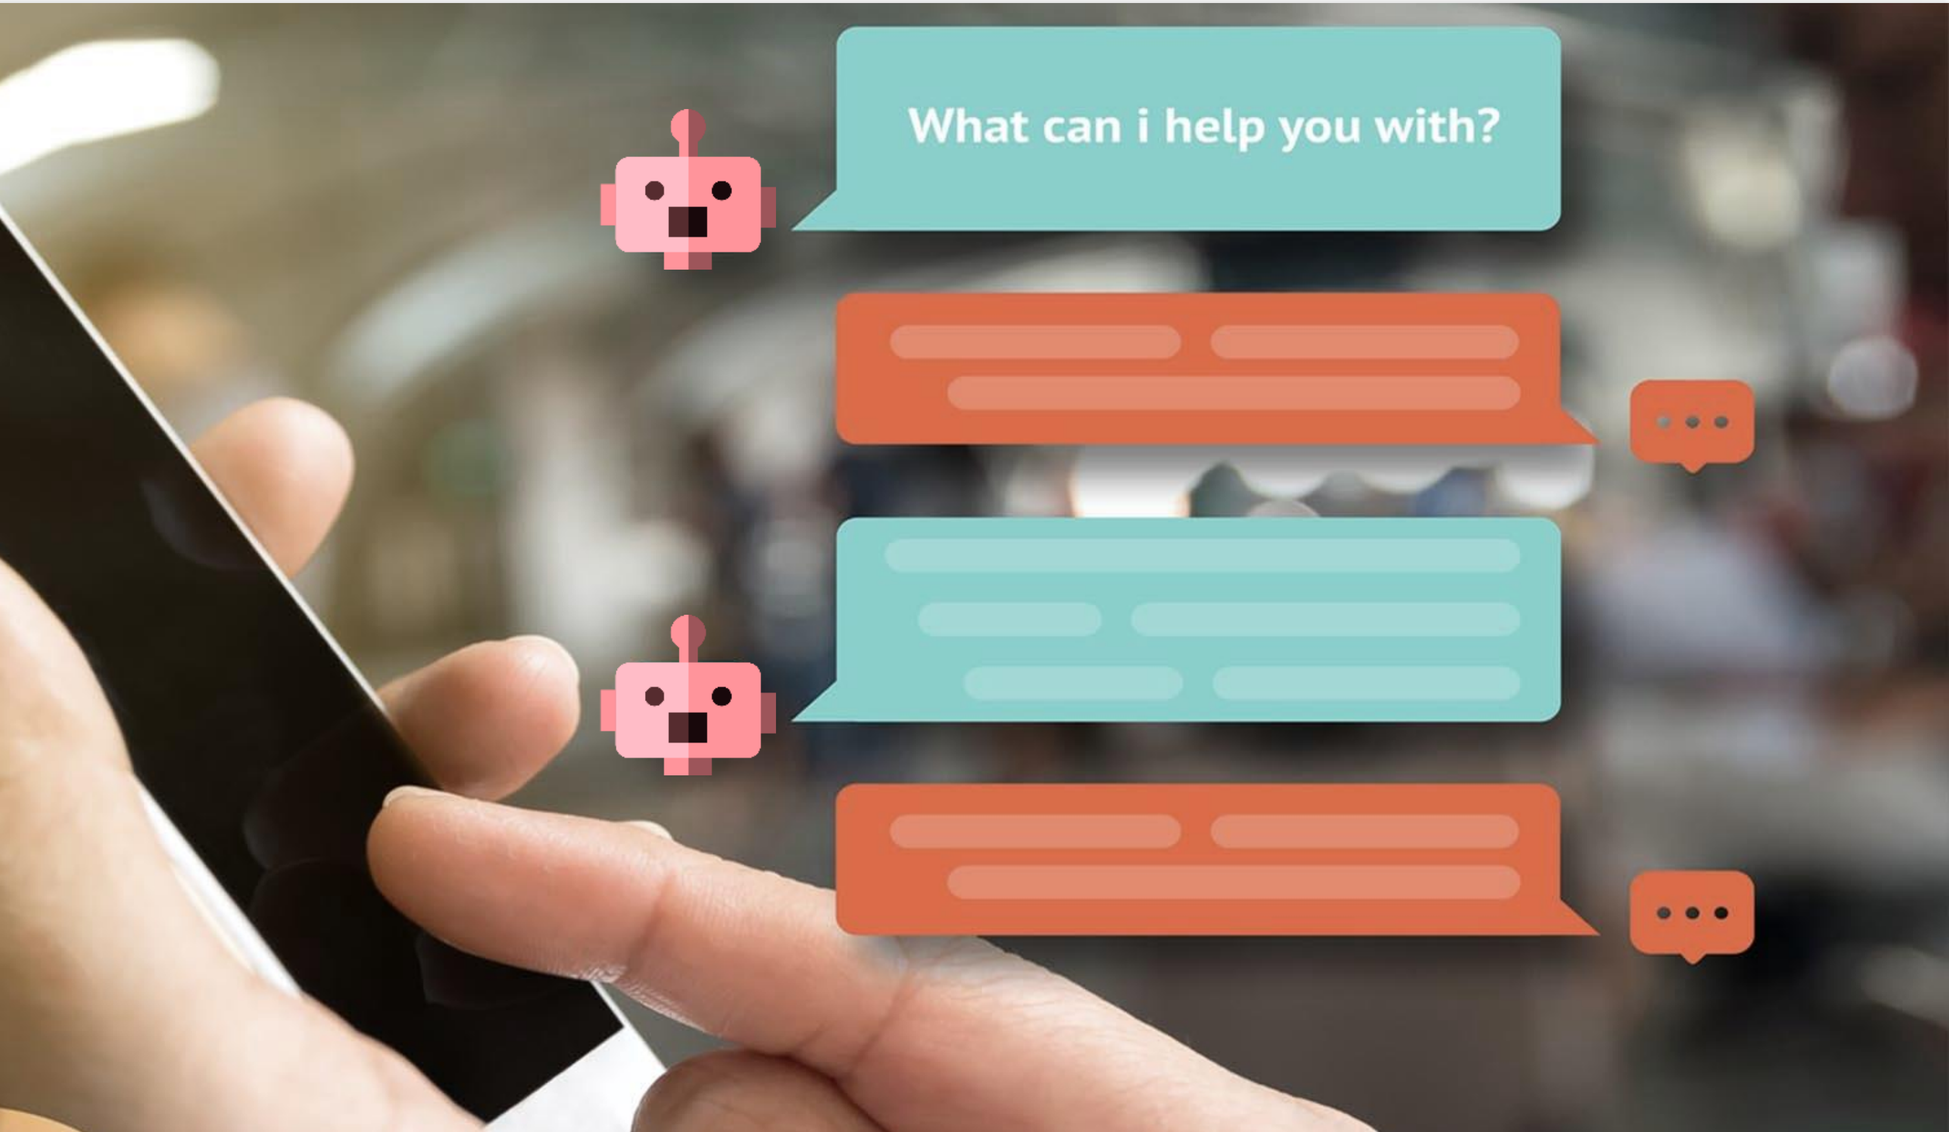

Deep Learning for Dialogue Systems

YUN-NUNG (VIVIAN) CHEN 陳縕儂

HTTP://VIVIANCHEN.IDV.TW

國立臺灣大學
National Taiwan University

Outline

- Introduction
- Background Knowledge
- Modular Dialogue System
- System Evaluation
- Recent Trends of Learning Dialogues

Outline

- **Introduction**
- Background Knowledge
- Modular Dialogue System
- System Evaluation
- Recent Trends of Learning Dialogues

Brief History of Dialogue Systems

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

Early 1990s

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

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

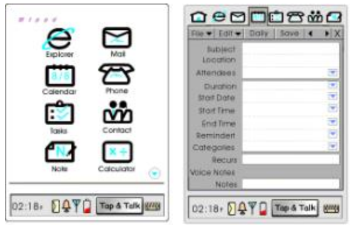

Early 2000s

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"

TV Voice Search
e.g., Bing on Xbox

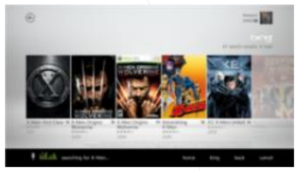

DARPA
CALO Project

2017

Virtual Personal Assistants

Language Empowering Intelligent Assistant

Apple Siri (2011)

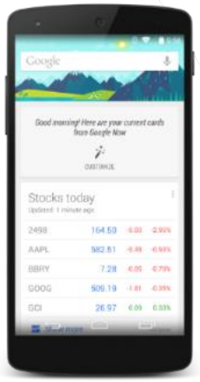

Google Now (2012)
Google Assistant (2016)

Microsoft Cortana (2014)

Amazon Alexa/Echo (2014)

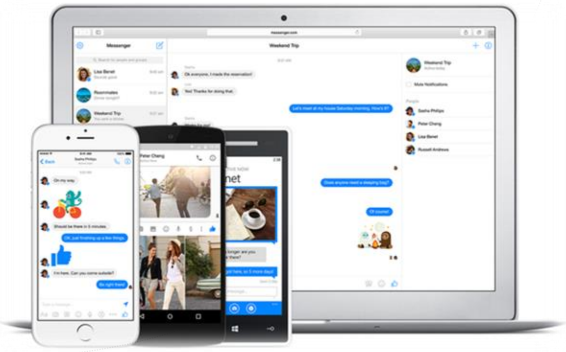

Facebook M & Bot (2015)

Google Home (2016)

Apple HomePod (2017)

Why We Need?

- Get things done
  - E.g. set up alarm/reminder, take note
- Easy access to structured data, services and apps
  - E.g. find docs/photos/restaurants
- Assist your daily schedule and routine
  - E.g. commute alerts to/from work
- Be more productive in managing your work and personal life




NTU MIULAB

Why Natural Language?

- Global Digital Statistics (2018 January)

Total Population
7.59B

Internet Users
4.02B
7%

Active Social Media Users
3.20B
13%

Unique Mobile Users
5.14B
4%

Active Mobile Social Users
2.96B
14%

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

Spoken Dialogue System (SDS)

- 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

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

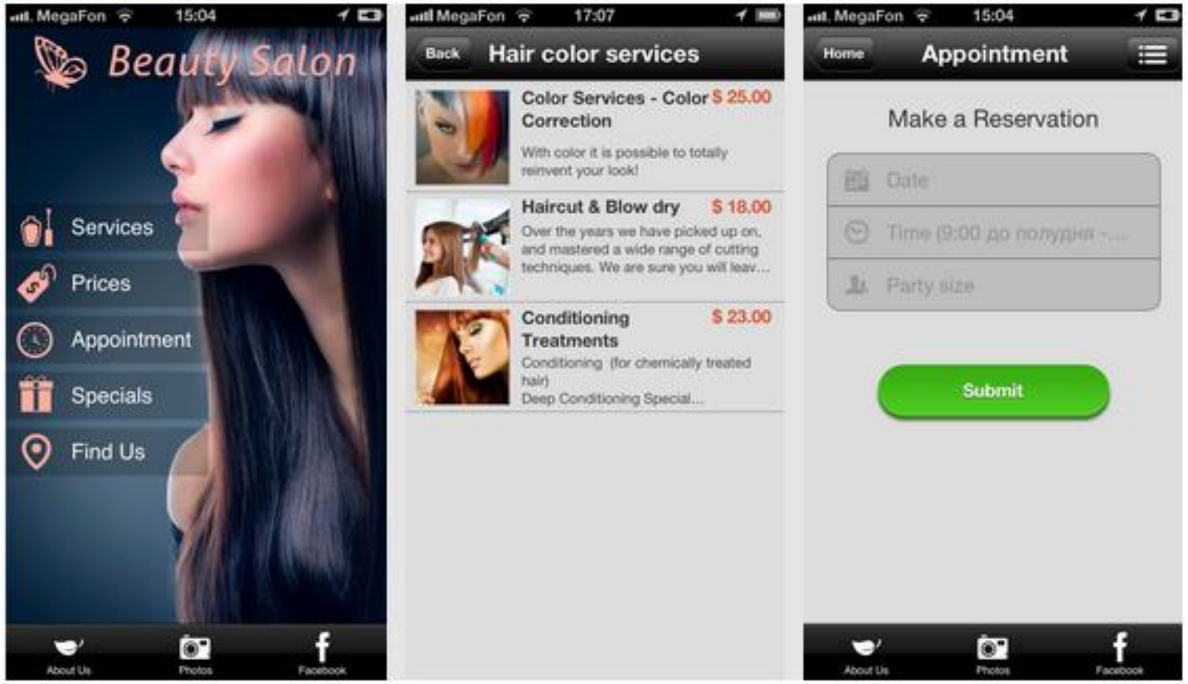

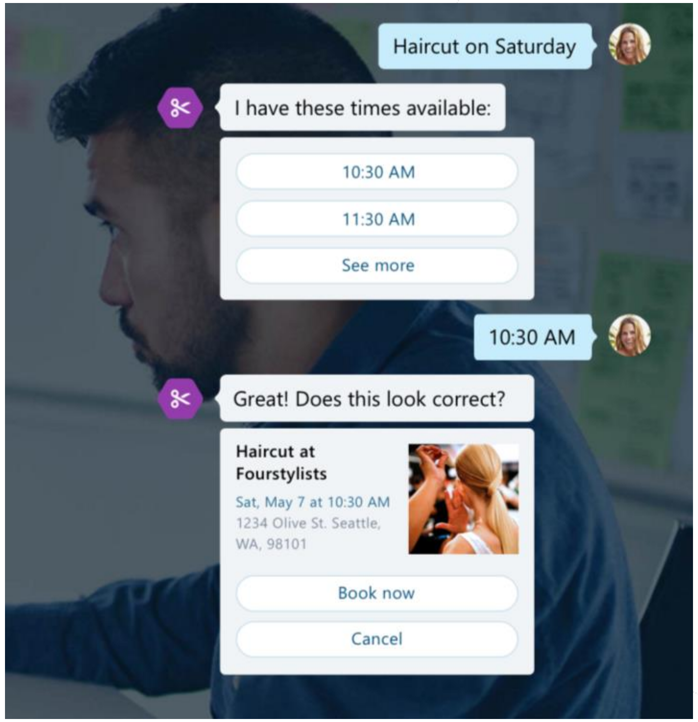

Users can initiate dialogues instead of following the GUI design

Material: http://deepdialogue.miulab.tw

GUI v.s. CUI (Conversational UI)

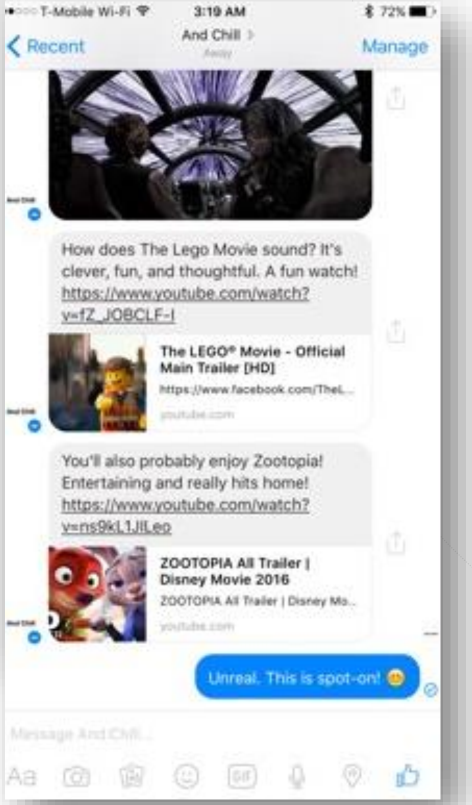

Material: http://deepdialogue.miulab.tw

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

Material: http://deepdialogue.miulab.tw

Conversational Agents

Chit-Chat

seq2seq models → Seq2seq with conversation contexts → Knowledge-grounded seq2seq models

Task-Oriented

Single-domain, system-initiative → Multi-domain, contextual, mixed-initiative → End-to-end learning, massively multi-domain

Challenges

- Variability in Natural Language
- Robustness
- Recall/Precision Trade-off
- Meaning Representation
- Common Sense, World Knowledge
- Ability to Learn
- Transparency

Task-Oriented Dialogue System (Young, 2000)

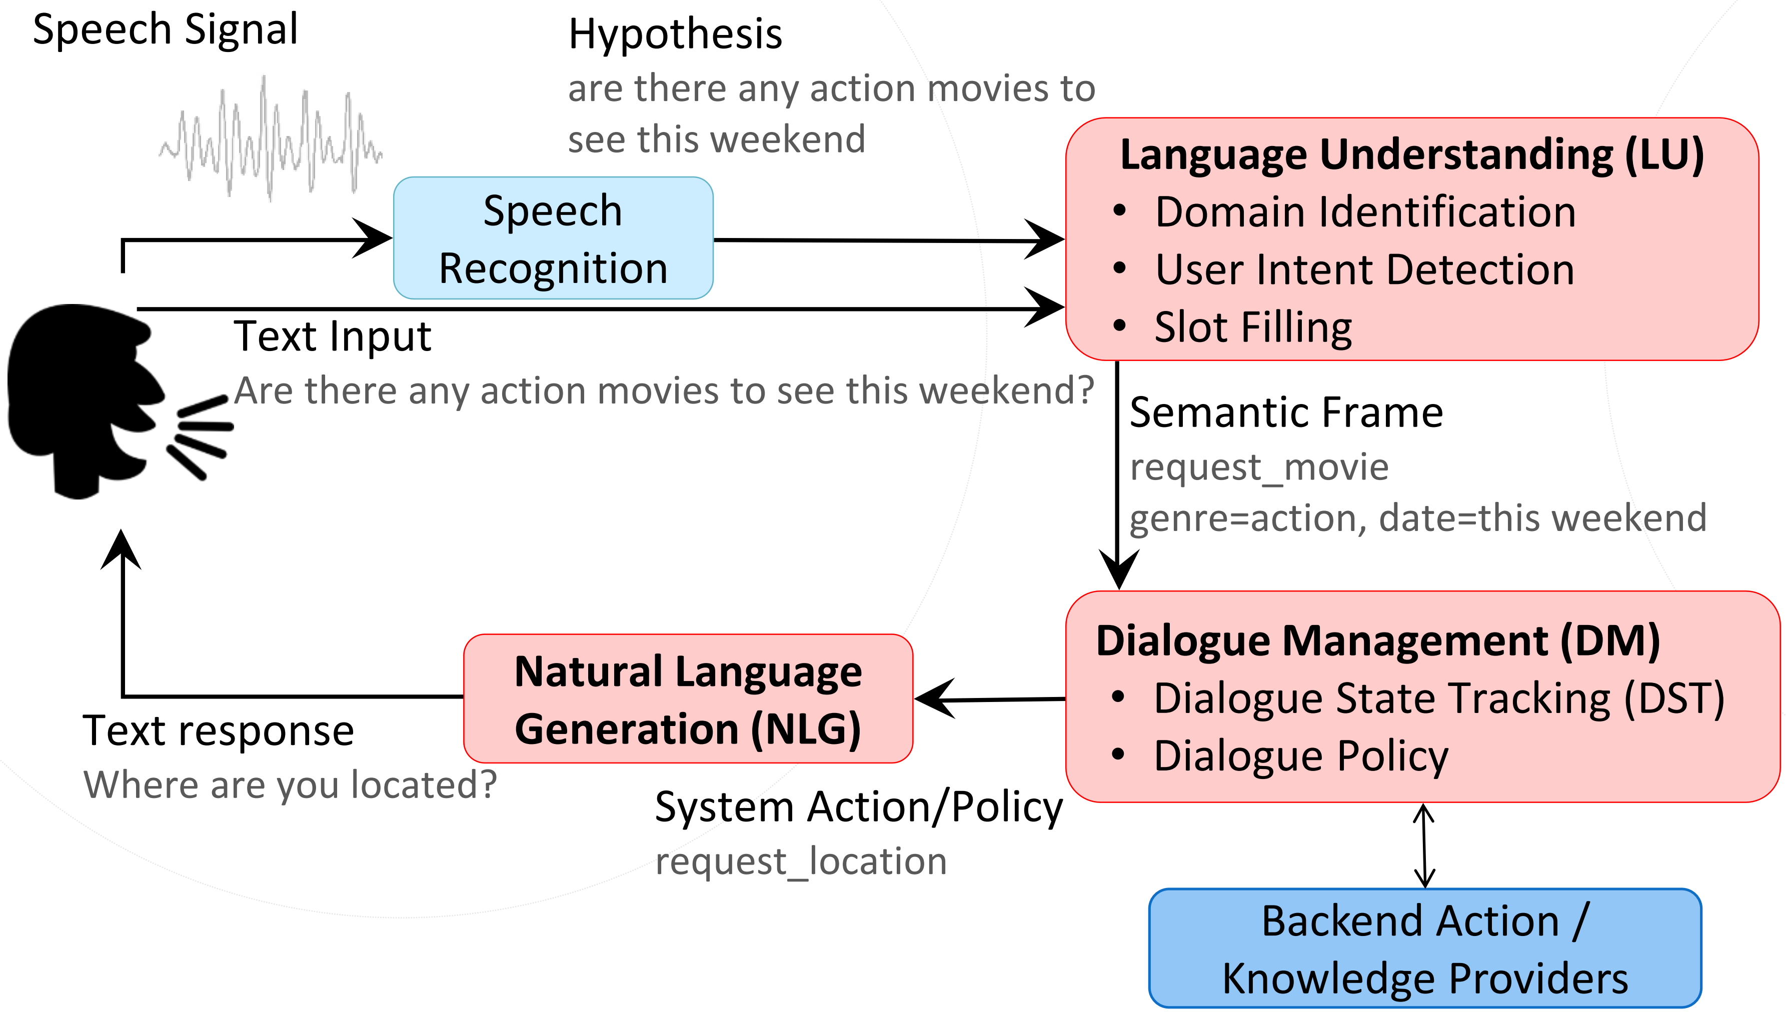

Material: http://deepdialogue.miulab.tw

Interaction Example

User

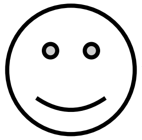

find a good eating place for taiwanese food

Intelligent Agent

Good Taiwanese eating places include Din Tai Fung, Boiling Point, etc. What do you want to choose? I can help you go there.

Q: How does a dialogue system process this request?

