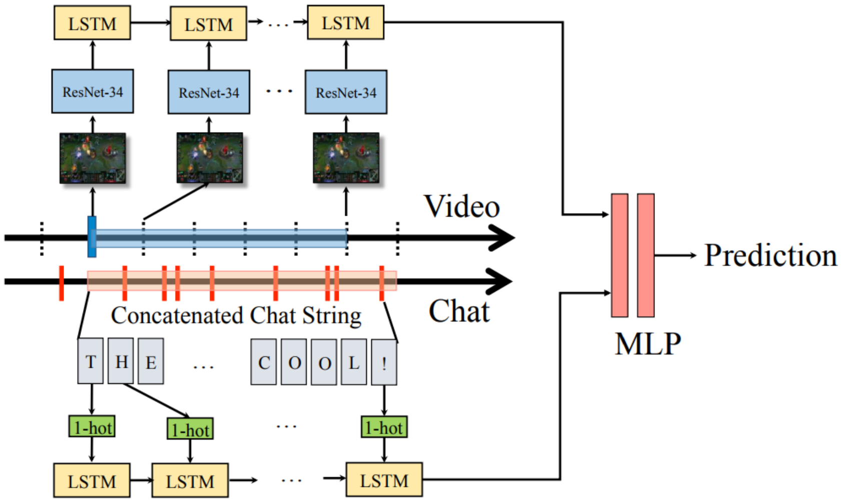

- Goal: predict highlight from the video
- Input : multi-modal and multi-lingual (real time text commentary from fans)
- Output: tag if a frame part of a highlight or not

Outline

- PART I. Introduction & Background Knowledge
- PART II. Task-Oriented Dialogue Systems
- PART III. Social Chat Bots
- PART IV. Evaluation
- **PART V. Recent Trends and Challenges**
  - Multimodality
  - ***Dialogue Breath***
  - Dialogue Depth