Material: http://deepdialogue.miulab.tw

Task-Oriented Dialogue System (Young, 2000)

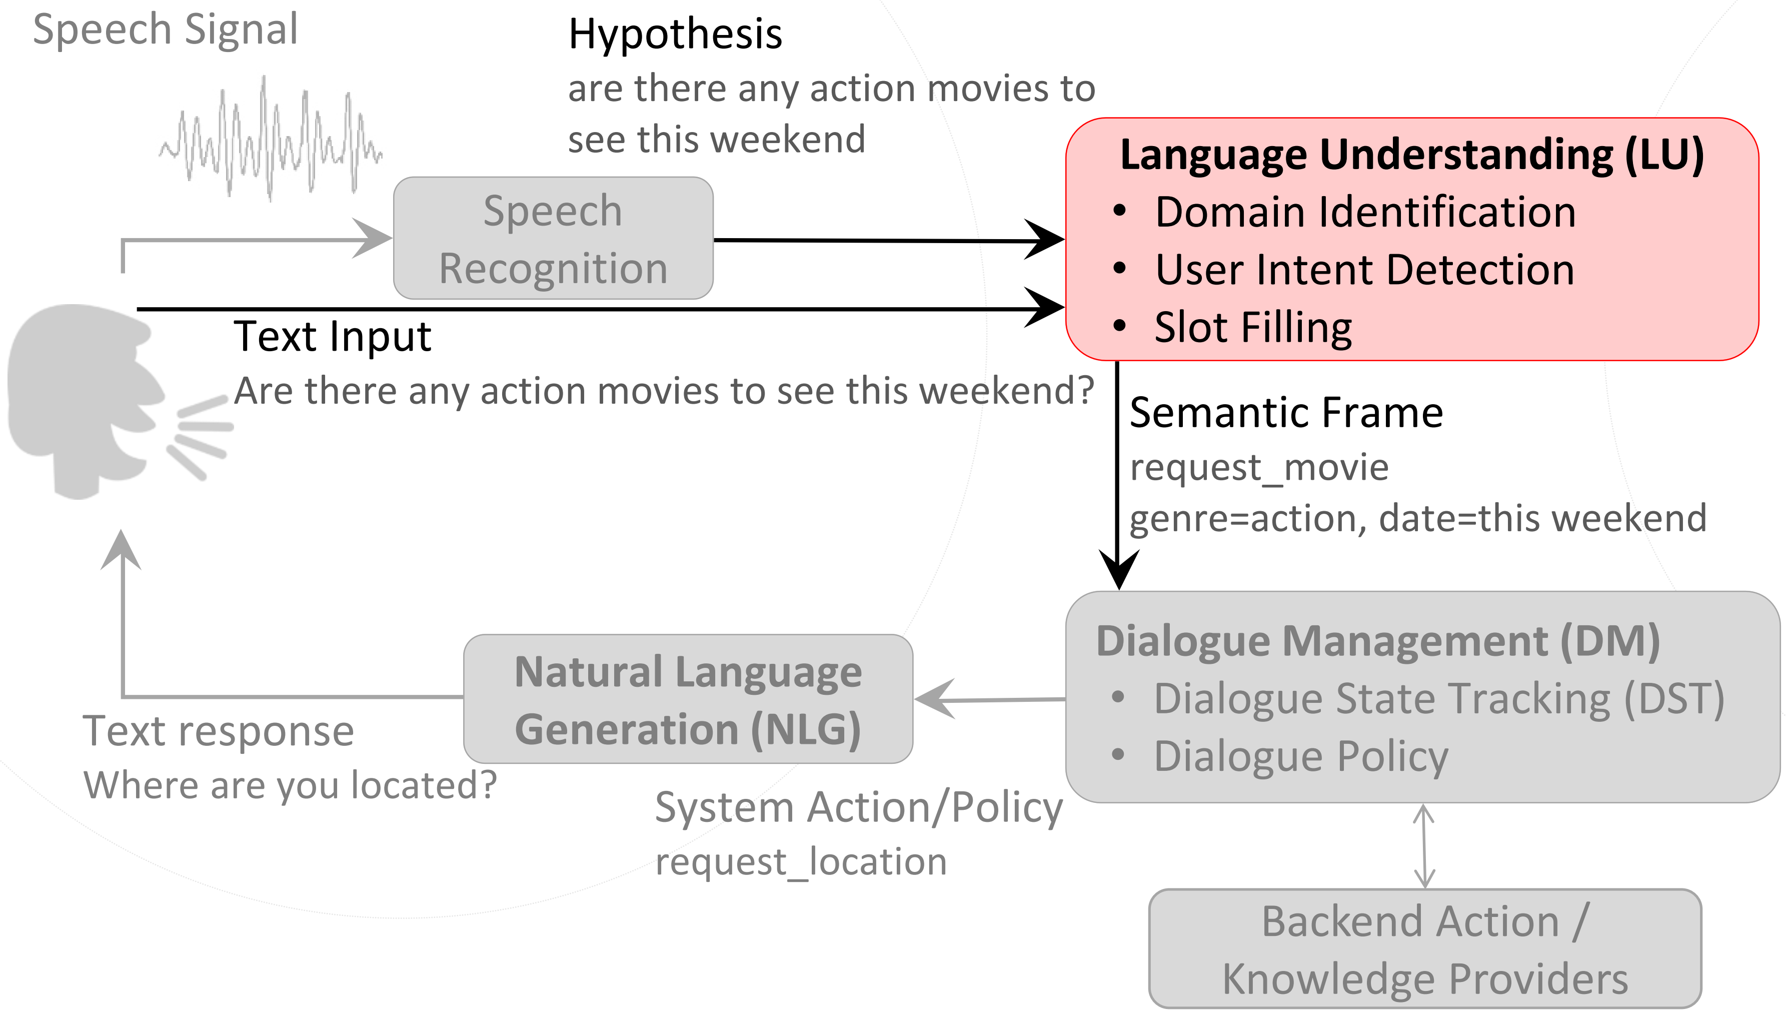

Material: http://deepdialogue.miulab.tw

1. Domain Identification

Requires Predefined Domain Ontology

User

find a good eating place for taiwanese food

Intelligent Agent

Restaurant DB

Taxi DB

Movie DB

Organized Domain Knowledge (Database)

Classification!

2. Intent Detection

Requires Predefined Schema

User

find a good eating place for taiwanese food

Intelligent Agent

Restaurant DB

FIND_RESTAURANT
FIND_PRICE
FIND_TYPE
:

Classification!

3. Slot Filling

Requires Predefined Schema

O O B-rating O O O B-type O

User

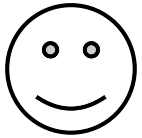

find a good eating place for taiwanese food

Restaurant	Rating	Type
Rest 1	good	Taiwanese
Rest 2	bad	Thai
:	:	:

Intelligent Agent

Restaurant DB

FIND_RESTAURANT
rating="good"
type="taiwanese"

Semantic Frame

SELECT restaurant {
rest.rating="good"
rest.type="taiwanese"
}

Sequence Labeling

Material: http://deepdialogue.miulab.tw

Task-Oriented Dialogue System (Young, 2000)

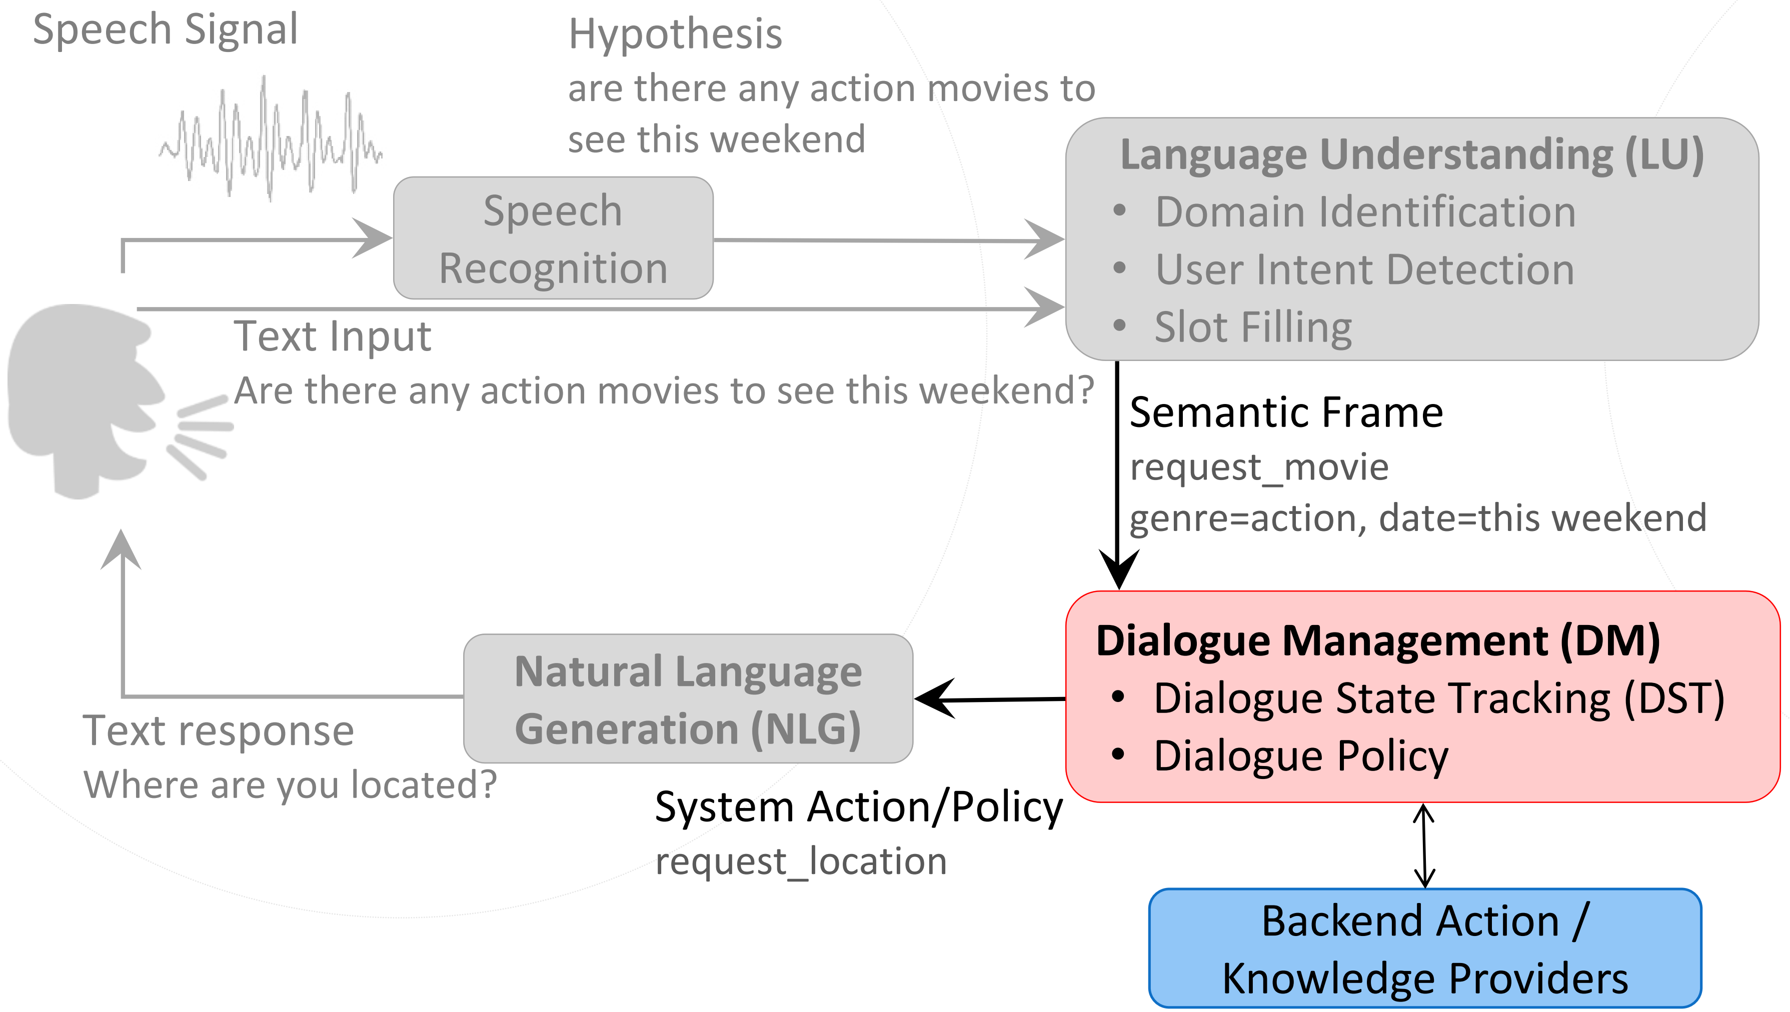

Material: http://deepdialogue.miulab.tw

State Tracking

Requires Hand-Crafted States

User

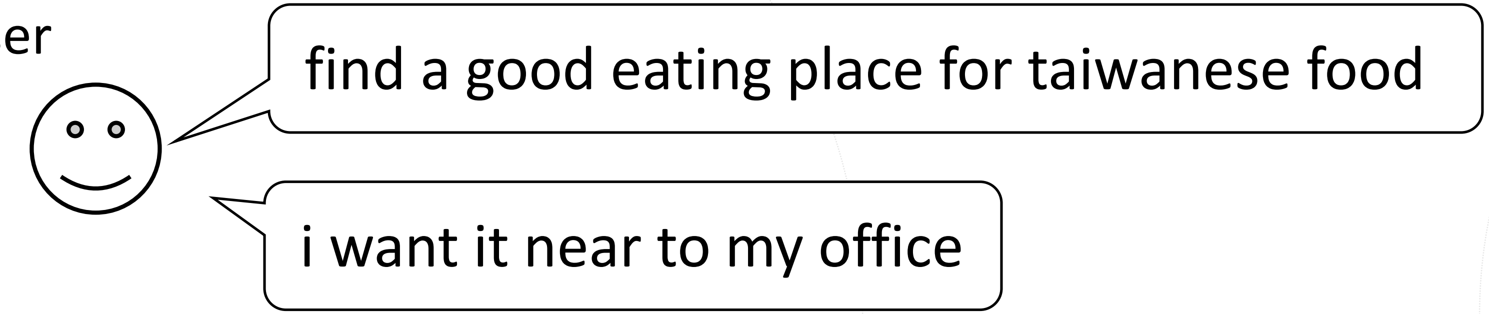

Intelligent Agent

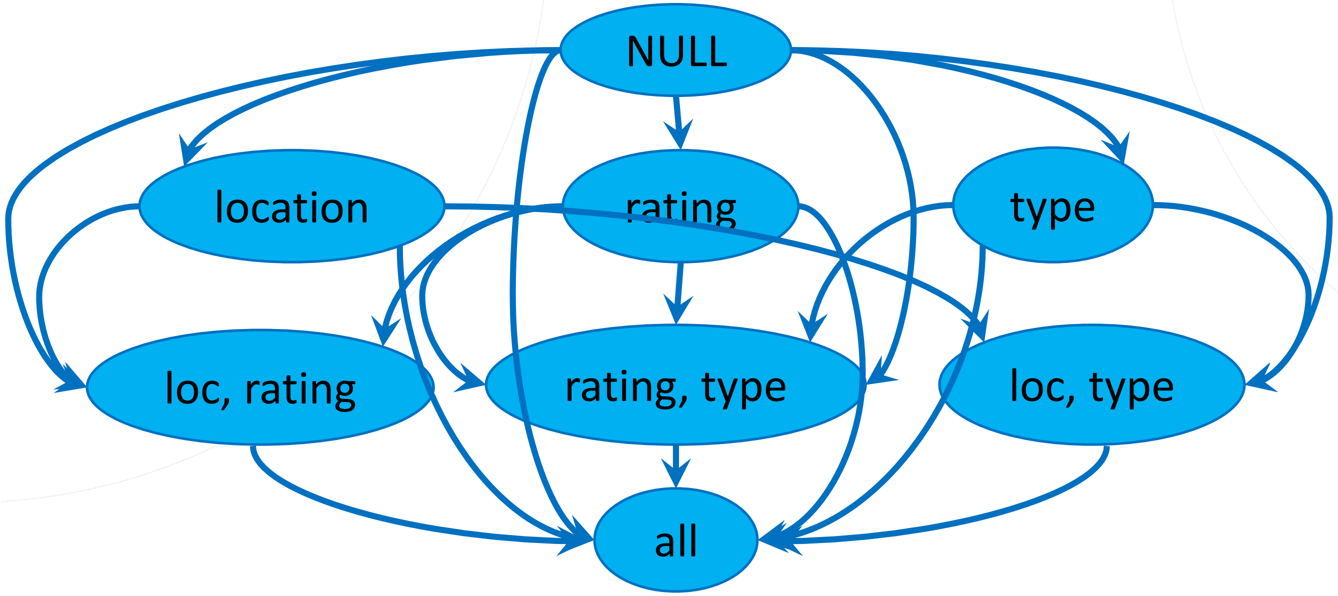

Material: http://deepdialogue.miulab.tw

State Tracking

Requires Hand-Crafted States

User

find a good eating place for taiwanese food

i want it near to my office

Intelligent Agent

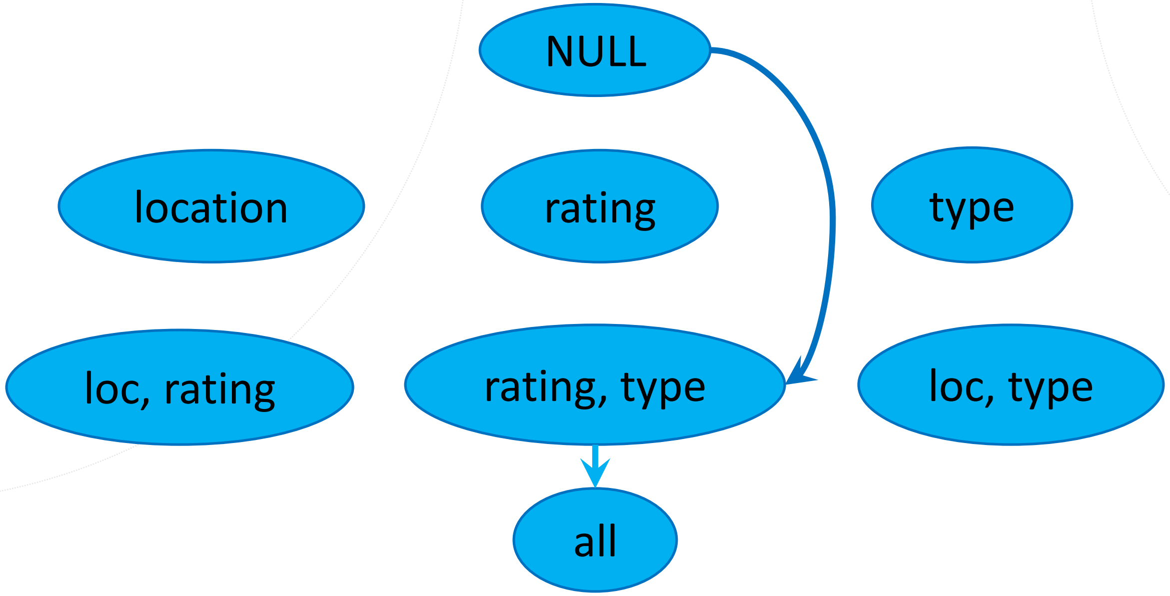

State Tracking

Handling Errors and Confidence

User

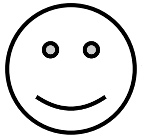

find a good eating place for taixxxx food

FIND_RESTAURANT
rating="good"
type="taiwanese"

FIND_RESTAURANT
rating="good"
type="thai"

FIND_RESTAURANT
rating="good"

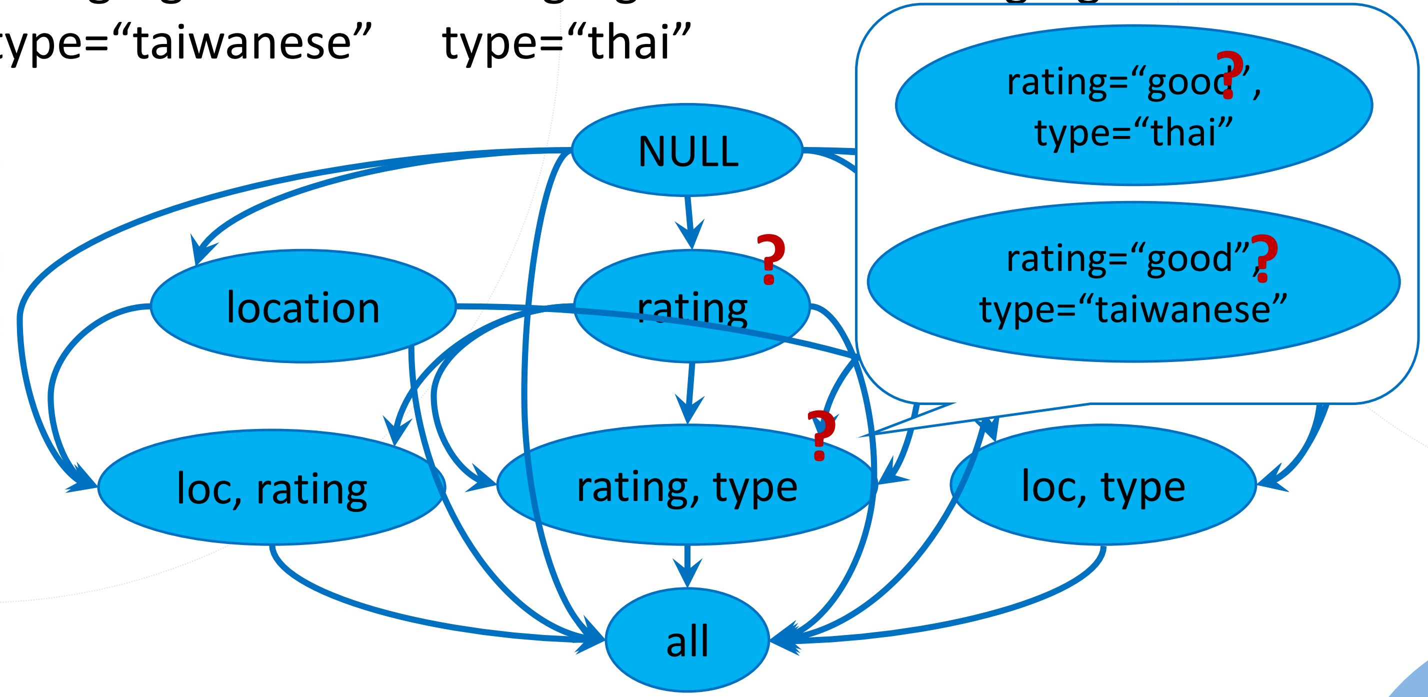

Intelligent
Agent

Dialogue Policy for Agent Action

- Inform(location="Taipei 101")
  - "The nearest one is at Taipei 101"
- Request(location)
  - "Where is your home?"
- Confirm(type="taiwanese")
  - "Did you want Taiwanese food?"

Material: http://deepdialogue.miulab.tw

Task-Oriented Dialogue System (Young, 2000)

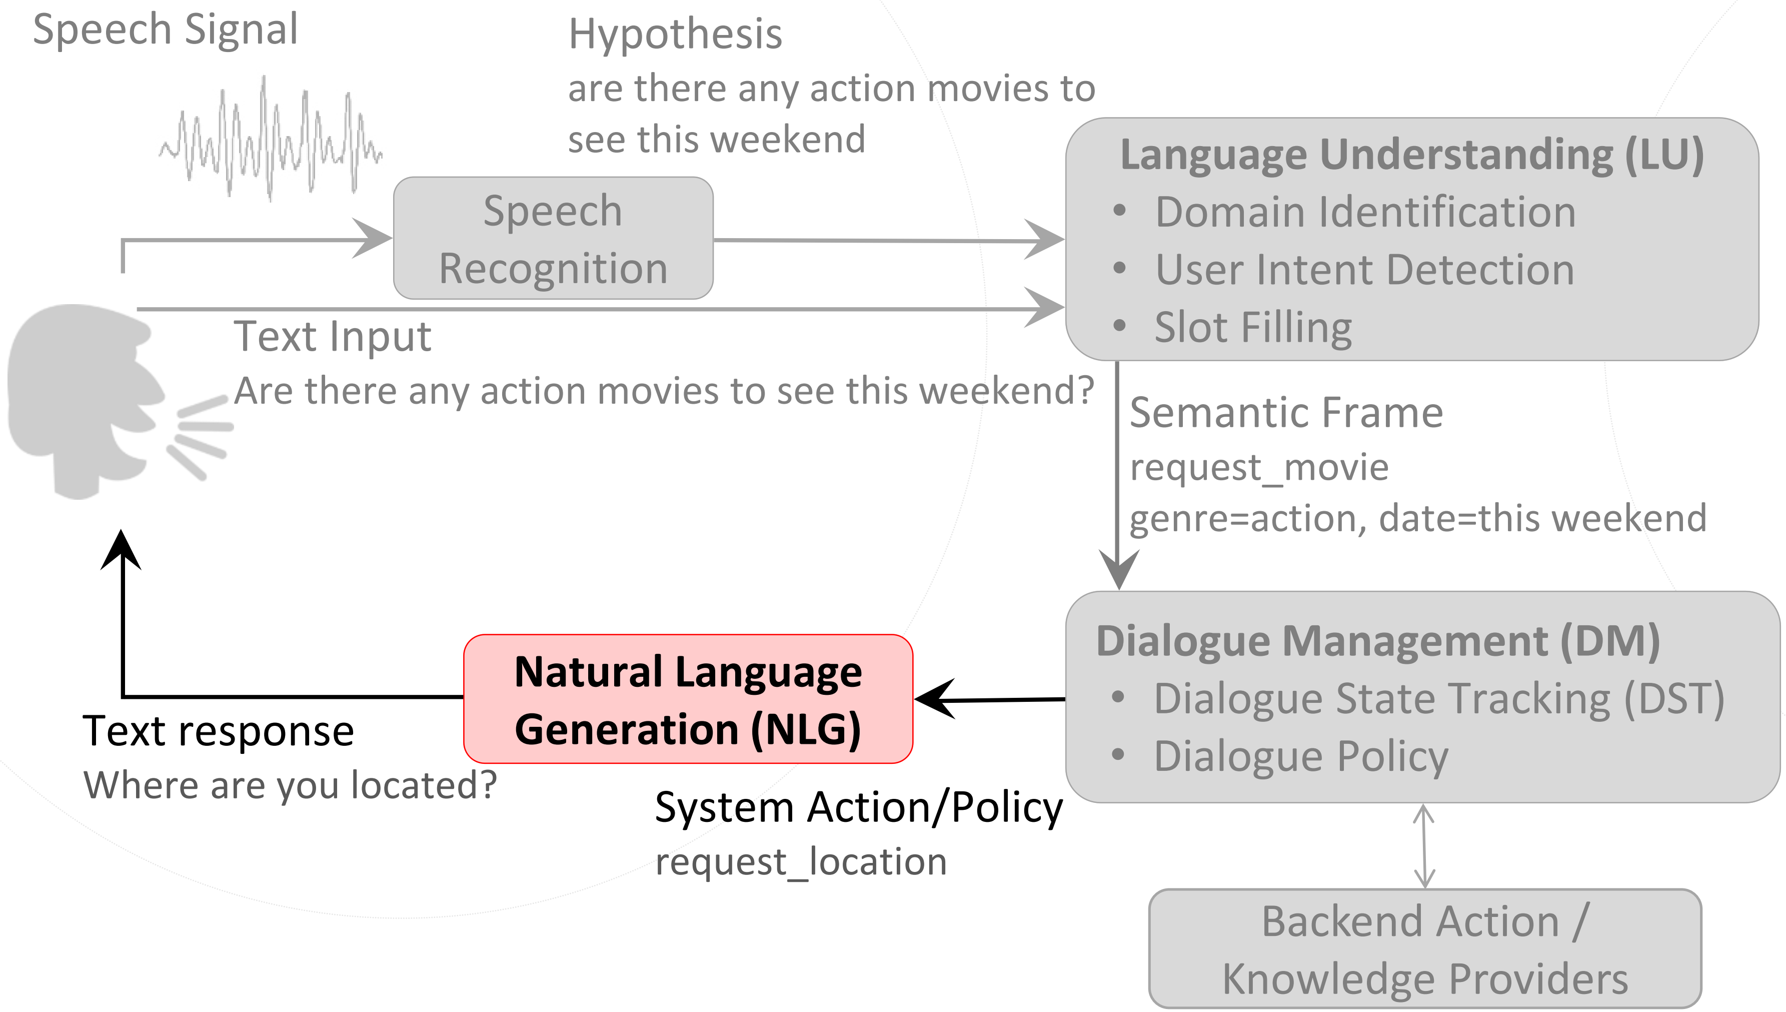

Output / Natural Language Generation

- Goal: generate natural language or GUI given the selected dialogue action for interactions

- Inform(location="Taipei 101")
  - "The nearest one is at Taipei 101" v.s.

- Request(location)
  - "Where is your home?" v.s.

- Confirm(type="taiwanese")
  - "Did you want Taiwanese food?" v.s.

Outline

- Introduction
- **Background Knowledge**
  - Neural Network Basics
  - Reinforcement Learning
- Modular Dialogue System
- System Evaluation
- Recent Trends of Learning Dialogues

Outline

- Introduction
- **Background Knowledge**
  - **Neural Network Basics**
  - Reinforcement Learning
- Modular Dialogue System
- System Evaluation
- Recent Trends of Learning Dialogues

Machine Learning ≈ Looking for a Function

- Speech Recognition

$f($ 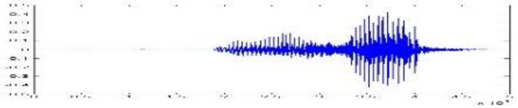 $)=$ "你好 (Hello) "

- Image Recognition

$f($  $)=$ cat

- Go Playing

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

- Chat Bot

$f($ "Where is KAIST?" $)=$ "The address is..."

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

Machine Learning

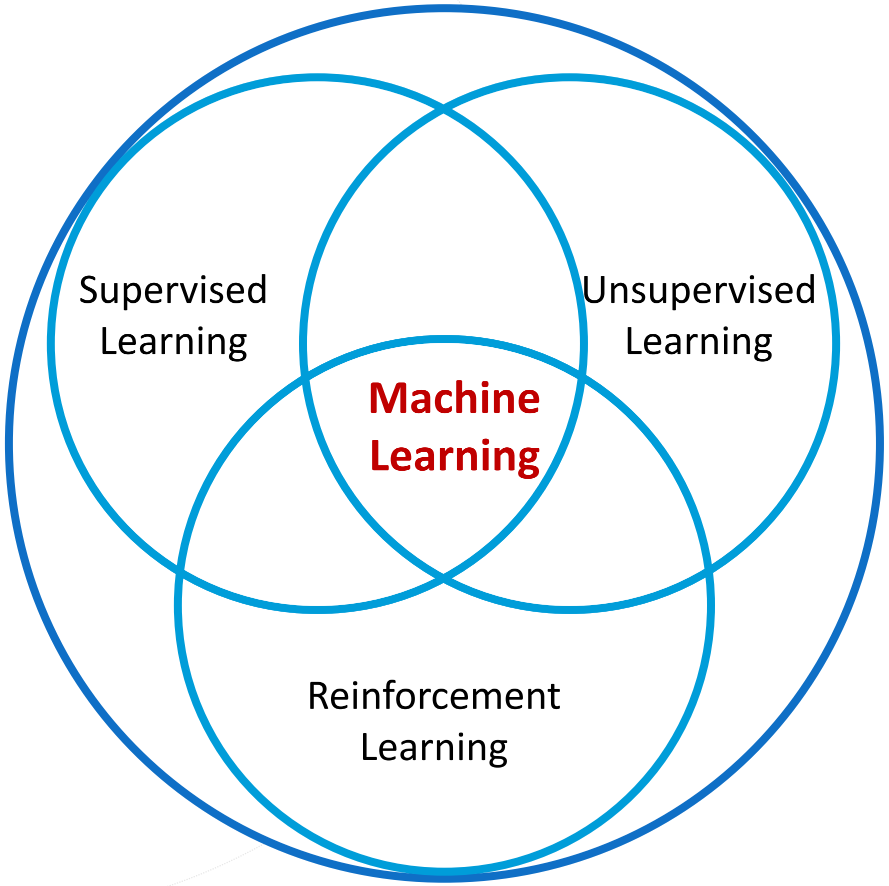

Deep learning is a type of machine learning approaches, called “neural networks”.

A Single Neuron

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

Activation function

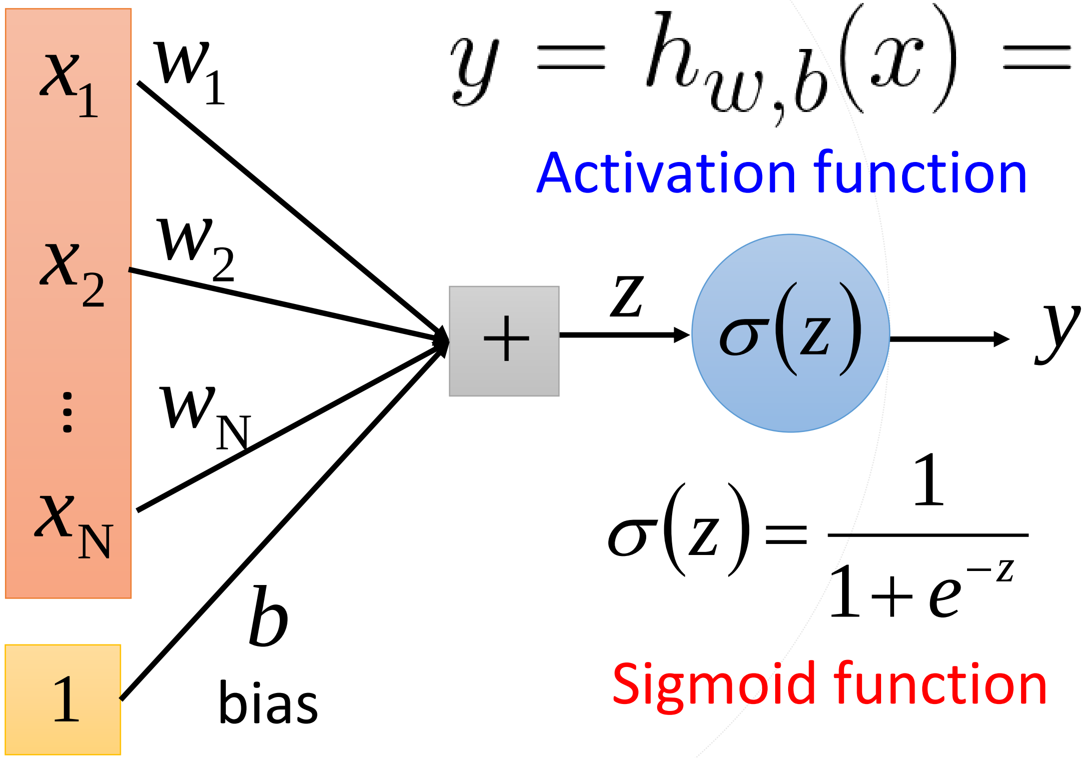

$$\sigma(z) = \frac{1}{1+e^{-z}}$$

Sigmoid function

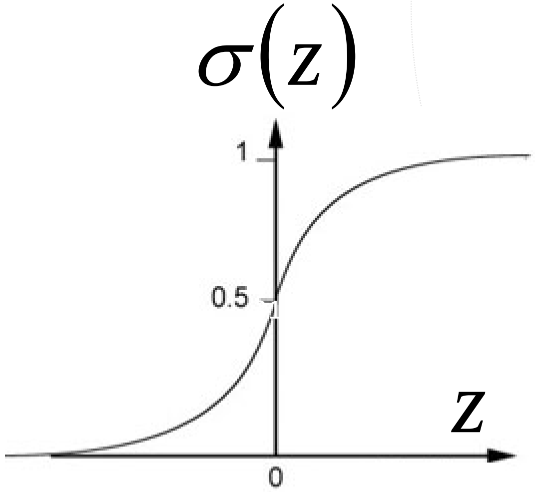

w, b are the parameters of this neuron

Material: http://deepdialogue.miulab.tw

A Single Neuron

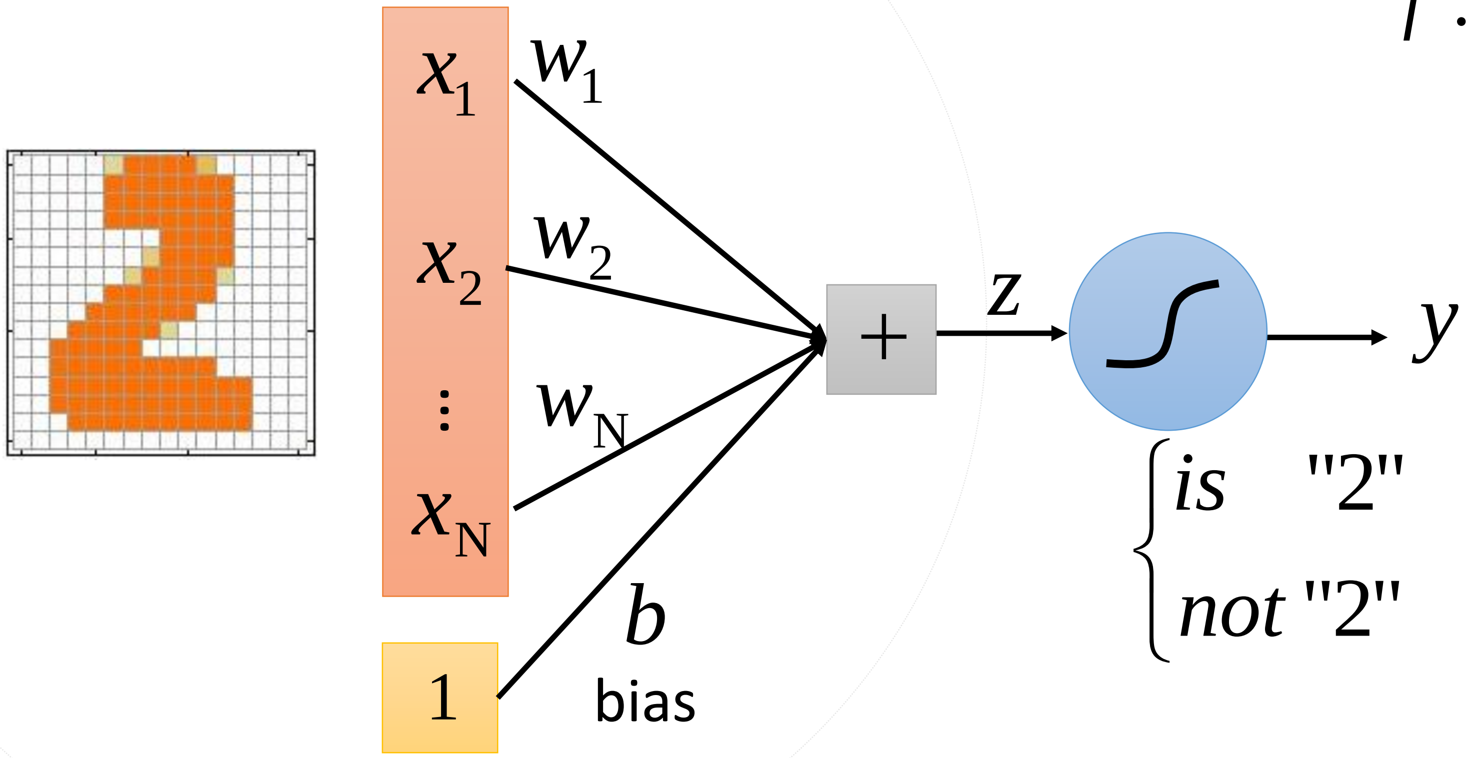

$$f : R^N \rightarrow R^M$$

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

A single neuron can only handle binary classification

A Layer of Neurons

- Handwriting digit classification

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

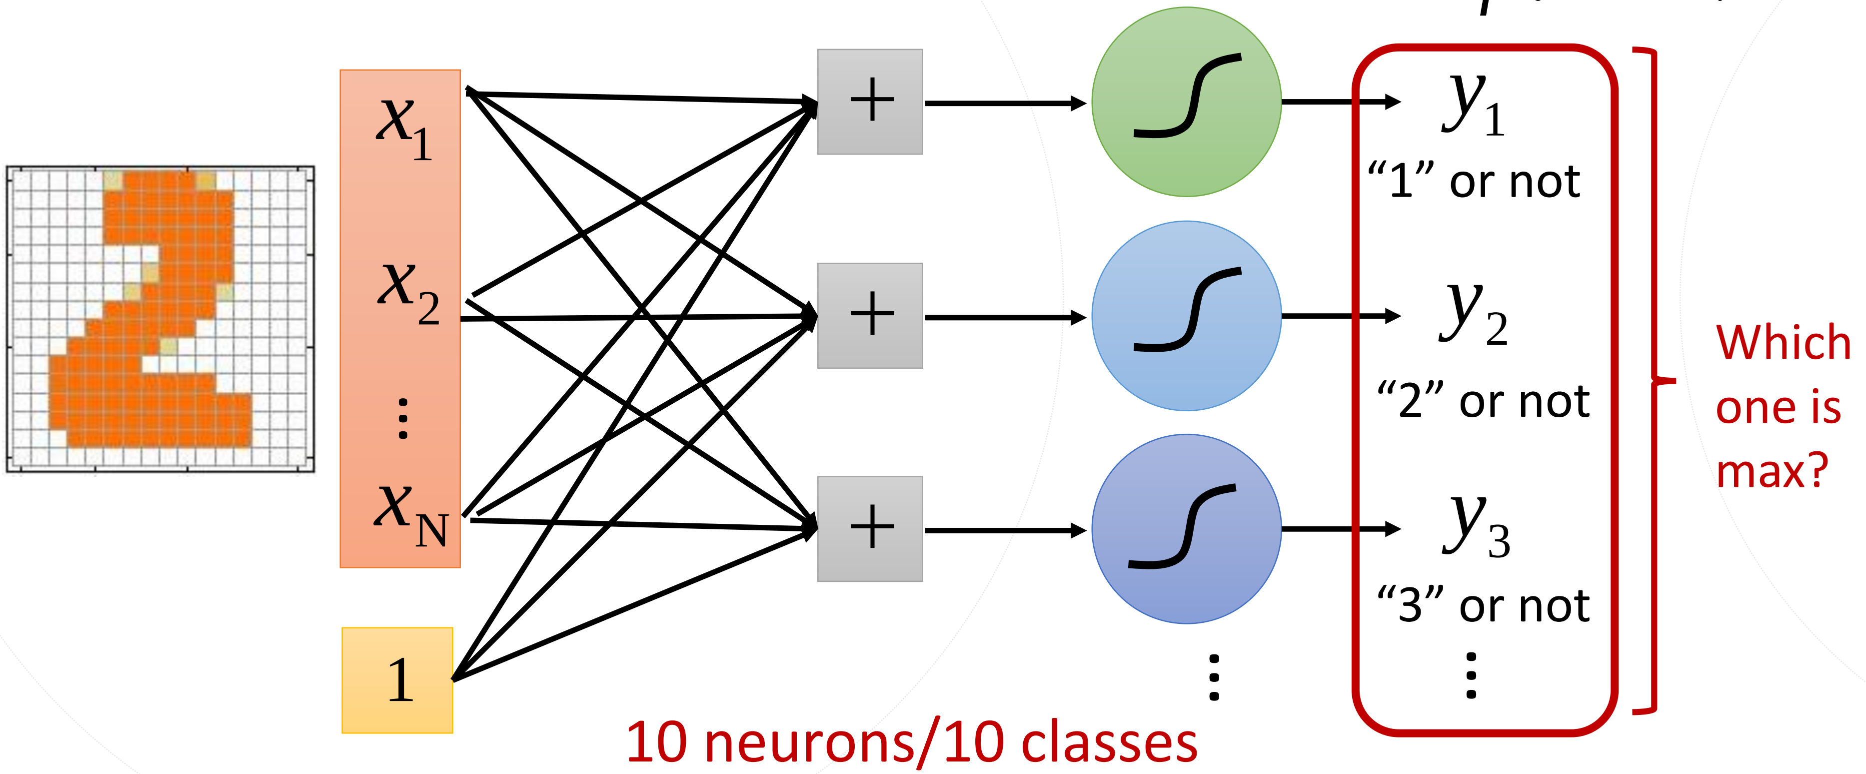

A layer of neurons can handle multiple possible output, and the result depends on the max one

Deep Neural Networks (DNN)

- Fully connected feedforward network

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

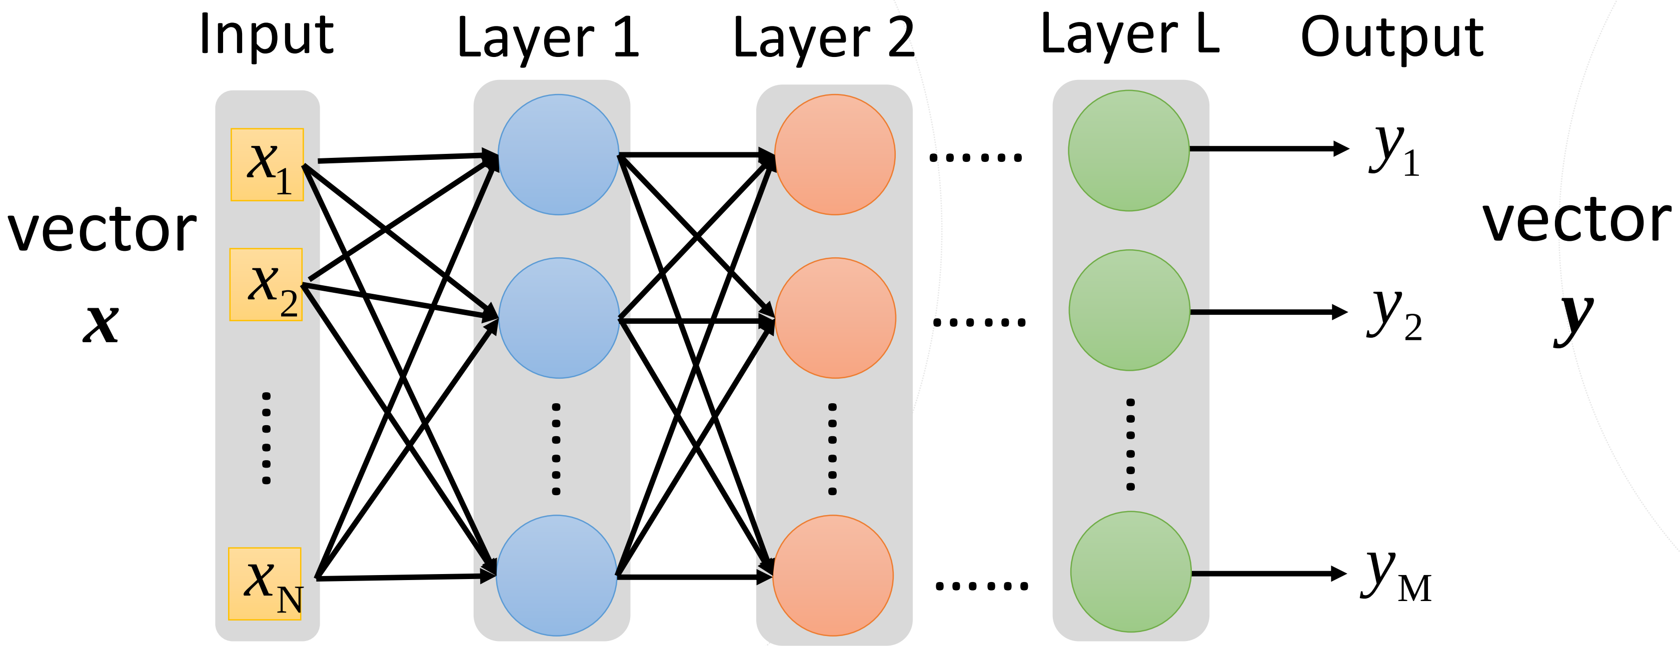

Deep NN: multiple hidden layers

Recurrent Neural Network (RNN)

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

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

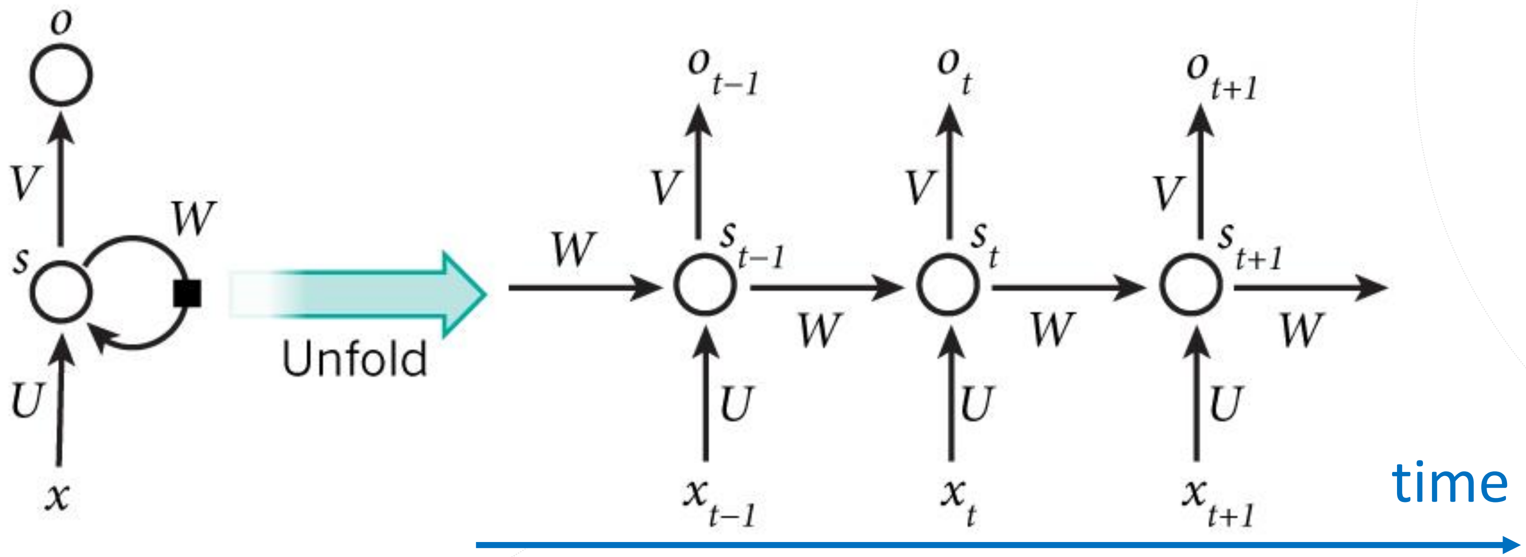

RNN can learn accumulated sequential information (time-series)