Evolution Roadmap

Material: http://opendialogue.miulab.tw

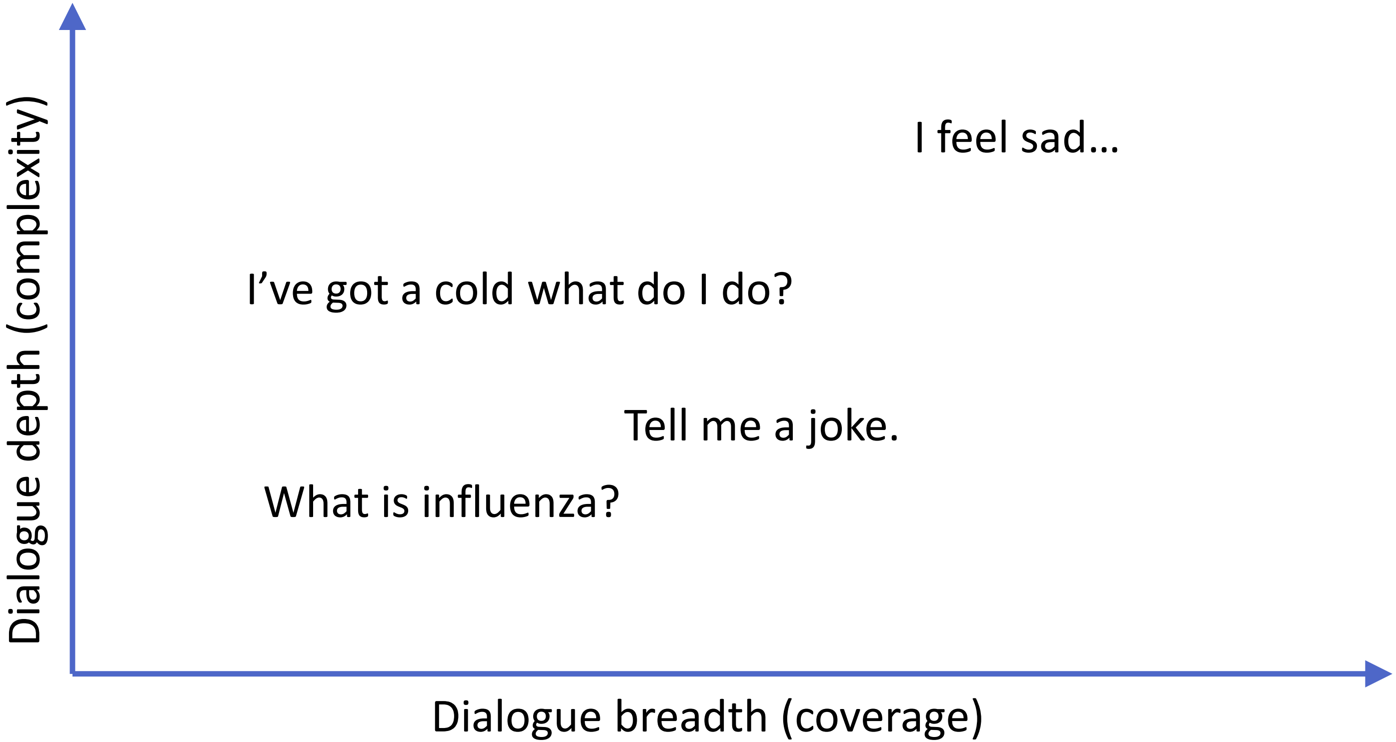

Evolution Roadmap

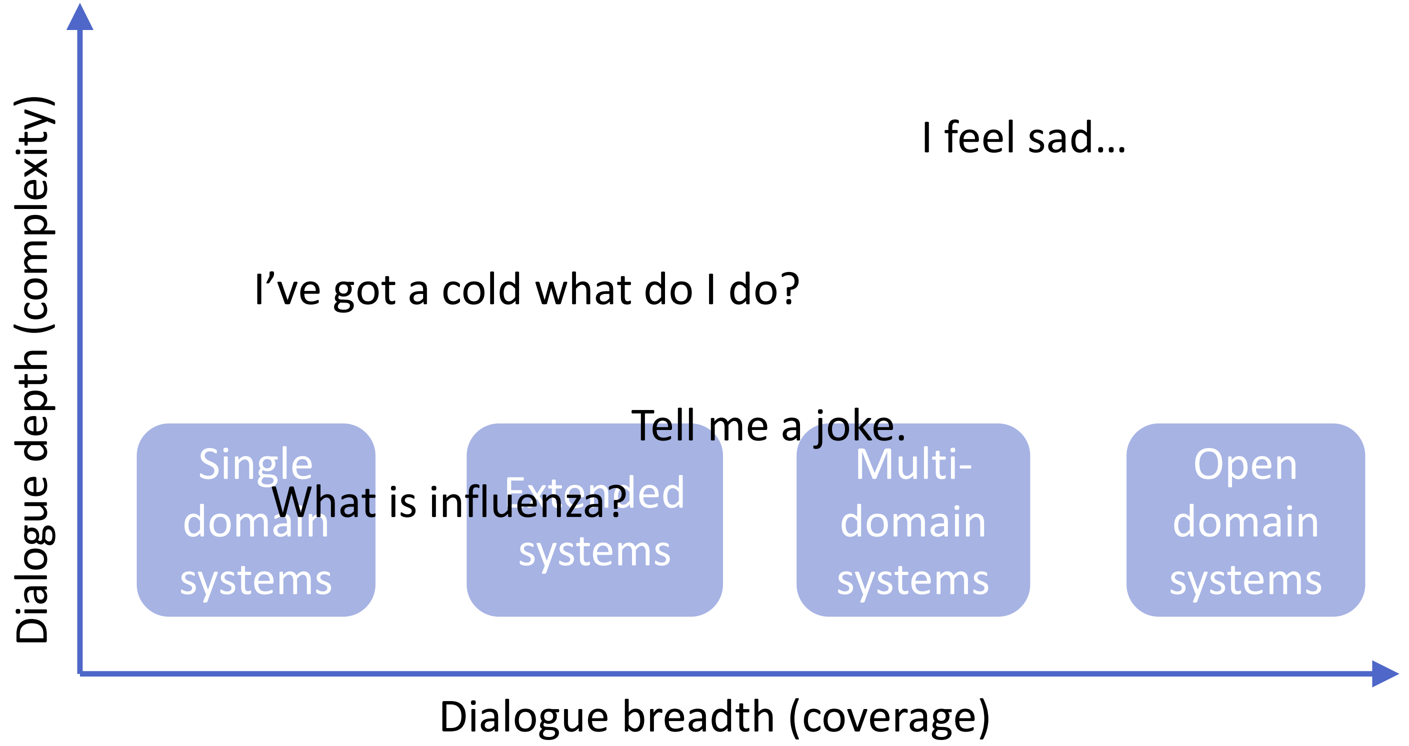

Intent Expansion (Chen et al., 2016)

- Transfer dialogue acts across domains
  - Dialogue acts are similar for multiple domains