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

Outline

- Introduction
- **Background Knowledge**
  - Neural Network Basics
  - **Reinforcement Learning**
- Modular Dialogue System
- System Evaluation
- Recent Trends of Learning Dialogues

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*

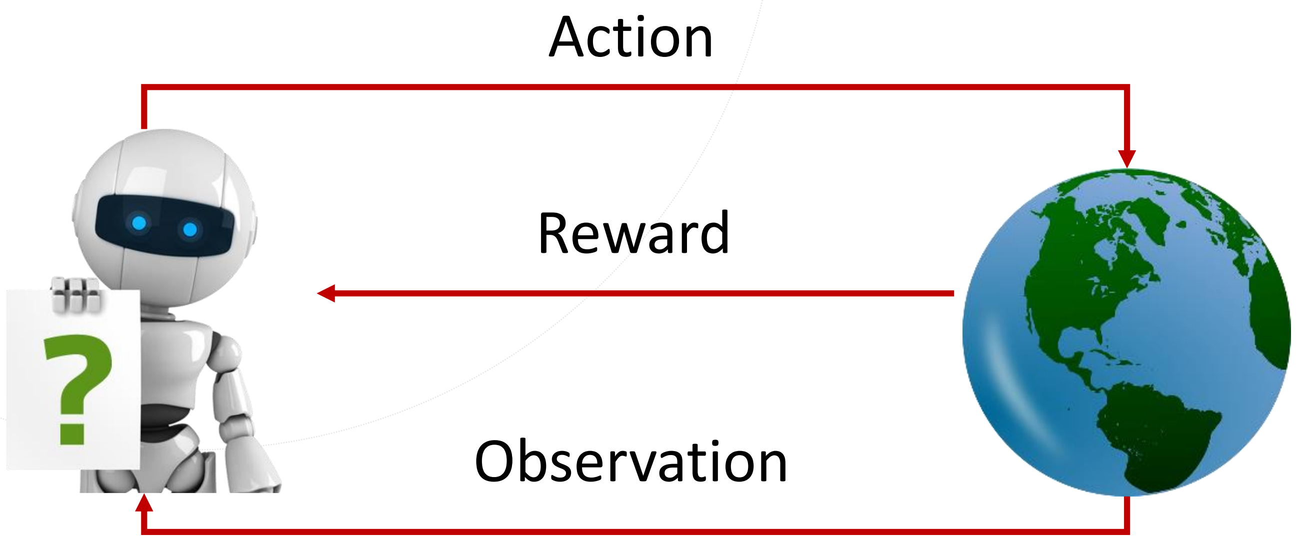

Scenario of Reinforcement Learning

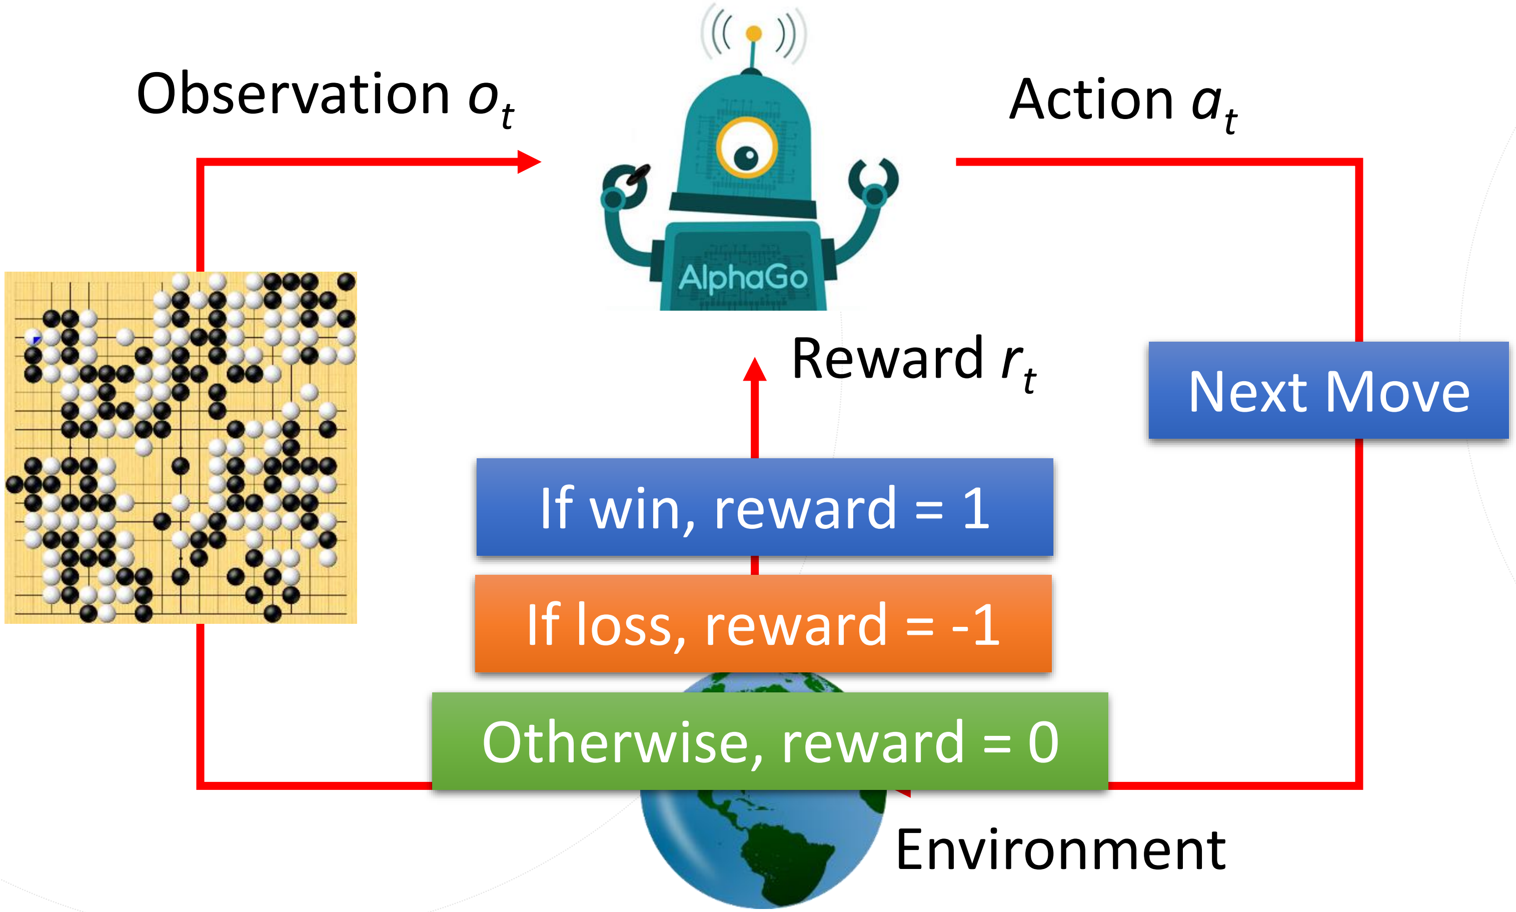

Agent learns to take actions to maximize expected reward.

Supervised v.s. Reinforcement

- Supervised

Learning from teacher

"Hello"

"Bye bye"

Say "Hi"

Say "Good bye"

- Reinforcement

…….

…….

…..

Bad

Hello ☺

……

Learning from critics

Sequential Decision Making

- 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

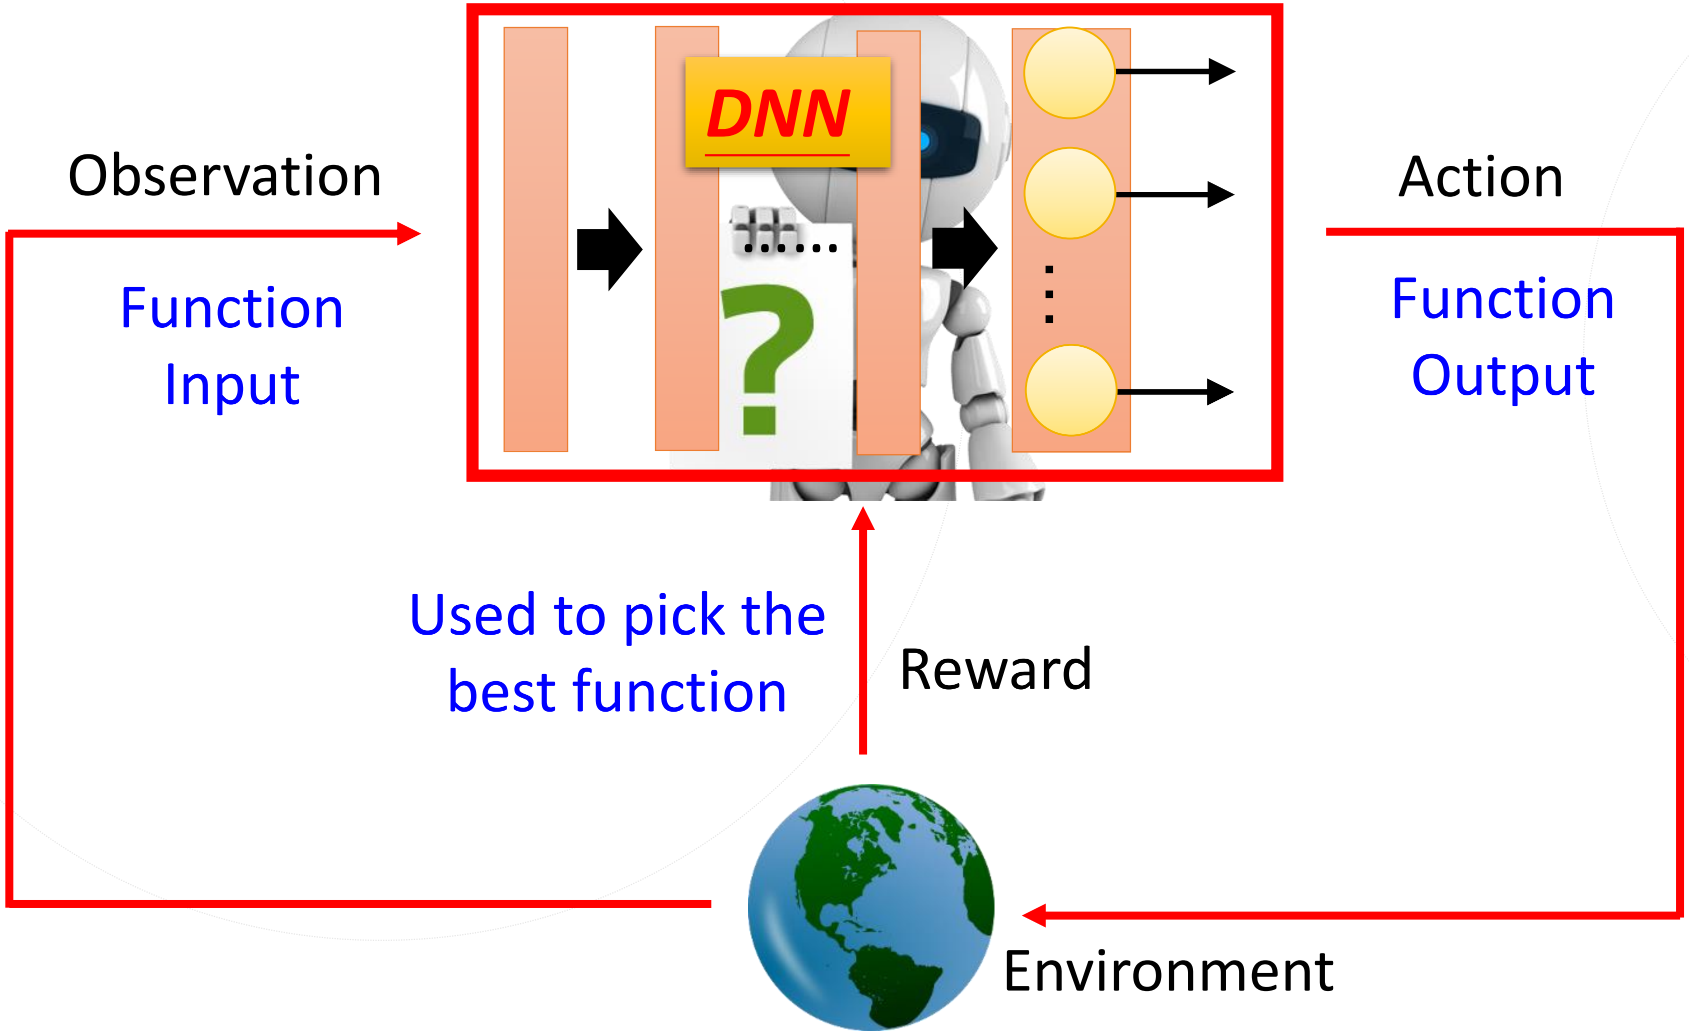

Reinforcing Learning

- 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

Outline

- Introduction
- Background Knowledge
- **Modular Dialogue System**
  - Spoken/Natural Language Understanding (SLU/NLU)
  - Dialogue Management
    - Dialogue State Tracking (DST)
    - Dialogue Policy Optimization
  - Natural Language Generation (NLG)
- System Evaluation
- Recent Trends of Learning Dialogues

Task-Oriented Dialogue System (Young, 2000)

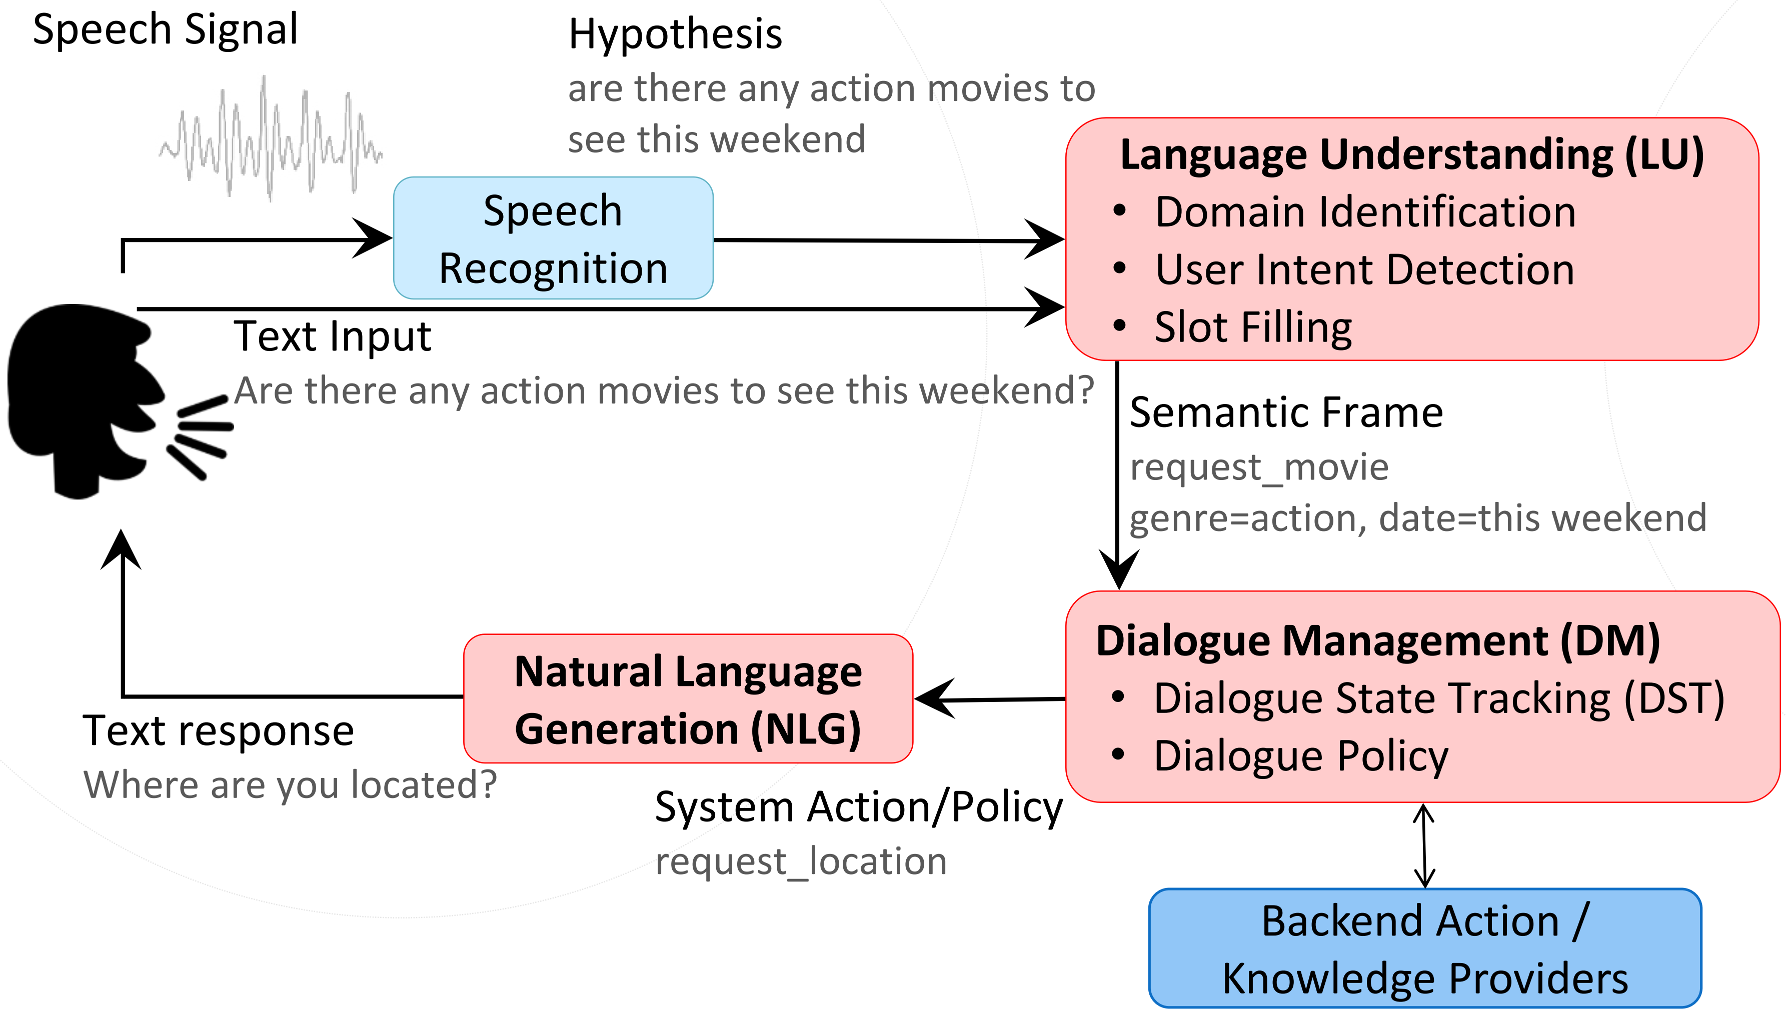

Outline

- Introduction
- Background Knowledge
- **Modular Dialogue System**
    - **Spoken/Natural Language Understanding (SLU/NLU)**
    - Dialogue Management
        - Dialogue State Tracking (DST)
        - Dialogue Policy Optimization
    - Natural Language Generation (NLG)
- System Evaluation
- Recent Trends of Learning Dialogues

Material: http://deepdialogue.miulab.tw

Language Understanding (LU)

- Pipelined

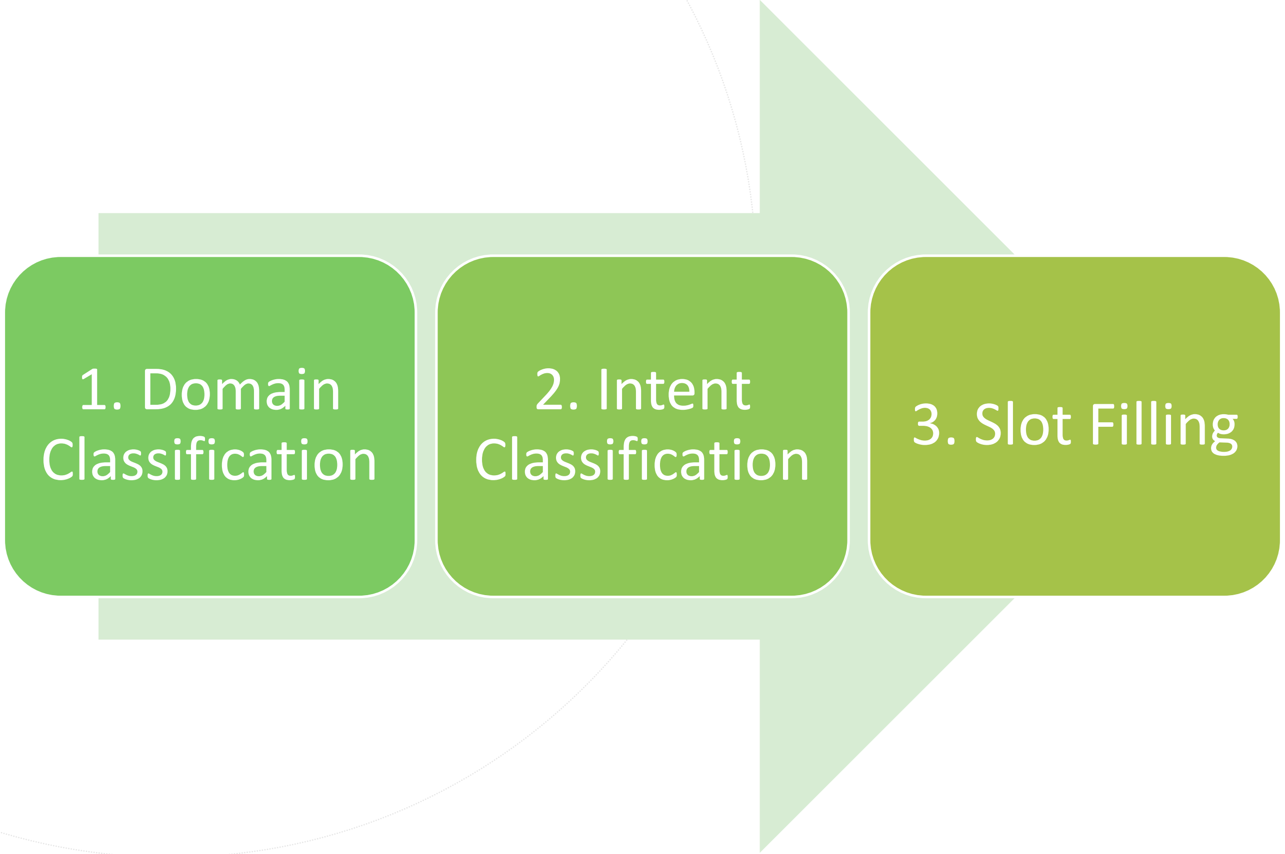

LU – Domain/Intent Classification

As an **utterance classification** task

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

find me a cheap taiwanese restaurant in oakland

Domain	Intent
Movies	find_movie, buy_tickets
Restaurants	find_restaurant, find_price, book_table
Music	find_lyrics, find_singer
Sports	…
…	