  - Learning new intents by information from other domains

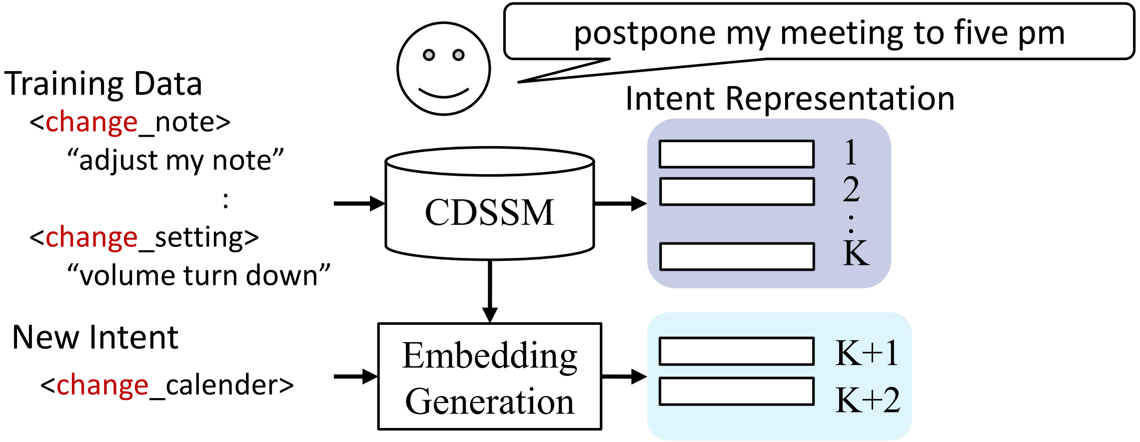

The dialogue act representations can be automatically learned for other domains

Policy for Domain Adaptation (Gašić et al., 2015)

- Bayesian committee machine (BCM) enables estimated Q-function to share knowledge across domains