Domain/Intent Classification (Sarikaya et al., 2011)

- Deep belief nets (DBN)
  - Unsupervised training of weights
  - Fine-tuning by back-propagation
  - Compared to MaxEnt, SVM, and boosting

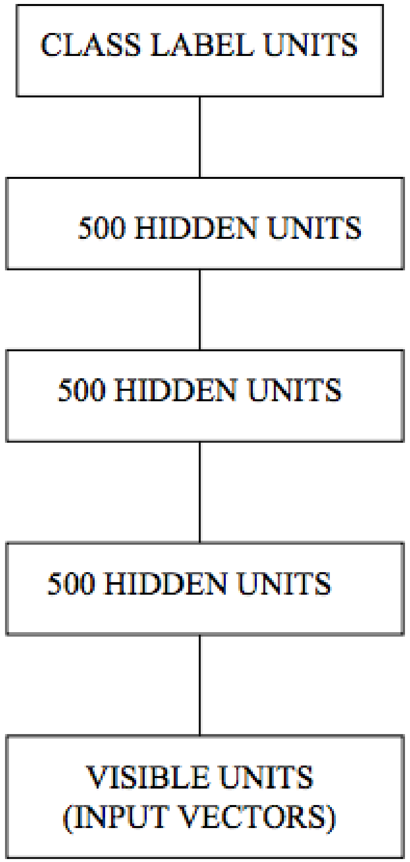

Domain/Intent Classification (Tur et al., 2012; Deng et al., 2012)

- Deep convex networks (DCN)
  - Simple classifiers are stacked to learn complex functions
  - Feature selection of salient n-grams
- Extension to kernel-DCN

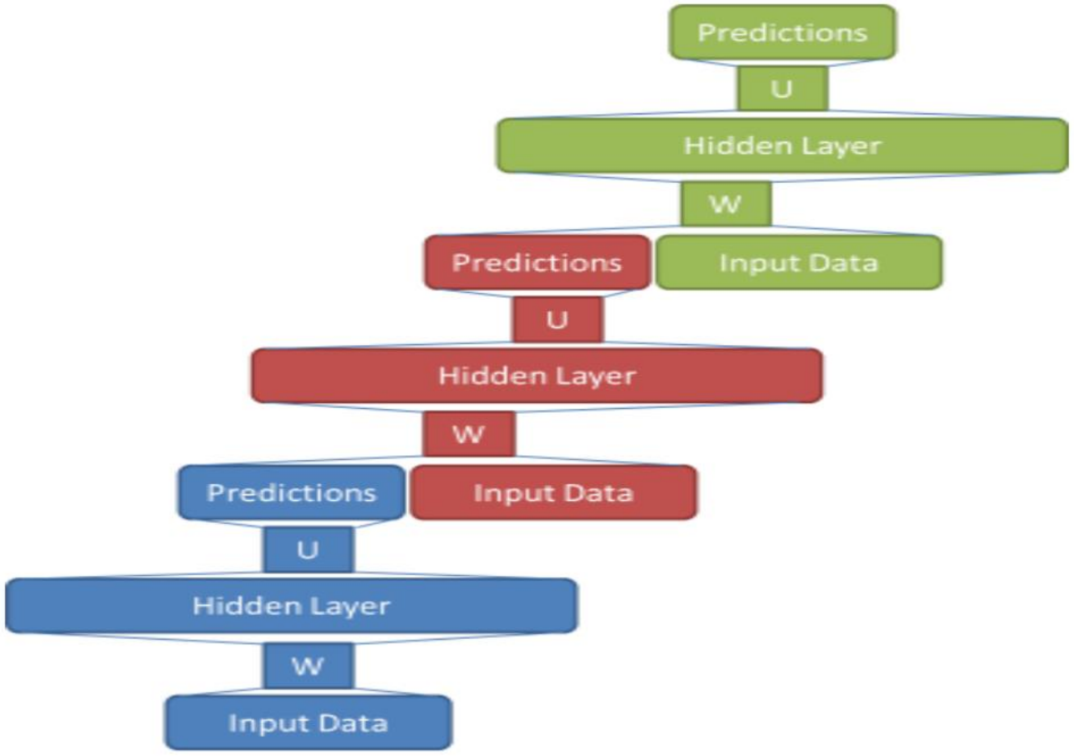

Domain/Intent Classification (Ravuri & Stolcke, 2015)

- RNN and LSTMs for utterance classification

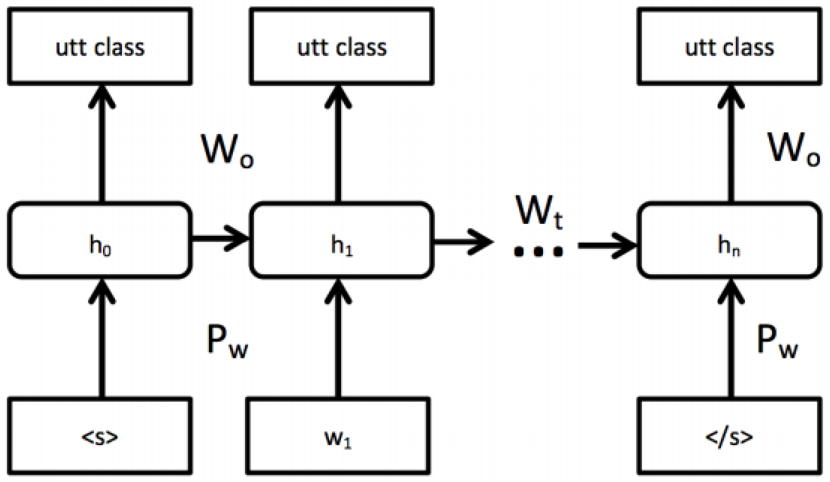

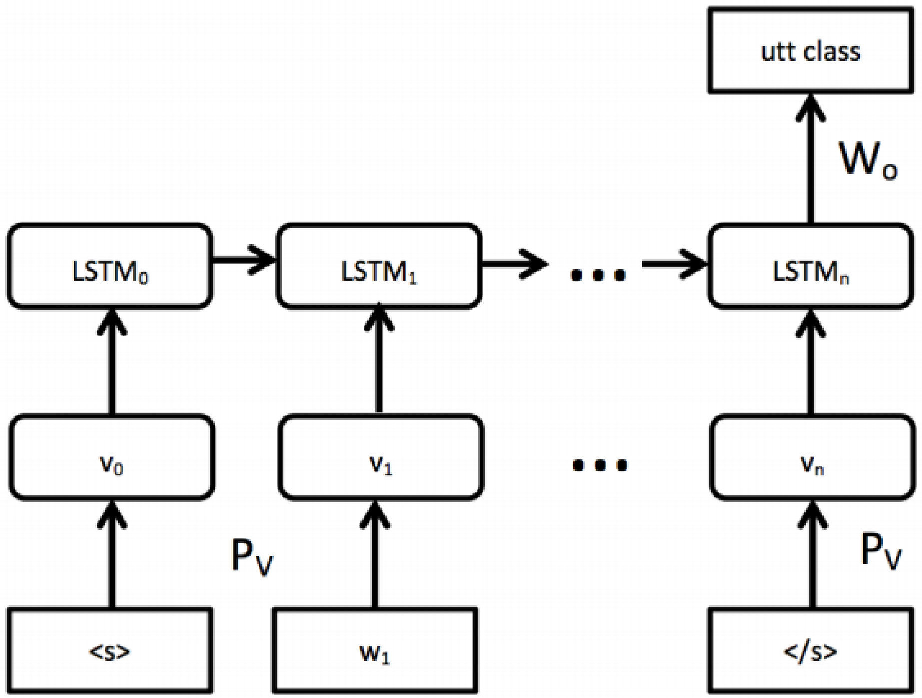

Intent decision after reading all words performs better

Dialogue Act Classification (Lee & Dernoncourt, 2016)

- RNN and CNNs for dialogue act classification

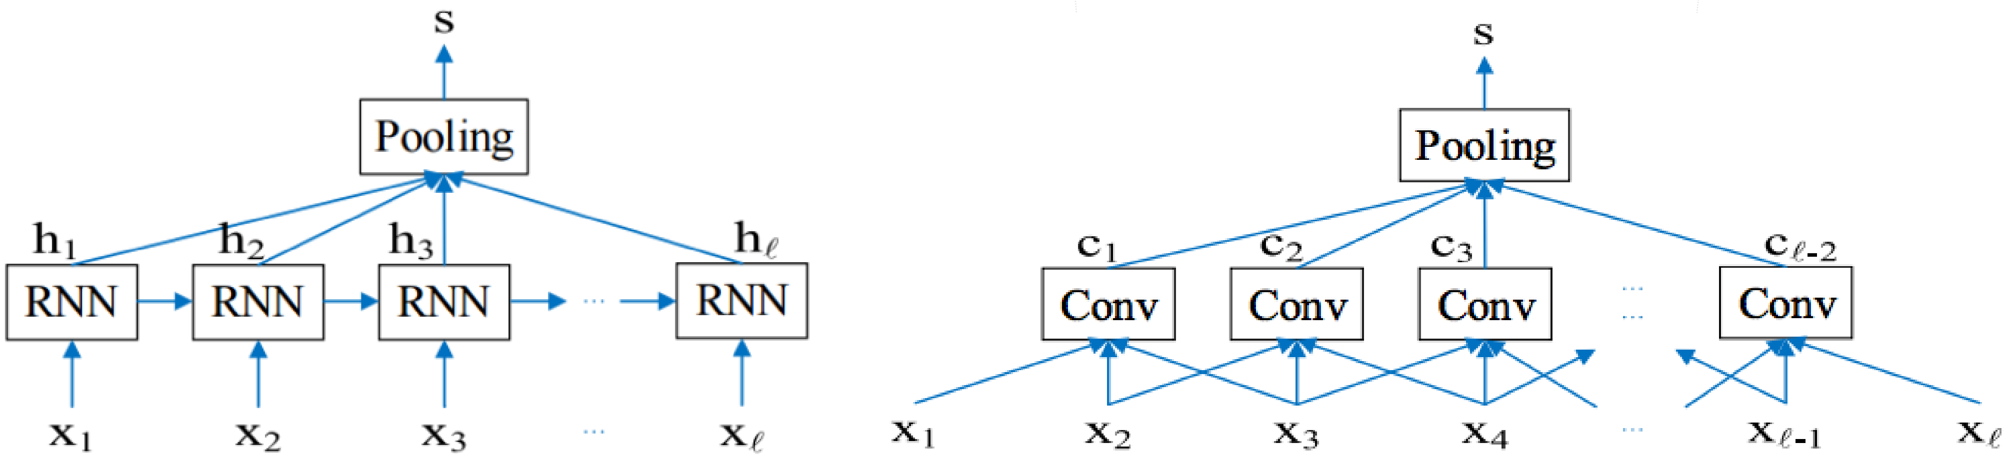

LU – Slot Filling

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

Slot Tagging (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

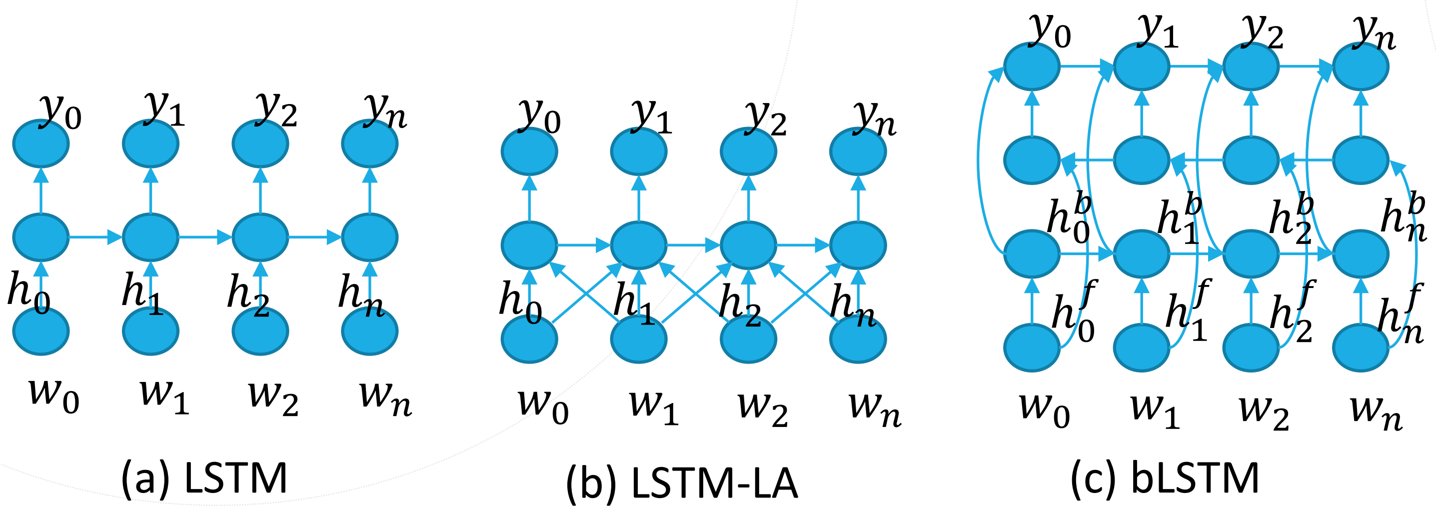

Slot Tagging (Kurata et al., 2016; Simonnet et al., 2015)

- Encoder-decoder networks
  - Leverages sentence level information

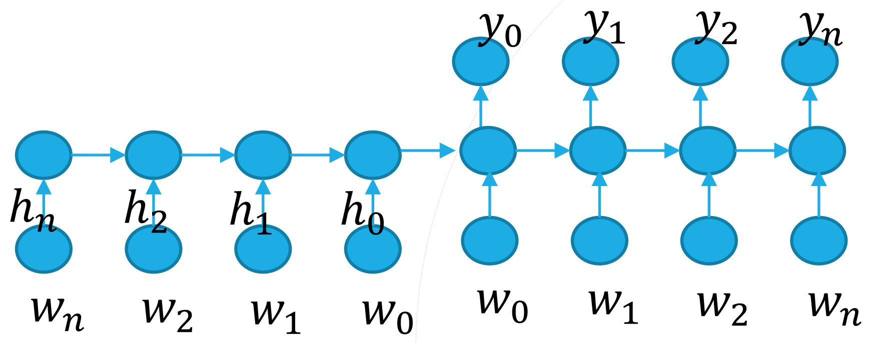

- Attention-based encoder-decoder
  - Use of attention (as in MT) in the encoder-decoder network
  - Attention is estimated using a feed-forward network with input: h_t and s_t at time t

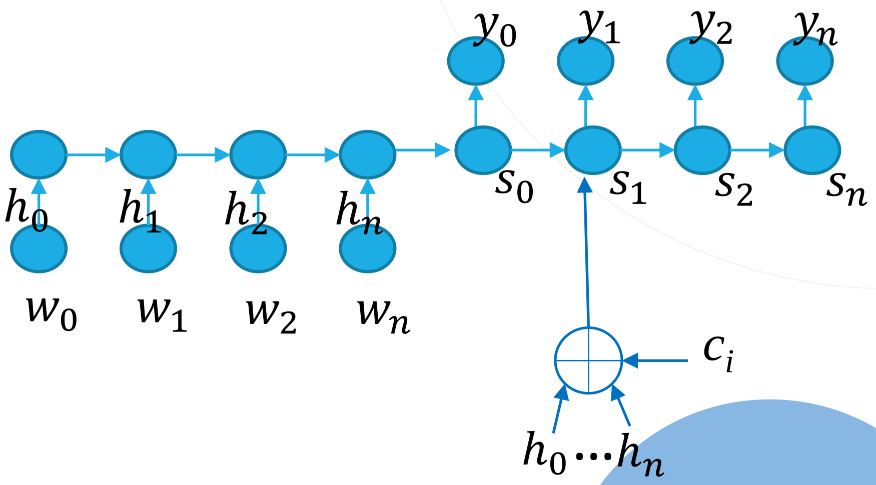

Joint Segmentation & Slot Tagging (Zhai+, 2017)

- Encoder that segments
- Decoder that tags the segments

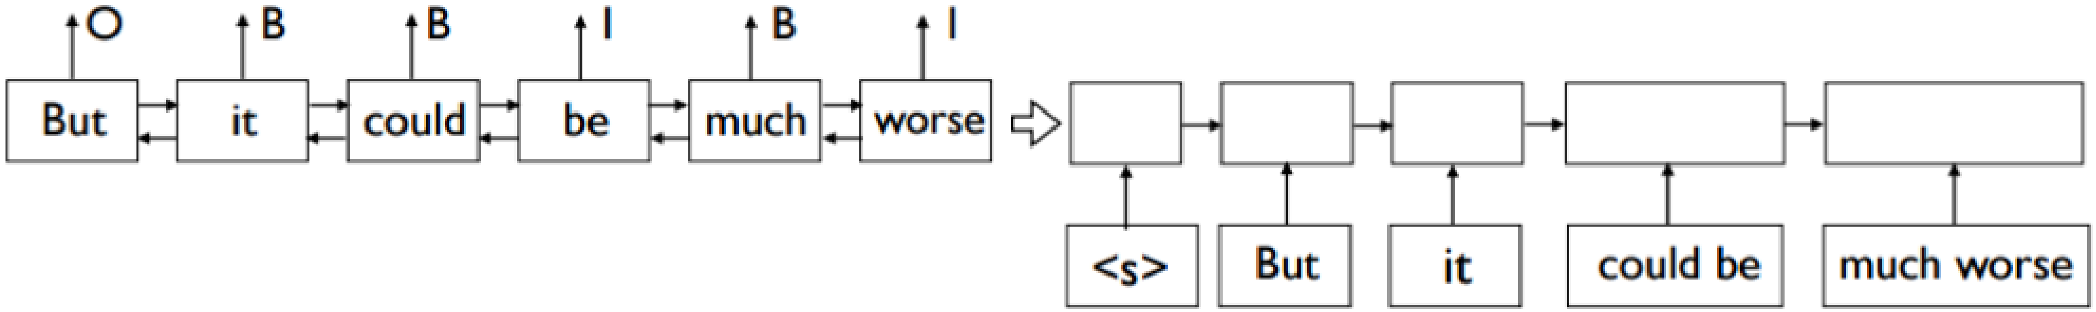

Multi-Task Slot Tagging (Jaech et al., 2016; Tafforeau et al., 2016)

- 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

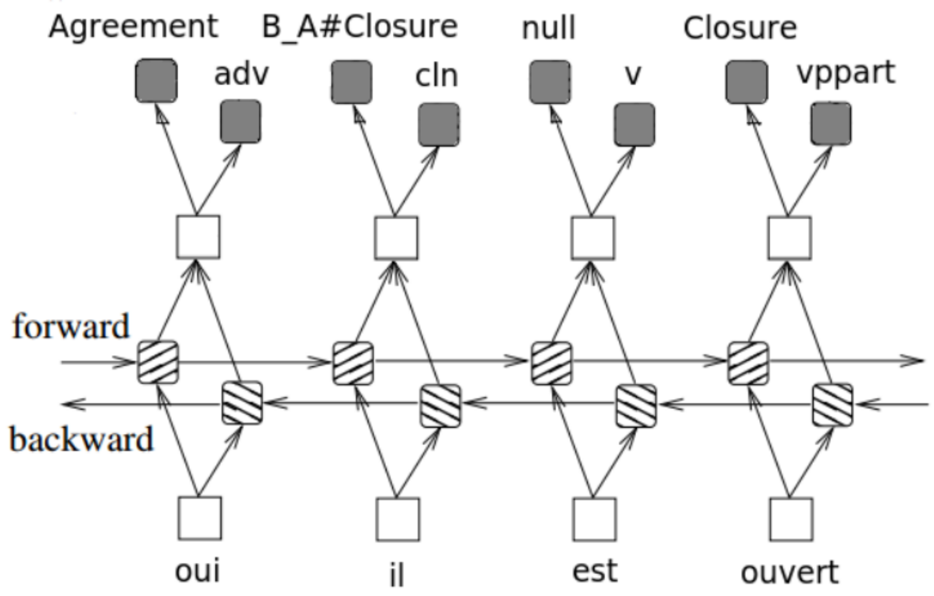

Semi-Supervised Slot Tagging (Lan+, 2018)

- Idea: language model objective can enhance other tasks

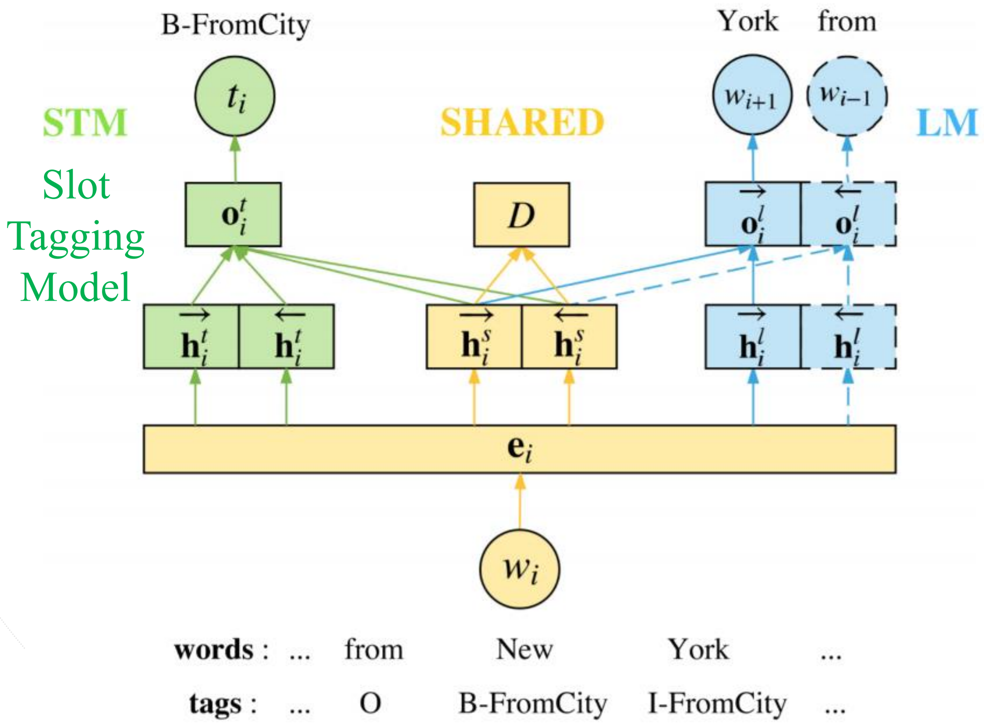

Algorithm 1: Adversarial Multi-task Learning for SLU

Input : Labeled training data $\{(\mathbf{w}^l, \mathbf{t}^l)\}$
Unlabeled data $\{\mathbf{w}^u\}$
Output: Adversarially enhanced slot tagging model

```
Initialize parameters {θ^s, θ^t, θ^l, θ^d} randomly.
repeat
    /* Sample from {(w^l, t^l)}                         */
    Train the STM and shared model by Eq.(8).
    Train the task discriminator and the shared model
      by Eq.(6) or Eq.(7) as slot tagging task (y = 1).
    /* Sample from {w^l} and {w^u}                      */
    Train the LM and shared models by Eq.(9) (and
      Eq.(10) for BLM).
    Train the task discriminator and the shared model
      by Eq.(6) or Eq.(7) as LM task (y = 0).
until convergence;
```

BLM exploits the *unsupervised knowledge*, the *shared-private framework* and *adversarial training* make the slot tagging model more generalized

LU Evaluation

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

Joint Semantic Frame Parsing

Sequence-based (Hakkani-Tur et al., 2016)

- Slot filling and intent prediction in the same output sequence

Parallel (Liu and Lane, 2016)