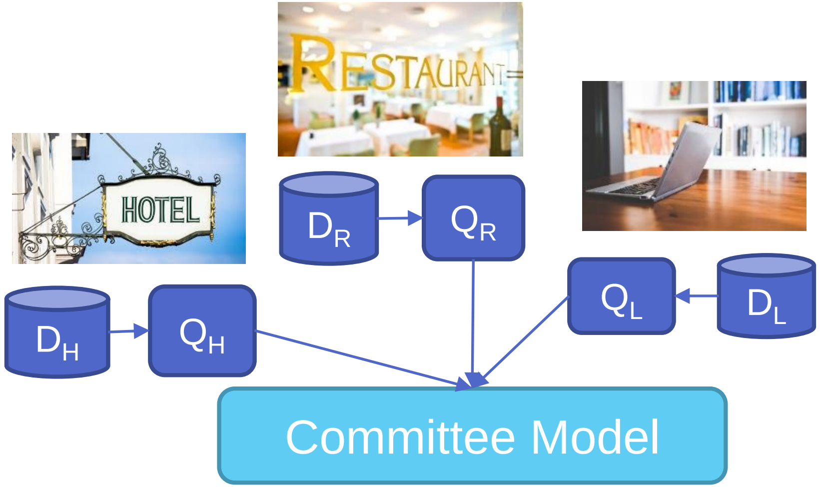

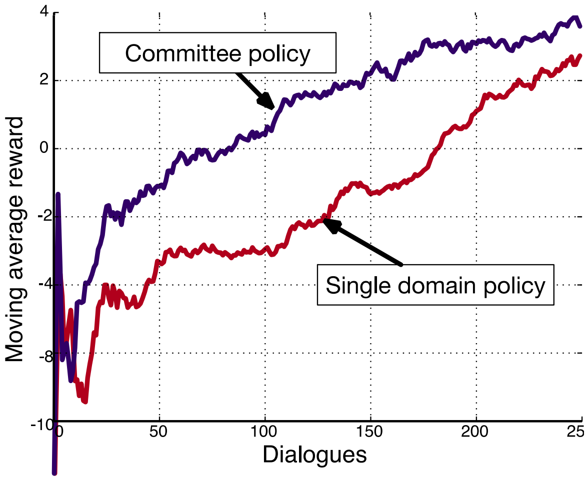

The policy from a new domain can be boosted by the committee policy

Outline

- PART I. Introduction & Background Knowledge
- PART II. Task-Oriented Dialogue Systems
- PART III. Social Chat Bots
- PART IV. Evaluation
- **PART V. Recent Trends and Challenges**
  - Multimodality
  - Dialogue Breath
  - ***Dialogue Depth***