- Intent prediction and slot filling are performed in two branches

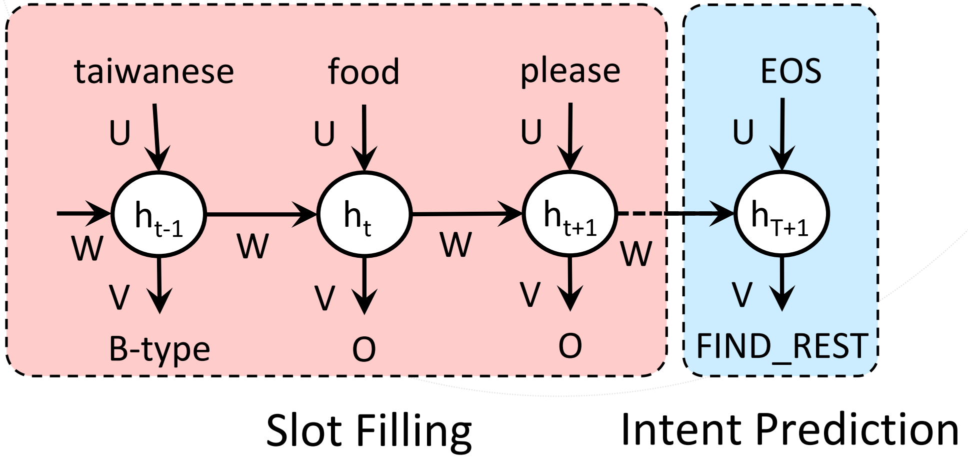

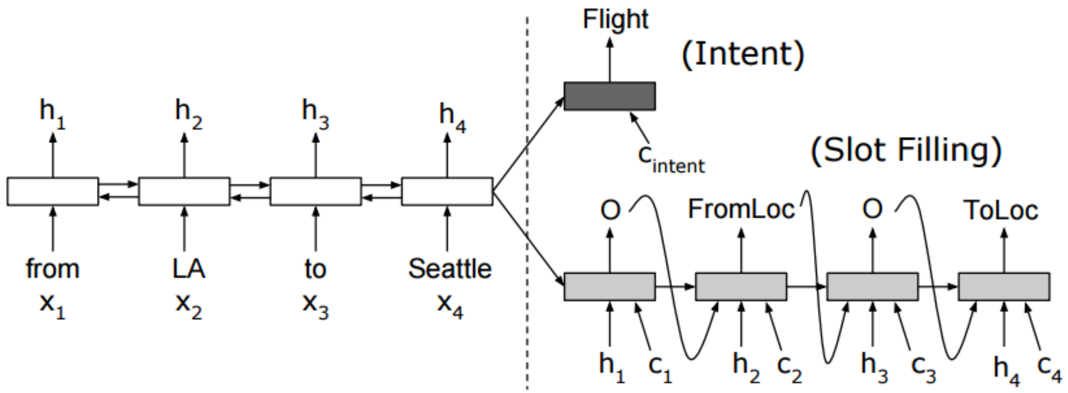

Slot-Gated Joint SLU (Goo+, 2018)

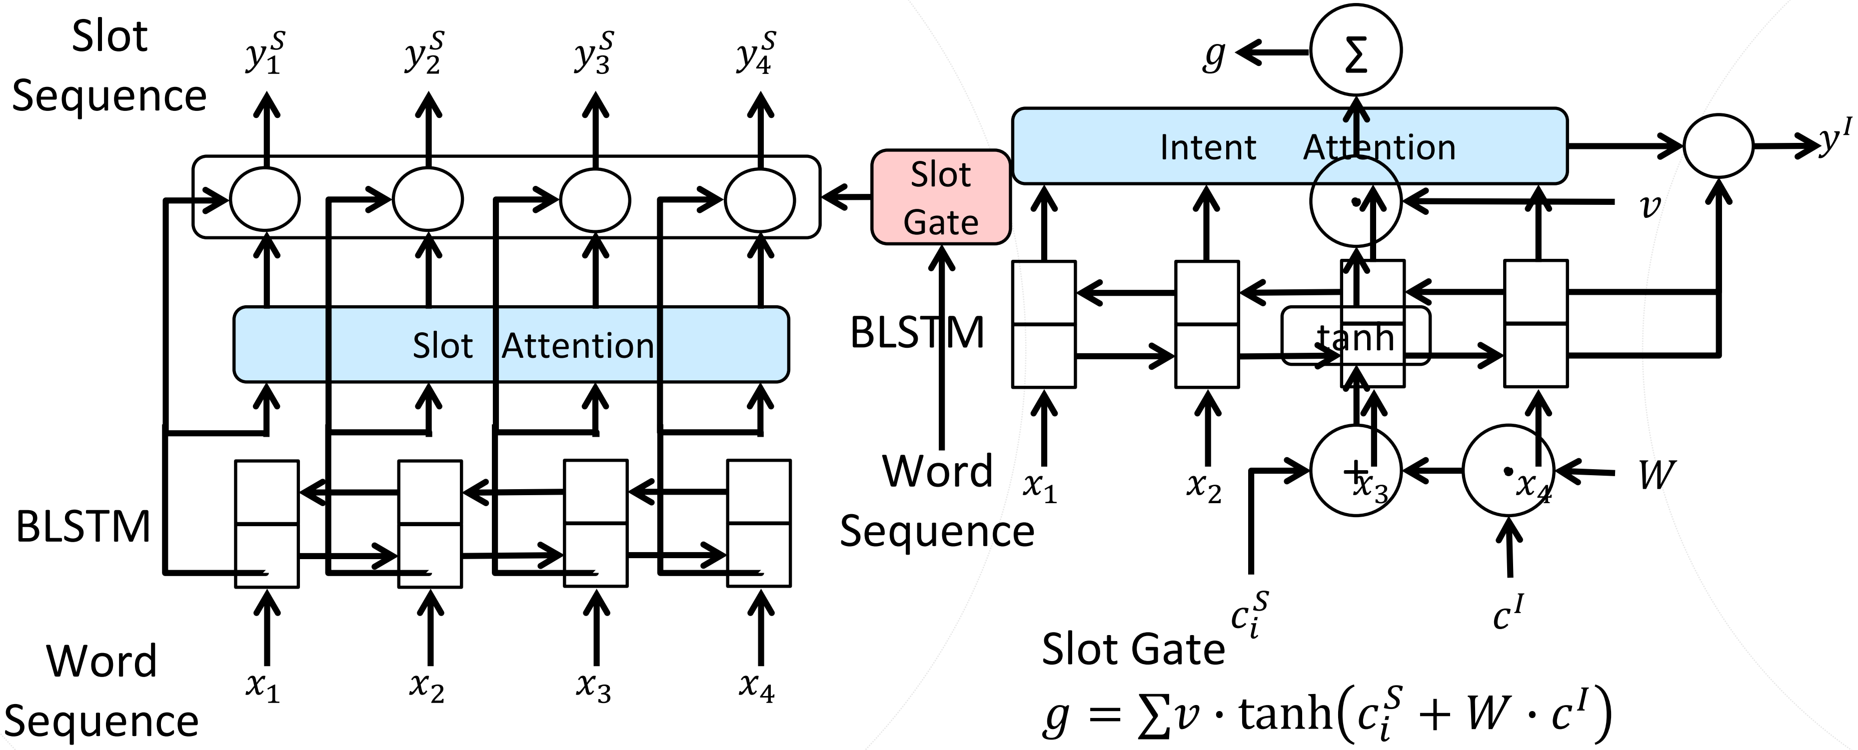

Slot Gate

$$g = \sum v \cdot \tanh\left(c_i^S + W \cdot c^I\right)$$

Slot Prediction

$$y_i^S = softmax\left(W^S\left(h_i + \boldsymbol{g} \cdot c_i^S\right) + b^S\right)$$

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

Material: http://deepdialogue.miulab.tw

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")

NTU MIULAB

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

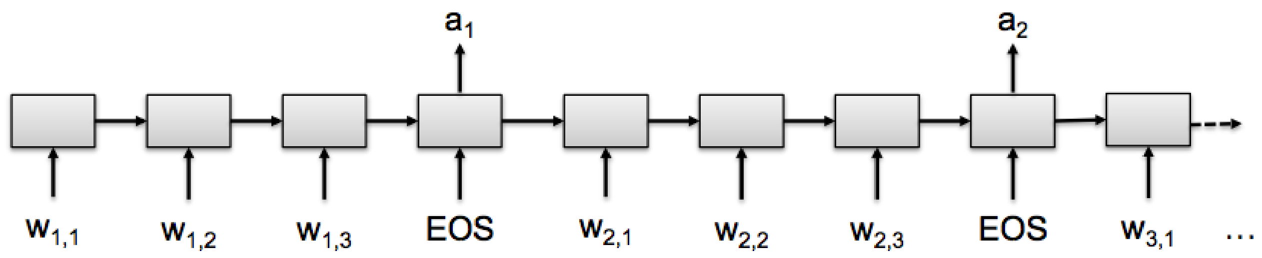

- Extension: LSTM with speaker role dependent layers

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

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"

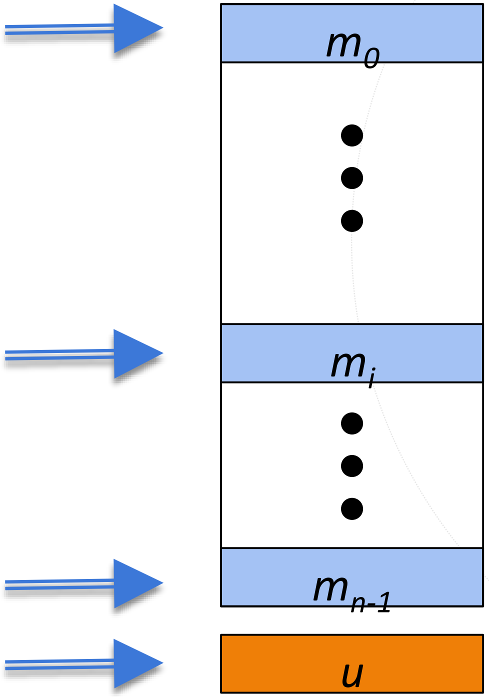

E2E MemNN for Contextual LU (Chen+, 2016)

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

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

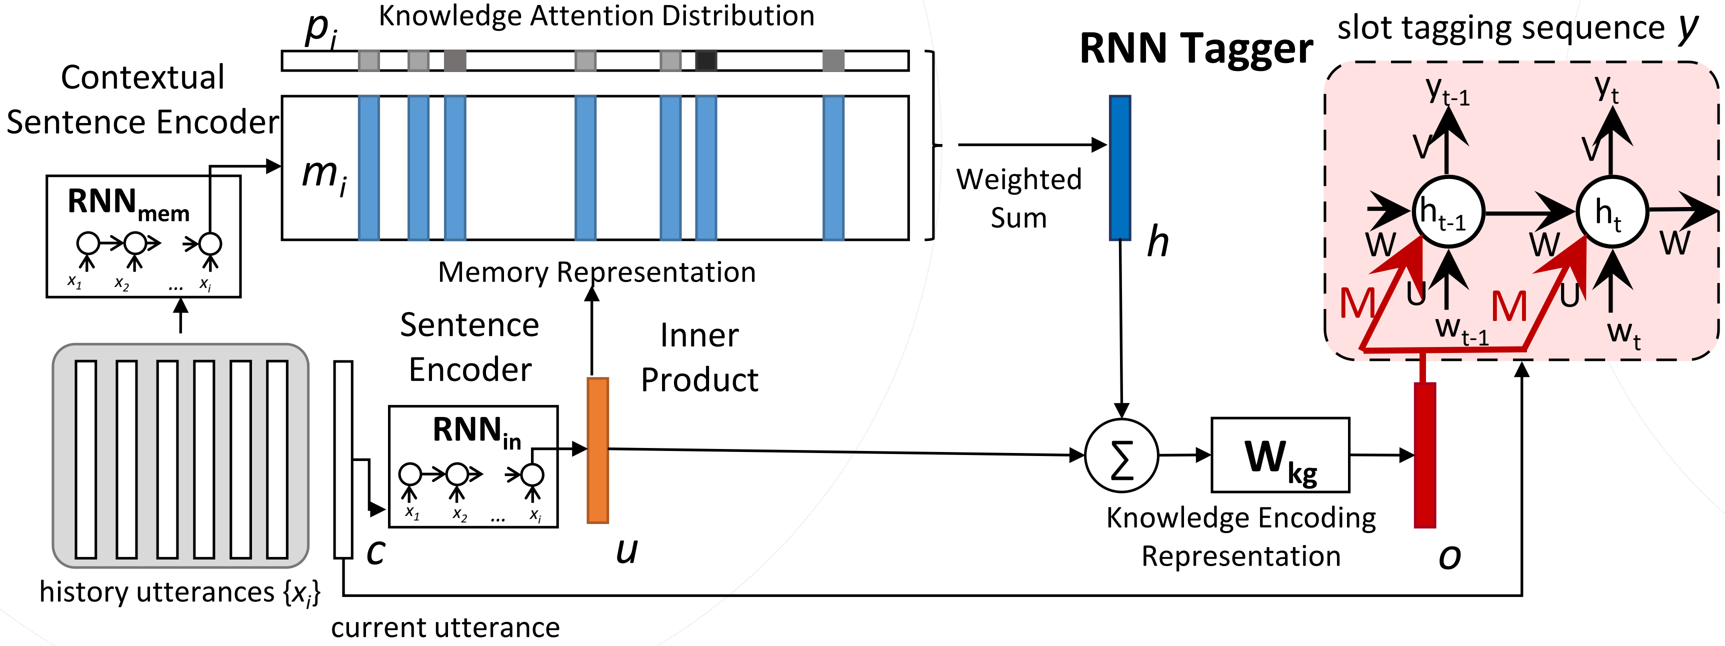

Idea: additionally incorporating contextual knowledge during slot tagging
→ track dialogue states in a latent way

Material: http://deepdialogue.miulab.tw

Analysis of Attention

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"

Role-Based & Time-Aware Attention (Su+, 2018)

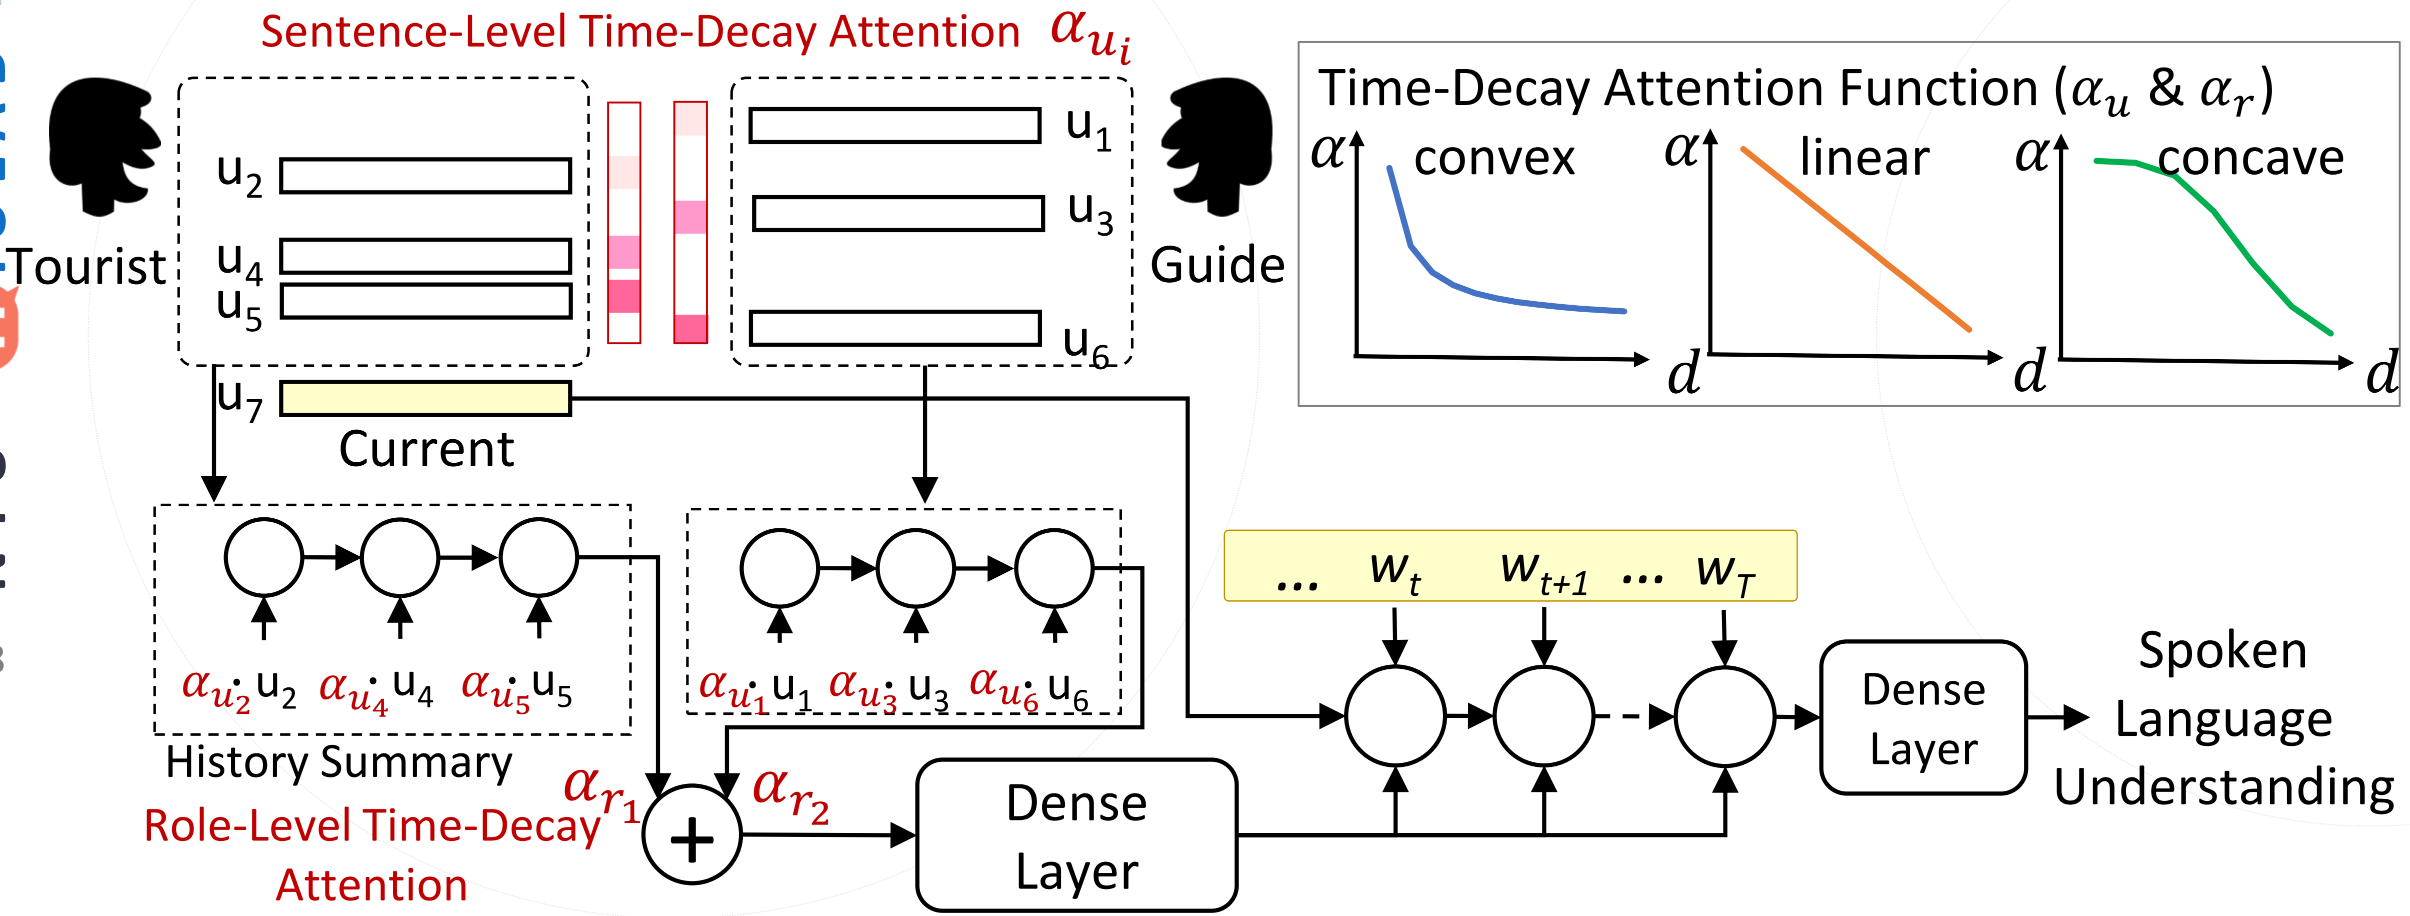

Learnable Time-Decay Attention (Su+, 2019)

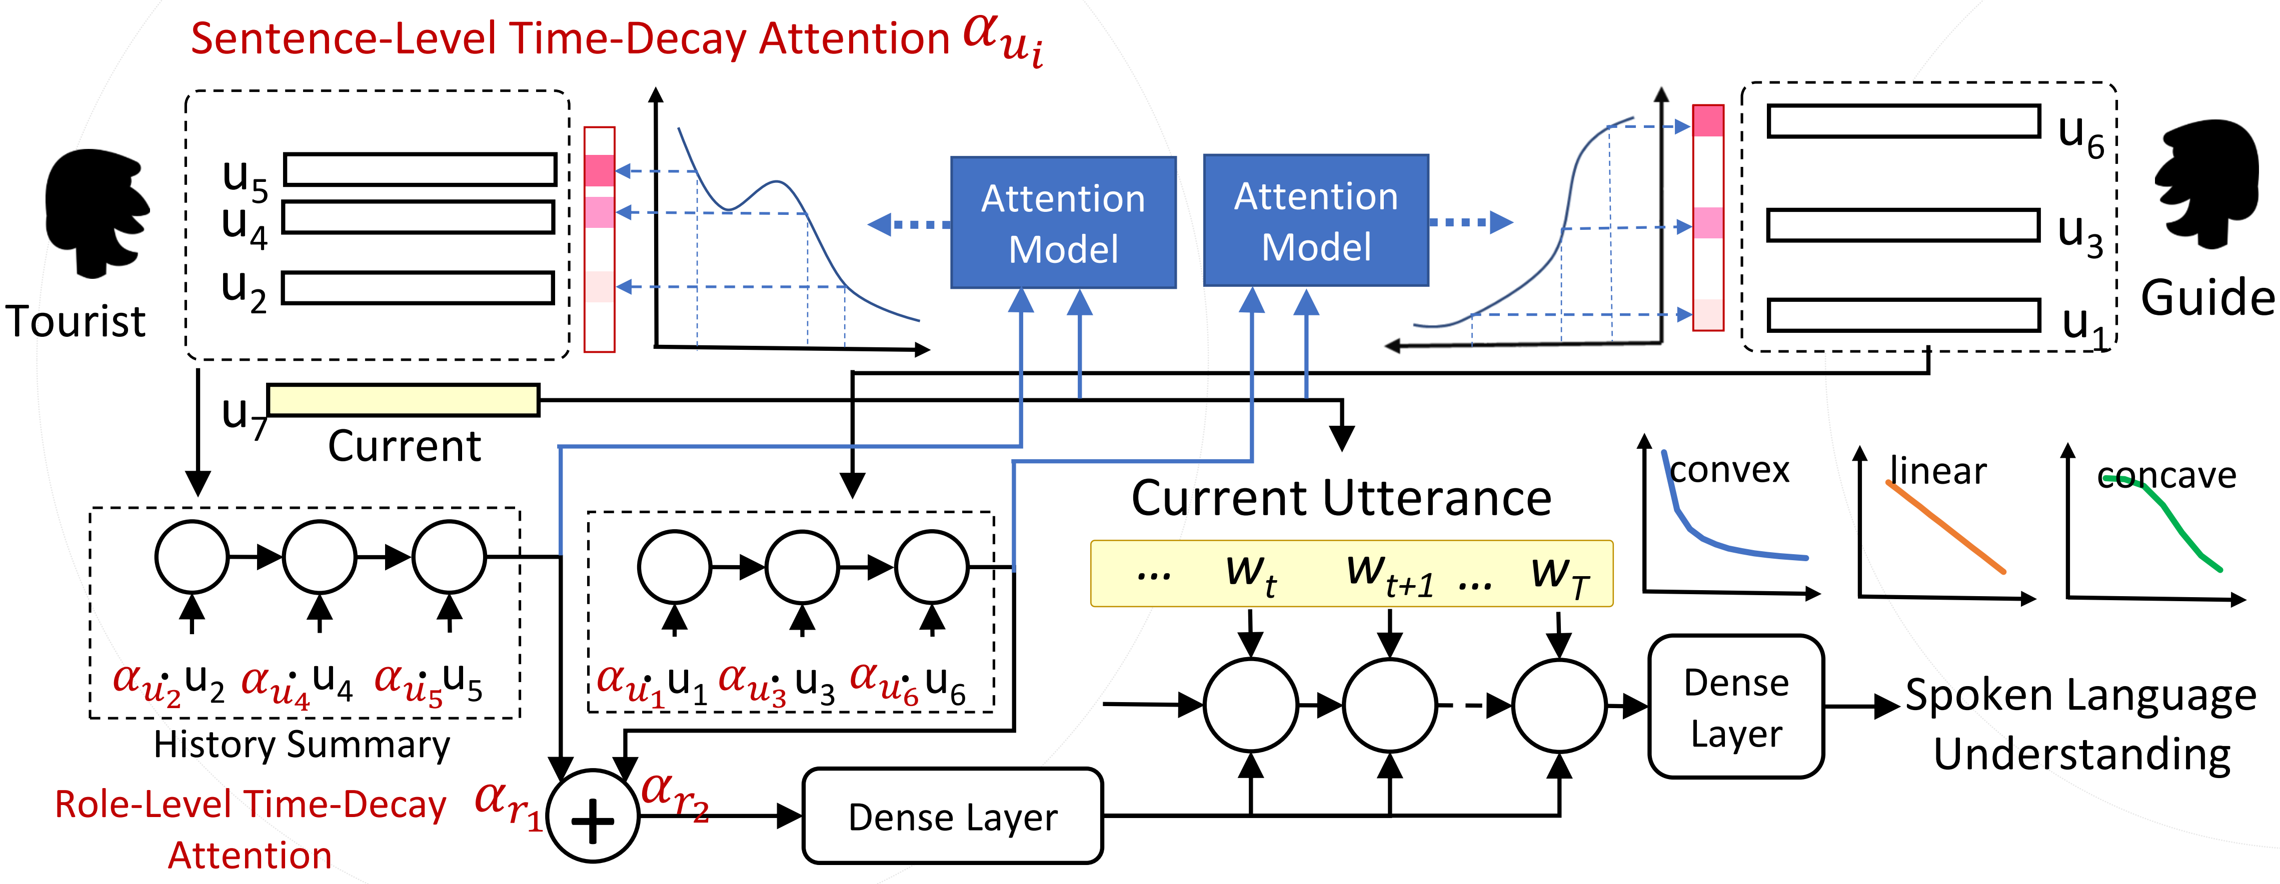

Structural LU (Chen et al., 2016)

- K-SAN: prior knowledge as a teacher

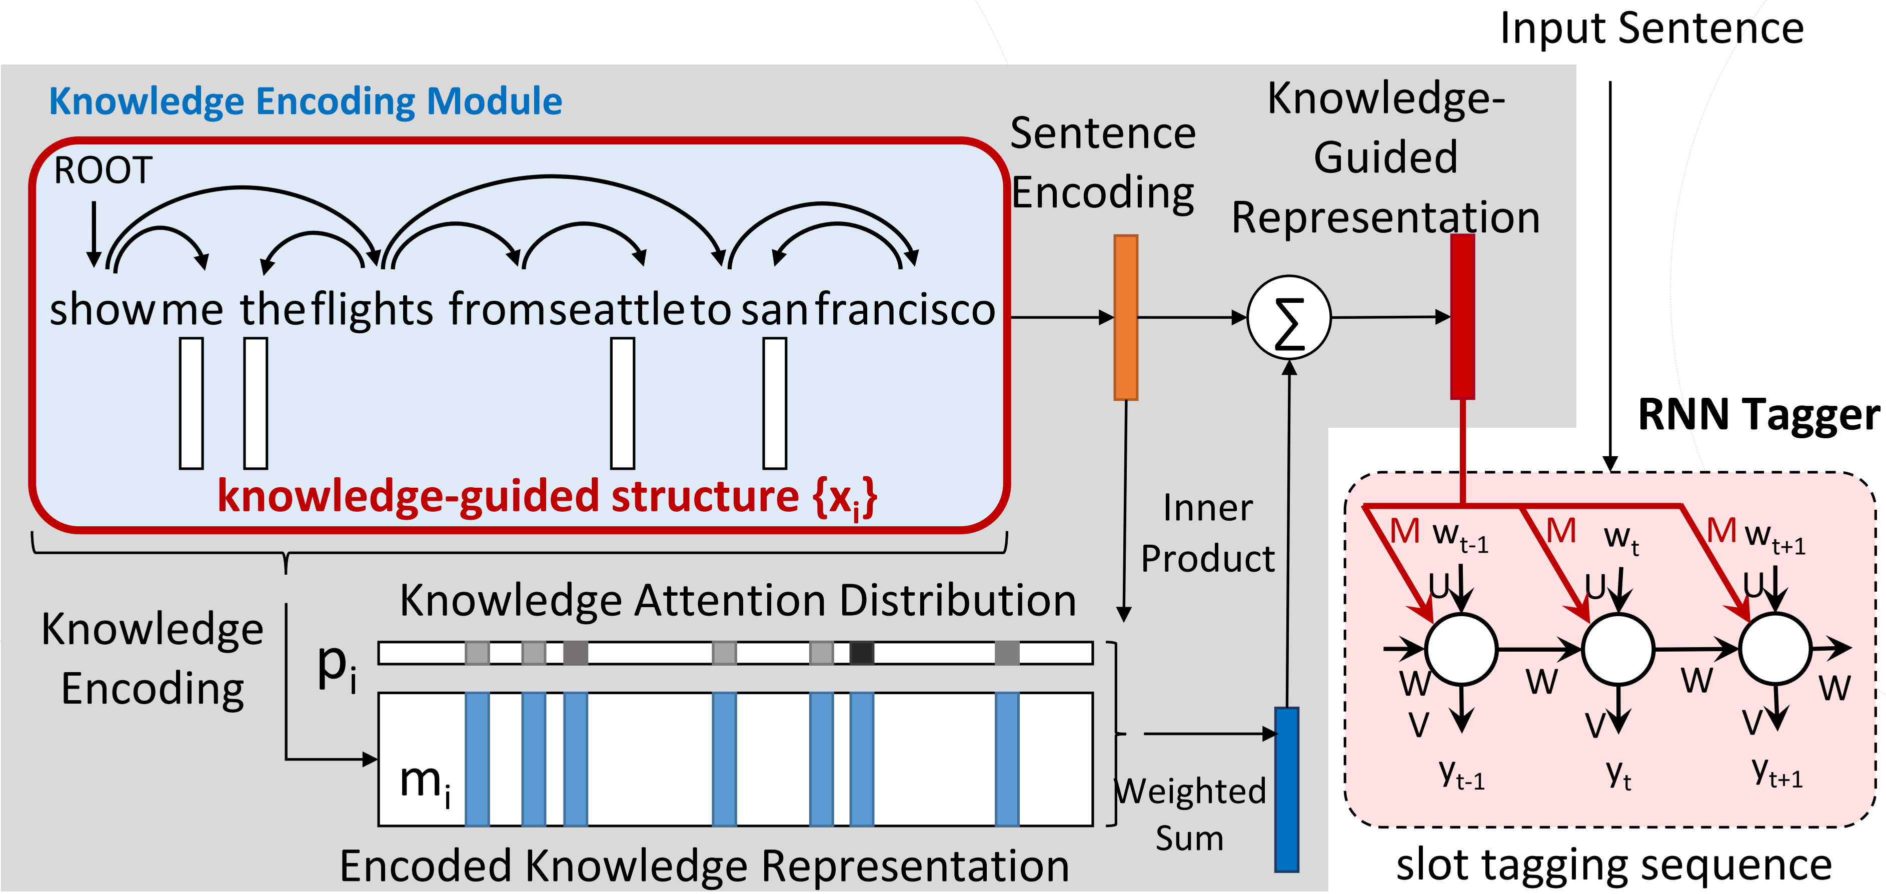

Structural LU (Chen et al., 2016)

- Sentence structural knowledge stored as memory

Sentence *s* show me the flights from seattle to san francisco

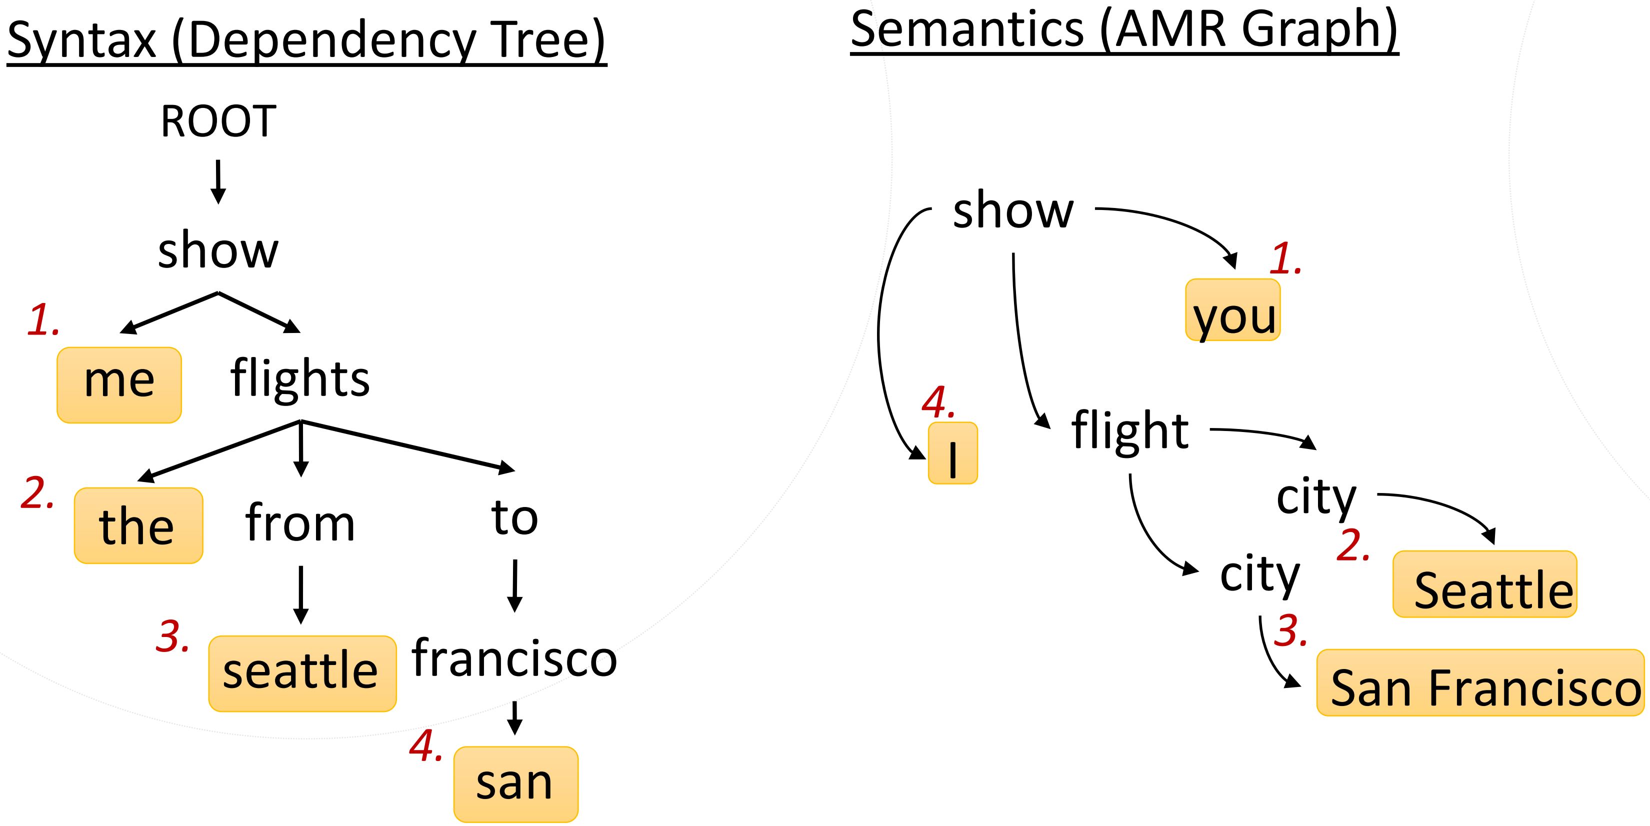

Structural LU (Chen et al., 2016)

- 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.

Semantic Frame Representation

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

Restaurant Domain

find me a cheap taiwanese restaurant in oakland

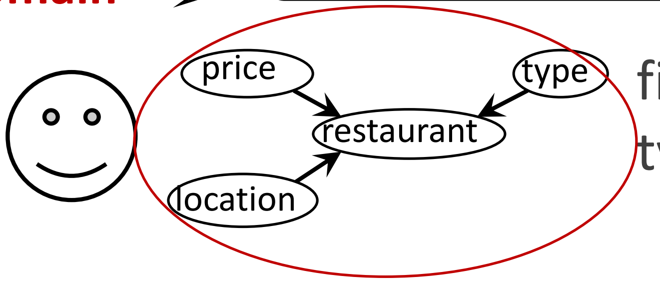

find_restaurant (price="cheap", type="taiwanese", location="oakland")

Movie Domain

show me action movies directed by james cameron

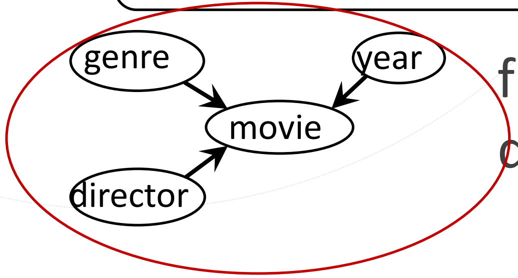

find_movie (genre="action", director="james cameron")

LU – Learning Semantic Ontology (Chen+, 2013)

- Learning key domain concepts from goal-oriented human-human conversations
  - Clustering with mutual information and KL divergence (Chotimongkol & Rudnicky, 2002)
  - Spectral clustering based slot ranking model (Chen et al., 2013)
    - Use a state-of-the-art frame-semantic parser trained for FrameNet
    - Adapt the generic output of the parser to the target semantic space

LU – Intent Expansion (Chen+, 2016)

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

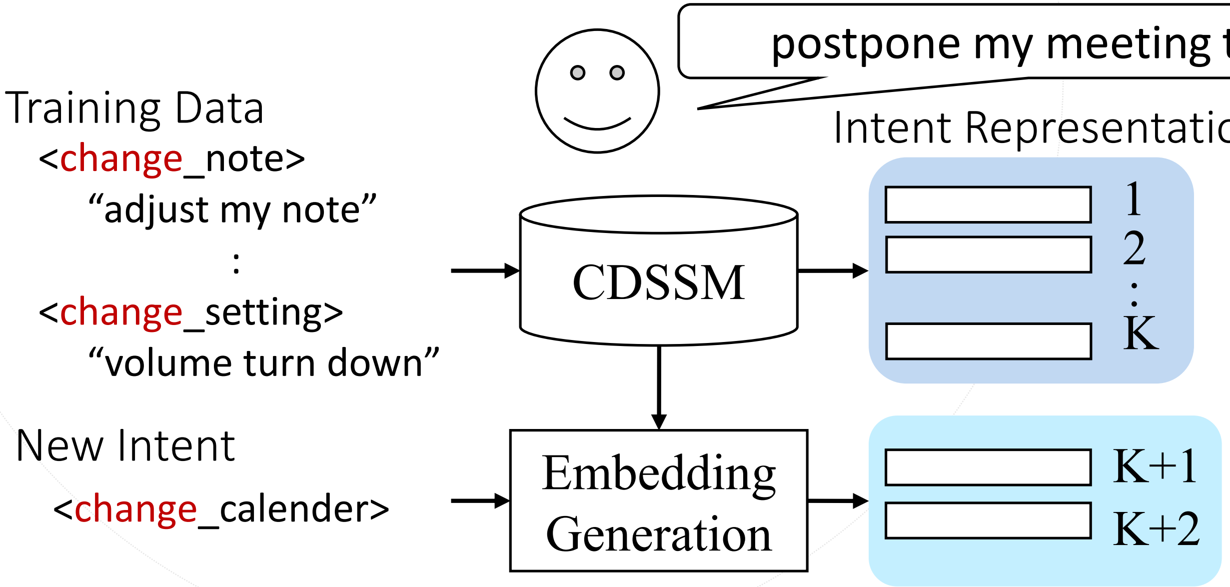

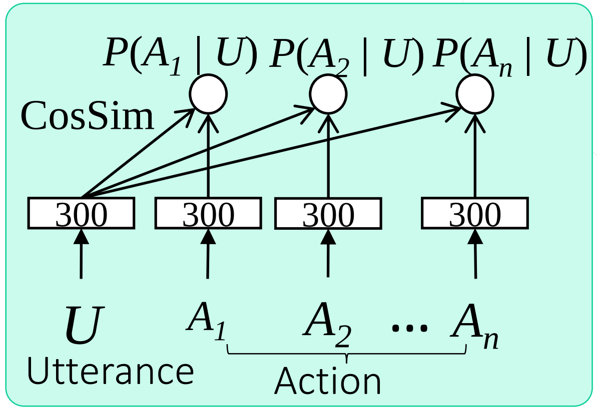

The dialogue act representations can be automatically learned for other domains

LU – Language Extension (Upadhyay+, 2018)

- Source language: English (full annotations)
- Target language: Hindi (limited annotations)

RT: round trip, FC: from city, TC: to city, DDN: departure day name

Utt:	find	a	one	way	flight	from	boston	to	atlanta	on	wednesday
Slots:	O	O	B-RT	I-RT	O	O	B-FC	O	B-TC	O	B-DDN

(a) English Utterance

Utt:	बुधवार	को	बोसटन	से	अटलांटा	तक	जाने	वाली	एकतरफ़ा	उड़ाने	खोजें
Slots:	B-DDN	O	B-FC	O	B-TC	O	O	O	B-RT	O	O

(b) Hindi Utterance

LU – Language Extension (Upadhyay+, 2018)

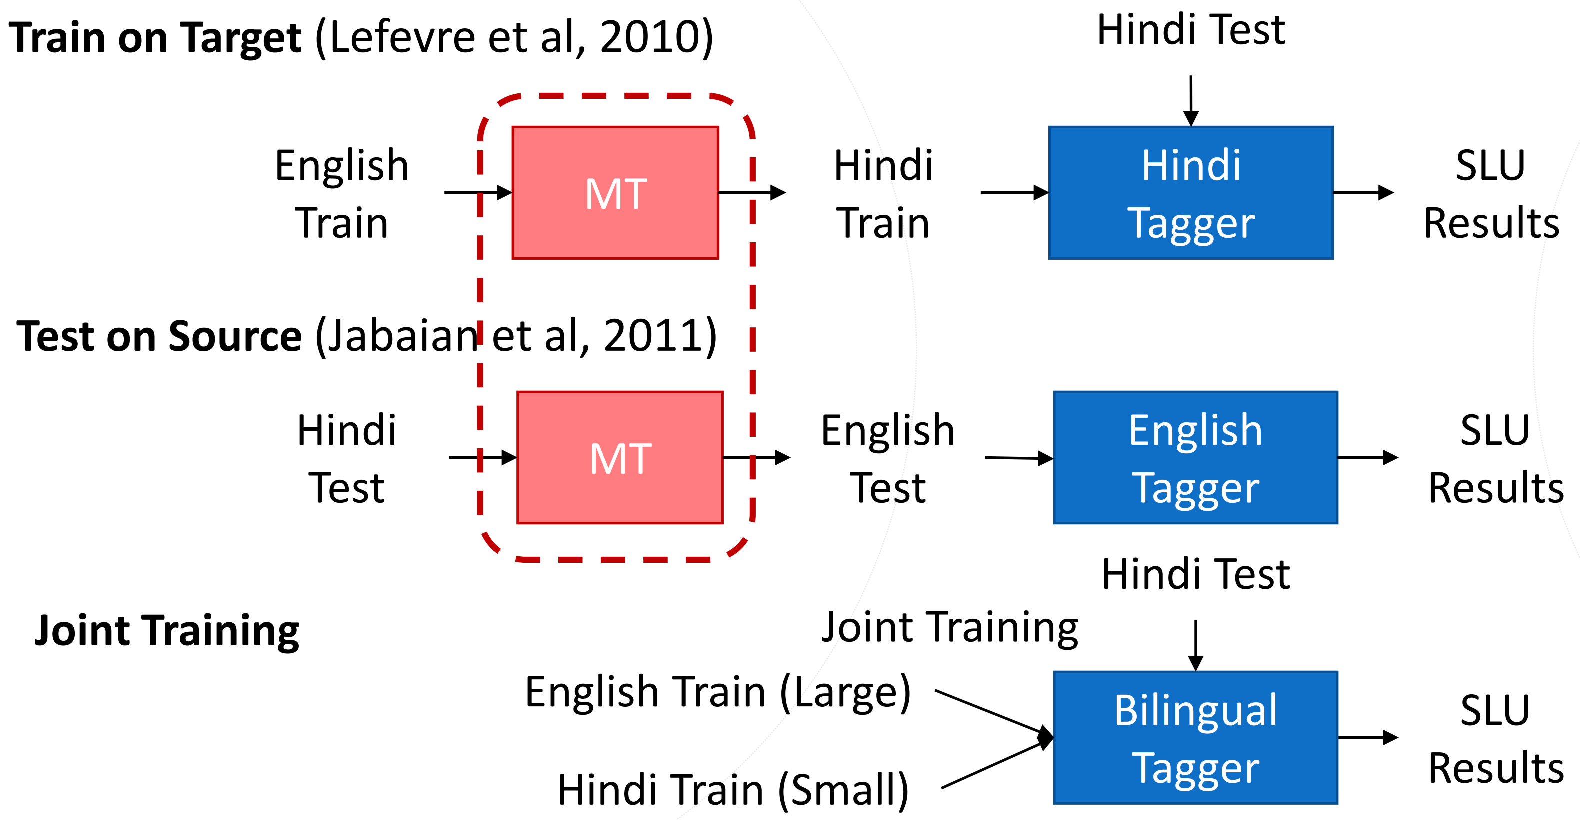

MT system is not required and both languages can be processed by a single model

Outline

- Introduction
- Background Knowledge
- **Modular Dialogue System**
  - Spoken/Natural Language Understanding (SLU/NLU)
  - **Dialogue Management**
    - **Dialogue State Tracking (DST)**
    - Dialogue Policy Optimization
  - Natural Language Generation (NLG)
- System Evaluation
- Recent Trends of Learning Dialogues