Evolution Roadmap

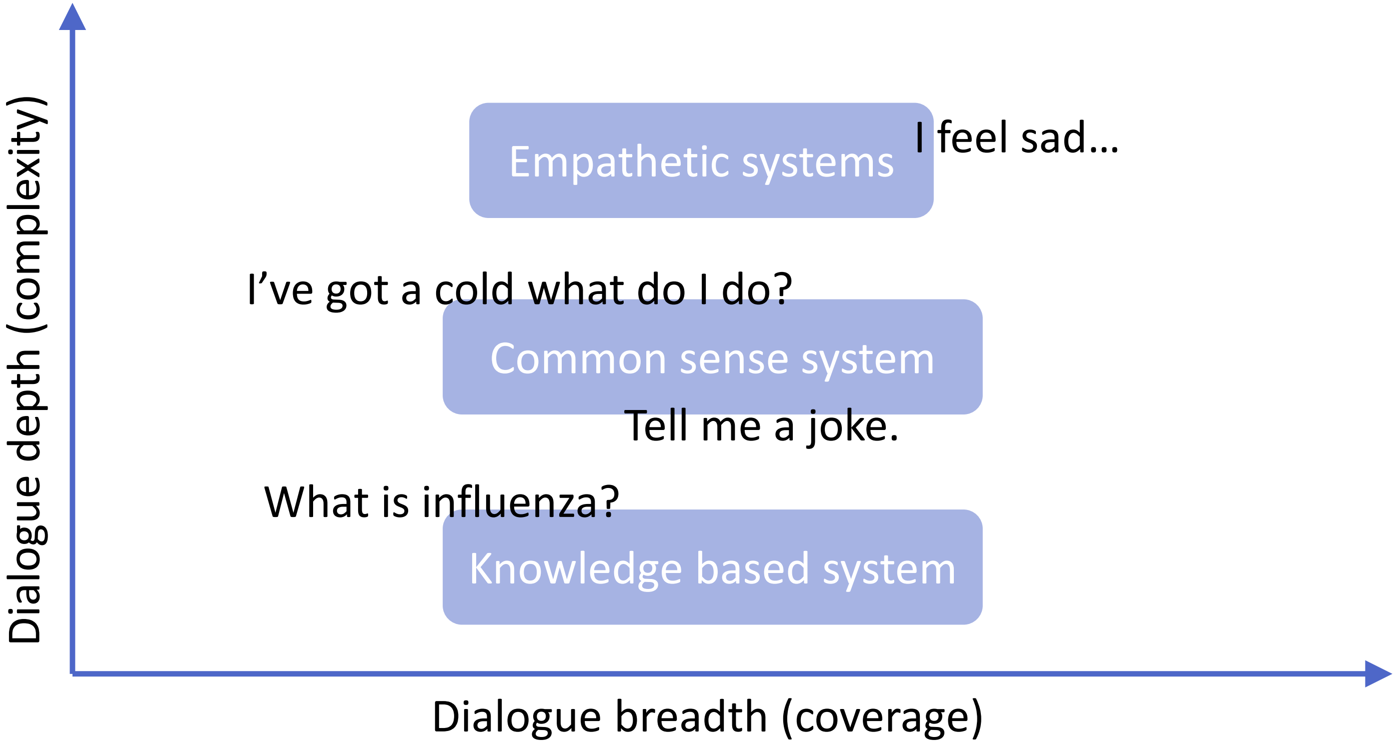

High-Level Intention for Dialogue Planning (Sun et al., 2016)

- High-level intention may span several domains

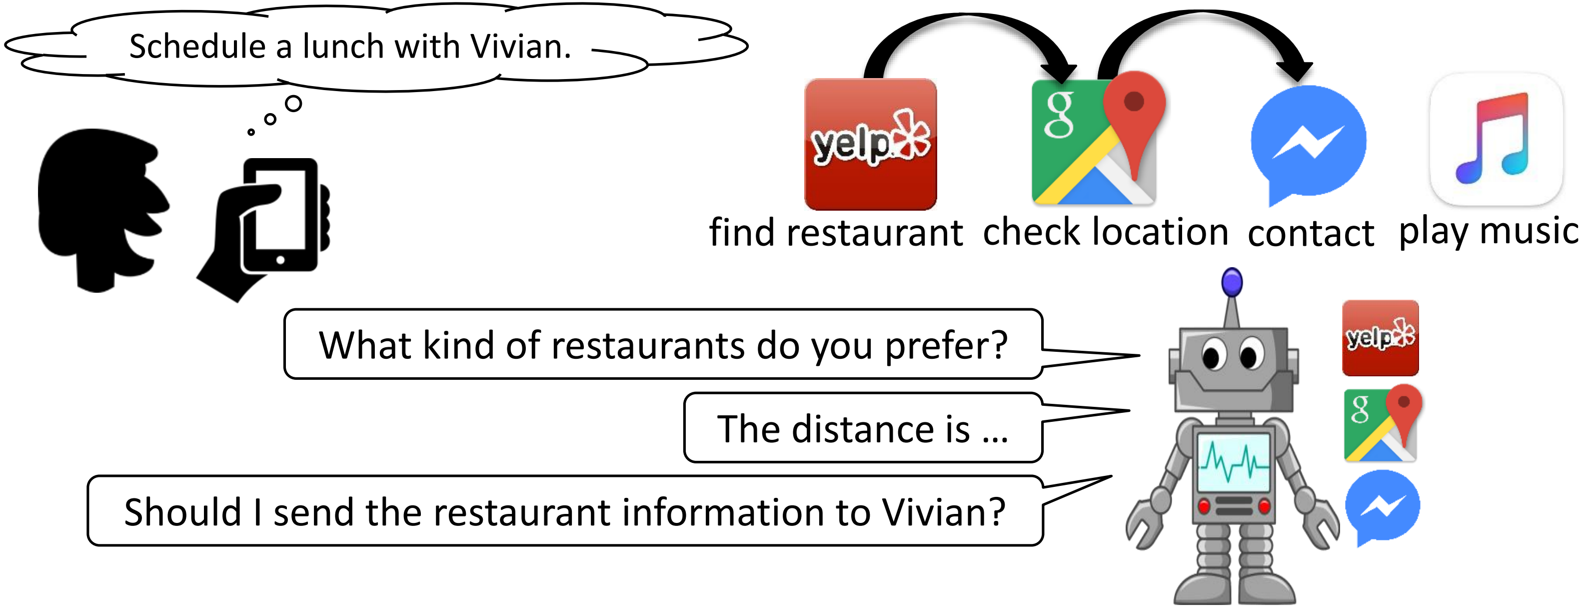

Users can interact via high-level descriptions and the system learns how to plan the dialogues

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

Material: http://opendialogue.miulab.tw

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

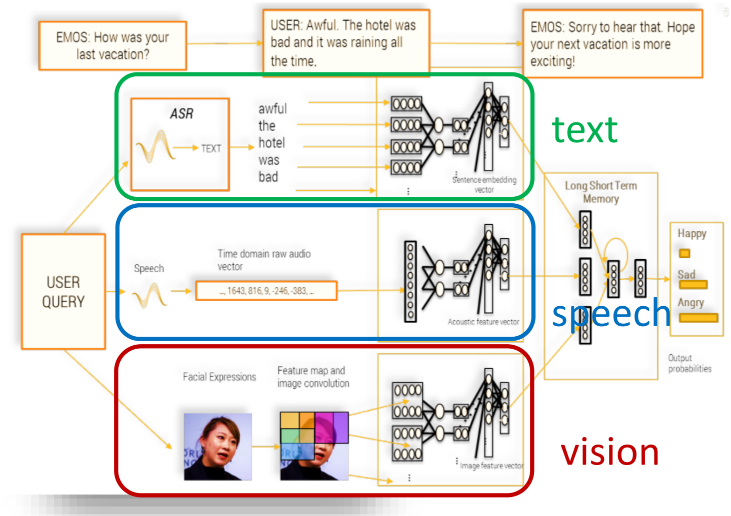

Emotion Recognizer

Zara - The Empathetic Supergirl

Face recognition output (index):1728
(index):1729
{
"recognition": "Race: Asian Confidence: 65.42750000000001 Smiling: 3.95896 Gender: Female Confidence: 88.9369",
"race": "Asian",
"race_confidence": "65.42750000000001",
"smiling": "3.95896",
"gender": "Female",
"gender_confidence": "88.9369"
}

Visual Object Discovery through Dialogues (Vries et al., 2017)

- Recognize objects using “Guess What?” game
- Includes “spatial”, “visual”, “object taxonomy” and “interaction”

Is it a person?	No
Is it an item being worn or held?	Yes
Is it a snowboard?	Yes
Is it the red one?	No
Is it the one being held by the person in blue?	Yes

Is it a cow?	Yes
Is it the big cow in the middle?	No
Is the cow on the left?	No
On the right ?	Yes
First cow near us?	Yes

Conclusion

Summarized Challenges

- Human-machine interfaces is a hot topic but building a good one is challenging!
- Most state-of-the-art technologies are based on DNN
  - Requires huge amounts of labeled data
  - Several frameworks/models are available
- Leveraging structured knowledge and unstructured data
- Handling reasoning
- Data collection and analysis from un-structured data
- The capability of task-oriented and chit-chat dialogues should be integrated.

Brief Conclusions

- Introduce recent deep learning methods used in dialogue models
- Highlight main components of task-oriented dialogue systems and new deep learning architectures used for these components
- Highlight the challenges and trends for current chat bot research
- Talk about new avenues for current state-of-the-art dialogue research
- Provide all materials online!

http://opendialogue.miulab.tw

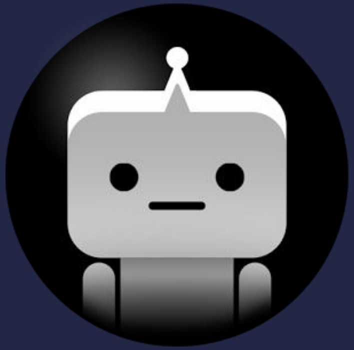

Thanks to Dilek Hakkani-Tur, Asli Celikyilmaz, Tsung-Hsien Wen, Pei-Hao Su, Li Deng, Sungjin Lee, Milica Gašić, Lihong Li, Xiujin Li, Abhinav Rastogi, Ankur Bapna, PArarth Shah and Gokhan Tur for sharing their slides.

THANKS FOR ATTENTION!

Q & A