Elements of Dialogue Management

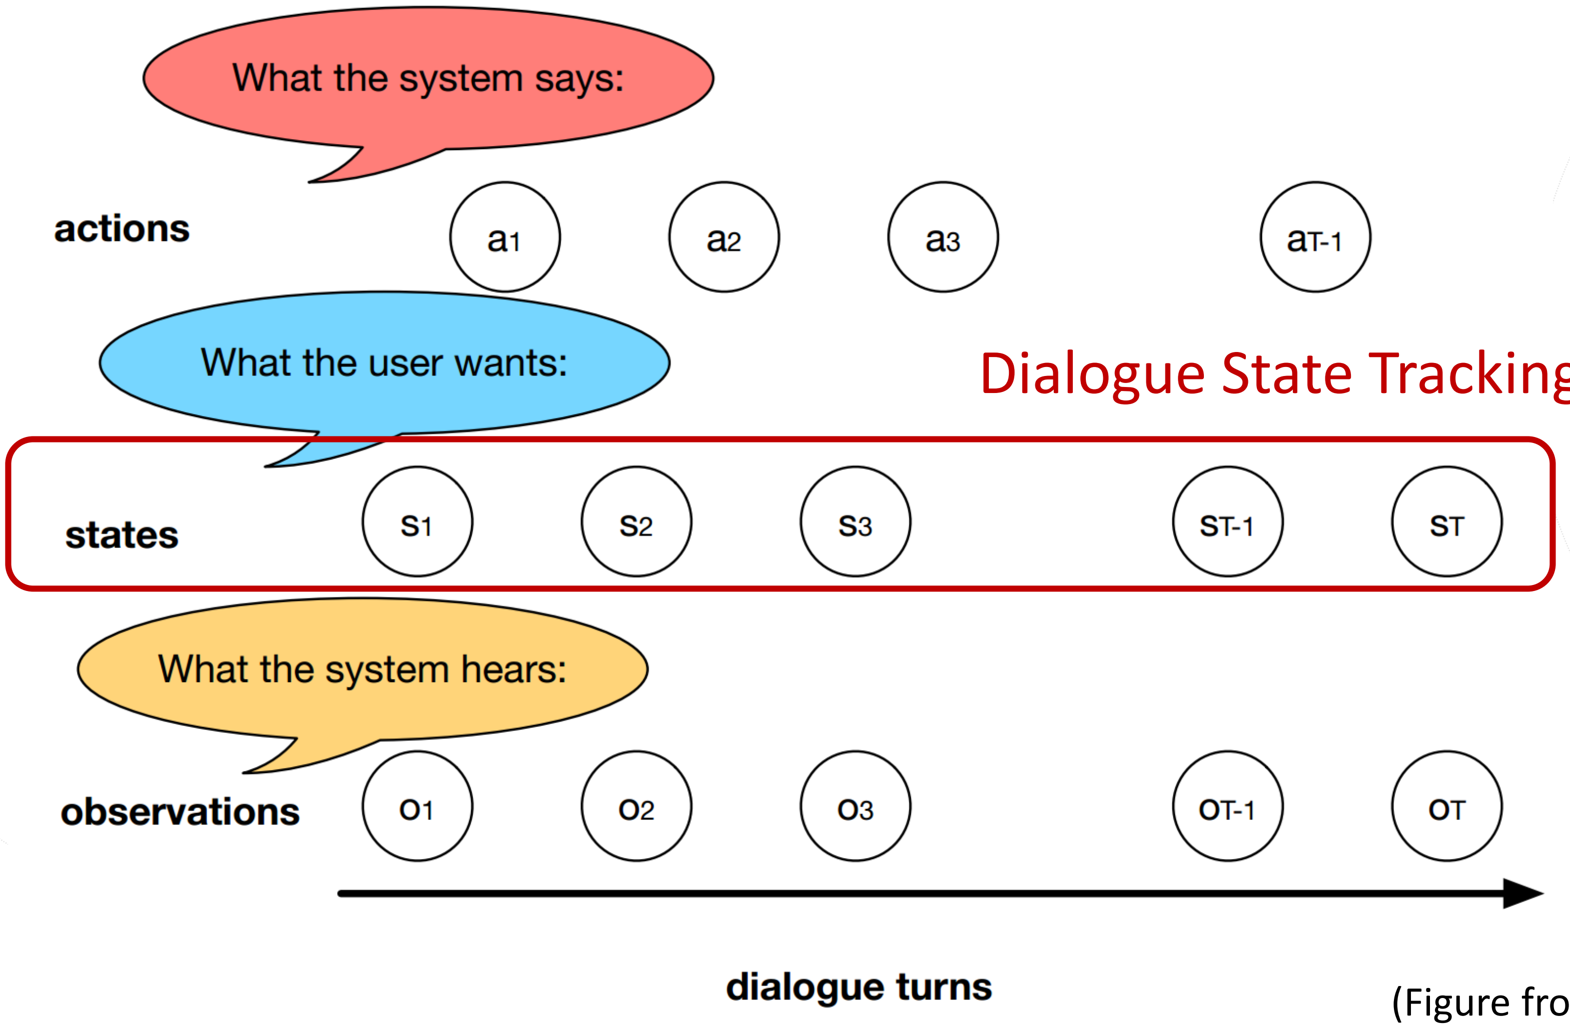

(Figure from Gašić)

Material: http://deepdialogue.miulab.tw

Dialogue State Tracking (DST)

- Maintain a probabilistic distribution instead of a 1-best prediction for better robustness

Turn 1
Kind
Android

Turn 2
Note
Android

Turn 1	
Kind	0.5
Android	0.3

Turn 2	
Note	0.4
Android	0.3

Incorrect for both!

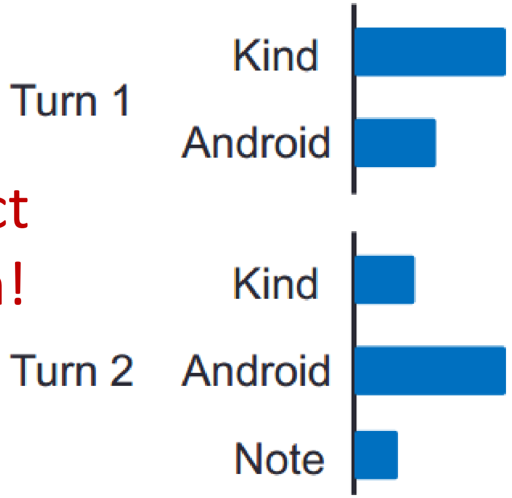

Dialogue State Tracking (DST)

- 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

- 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

Date	11/15/17			
Time	6 pm	7 pm	8 pm	9 pm
#People	2			
Theater	Century 16 Shoreline			
Movie	Inferno			

Less Likely — More Likely

Multi-Domain Dialogue State Tracking

- 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

Date	11/15/17			
Time	6:30 pm	7:30 pm	8:45 pm	9:45 pm
#People	2			
Theater	Century 16 Shoreline			
Movie	Inferno			

Restaurants

Date	11/15/17		
Time	6:00 pm	6:30 pm	7:00 pm
Restaurant	Cascal		
#People	2		

Less Likely … More Likely

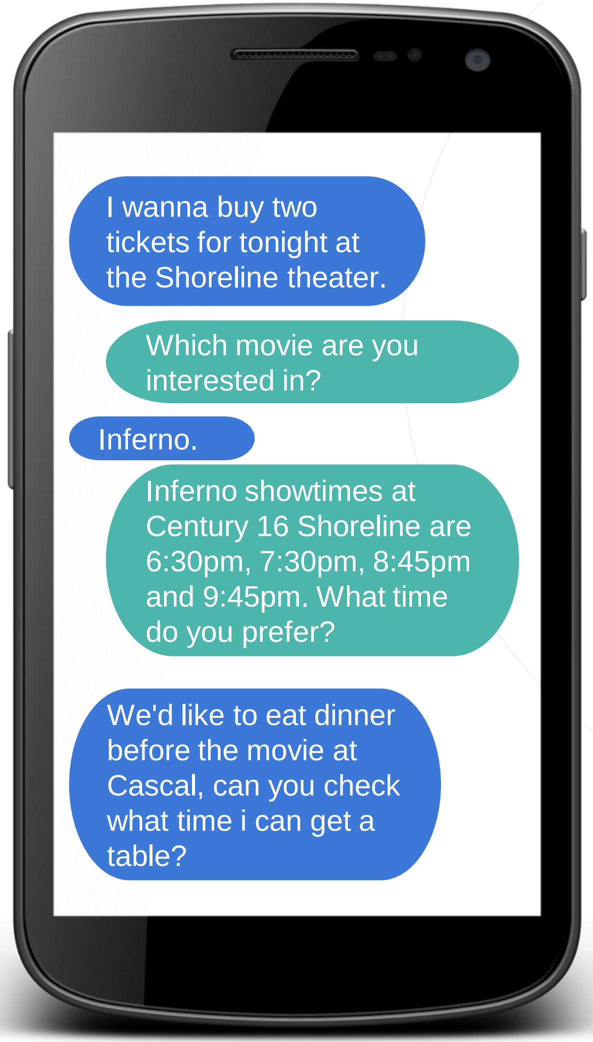

Material: http://deepdialogue.miulab.tw

Multi-Domain Dialogue State Tracking

- 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

Date	11/15/17			
Time	6:30 pm	7:30 pm	8:45 pm	9:45 pm
#People	2			
Theater	Century 16 Shoreline			
Movie	Inferno			

Restaurants

Date	11/15/17		
Time	6:00 pm	6:30 pm	7:00 pm
Restaurant	Cascal		
#People	2		

Less Likely — More Likely

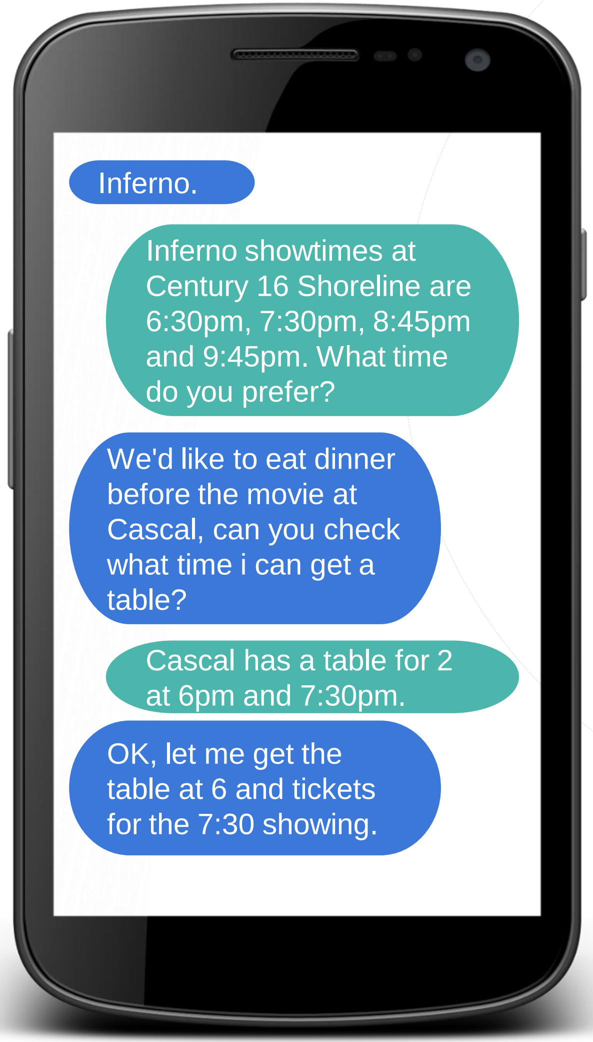

Material: http://deepdialogue.miulab.tw

RNN-CNN DST (Mrkšić+, 2015)

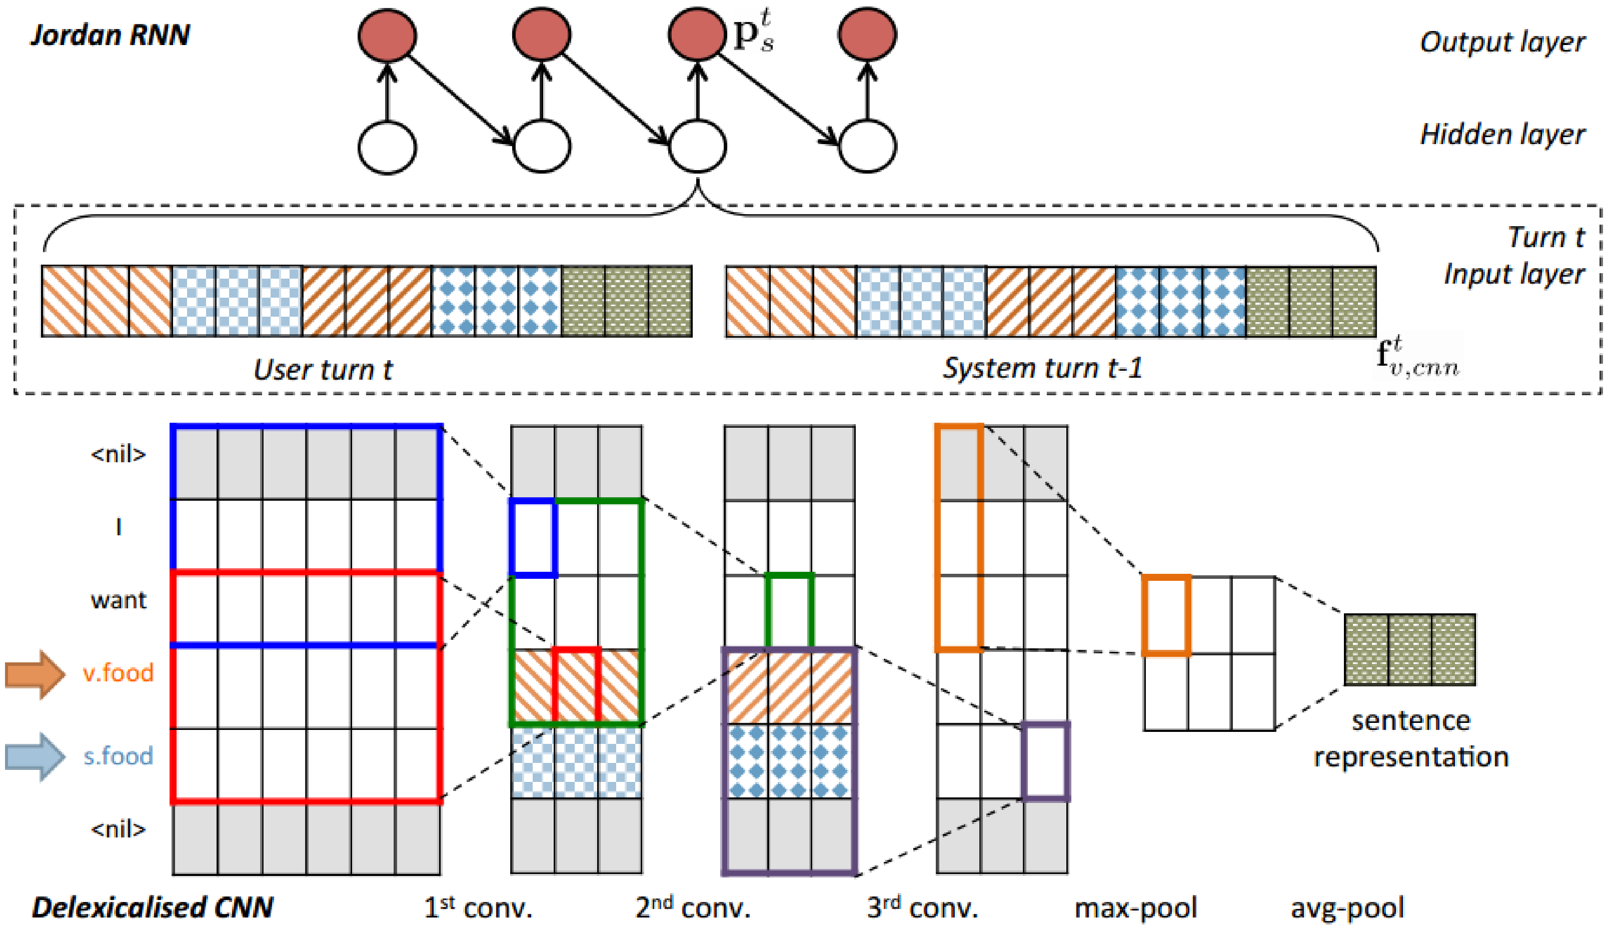

(Figure from Wen et al, 2016)

Neural Belief Tracker (Mrkšić+, 2016)

- Candidate pairs are considered

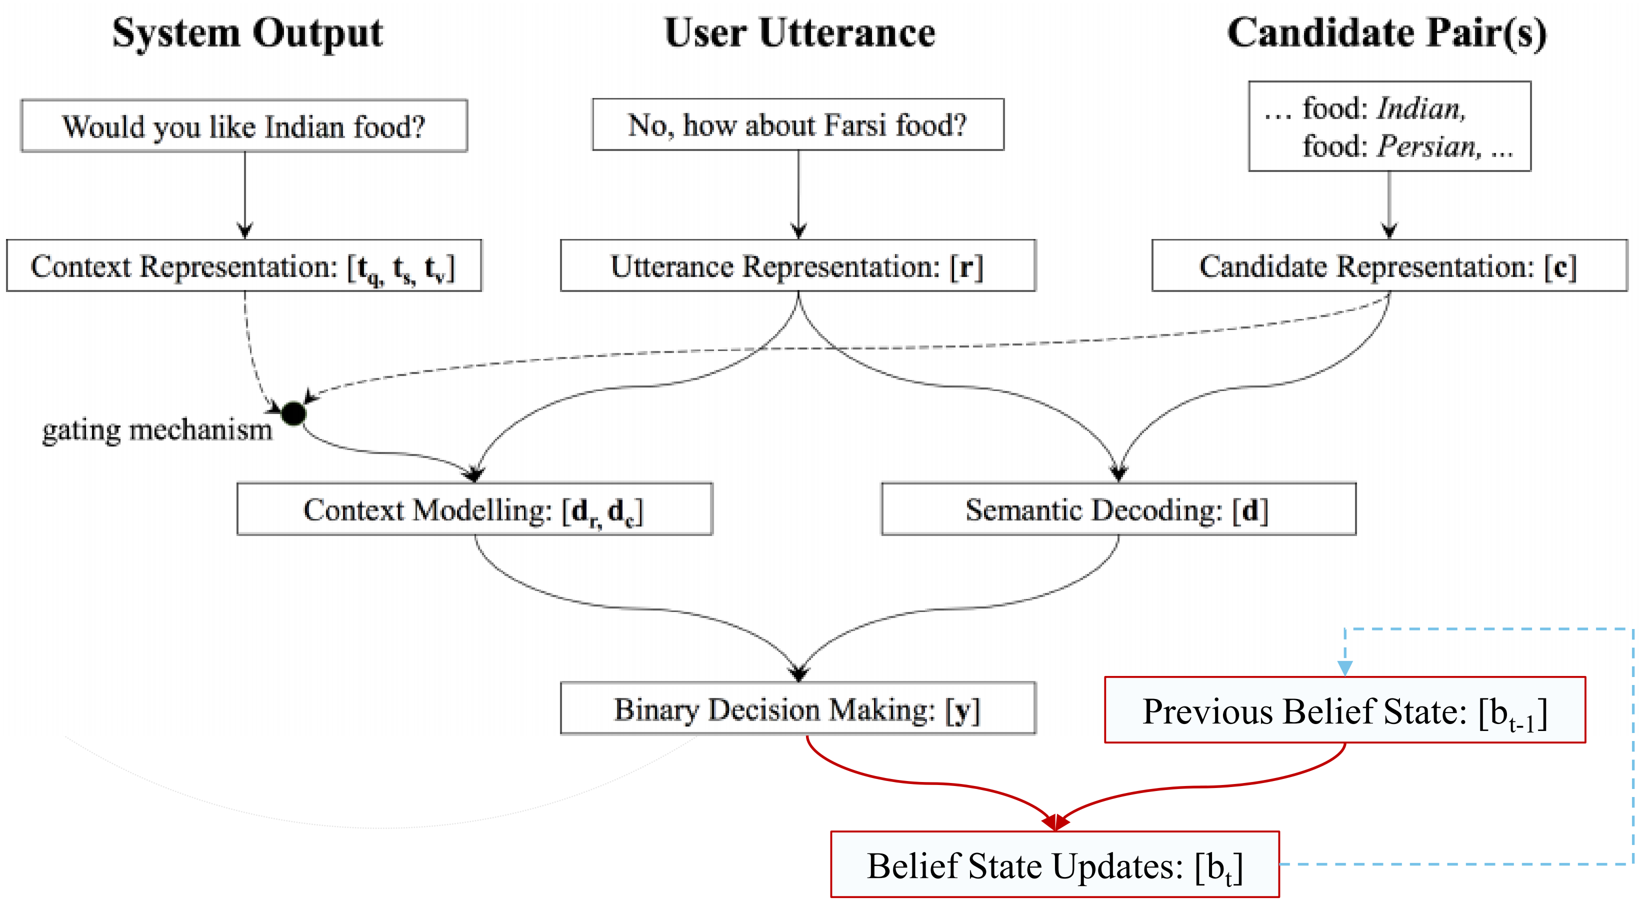

Global-Locally Self-Attentive DST (Zhong+, 2018)

- More advanced encoder
  - Global modules share parameters for all slots
  - Local modules learn slot-specific feature representations

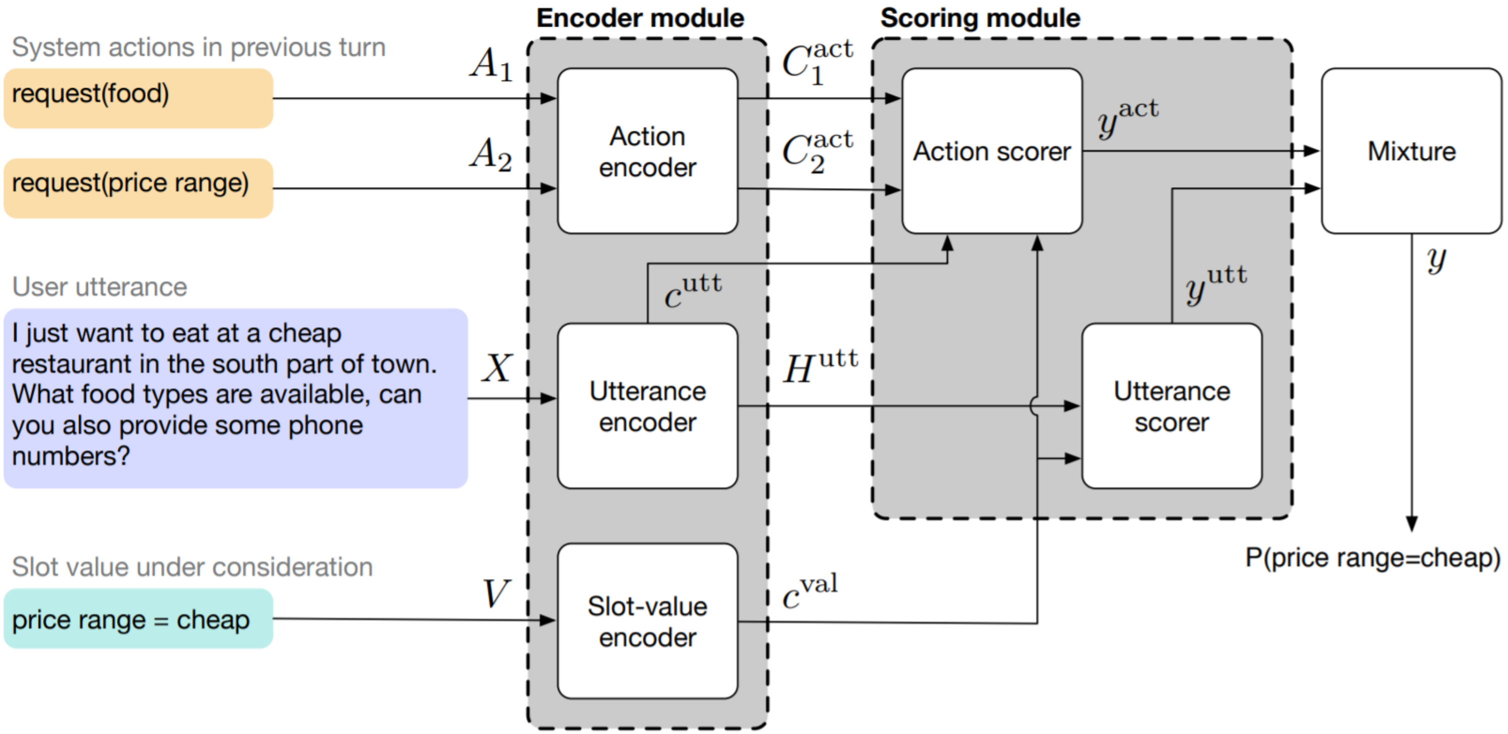

Material: http://deepdialogue.miulab.tw

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

DST Evaluation

- 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

DST – Language Extension (Shi+, 2016)

- Training a multichannel CNN for each slot
  - Chinese character CNN
  - Chinese word CNN
  - English word CNN

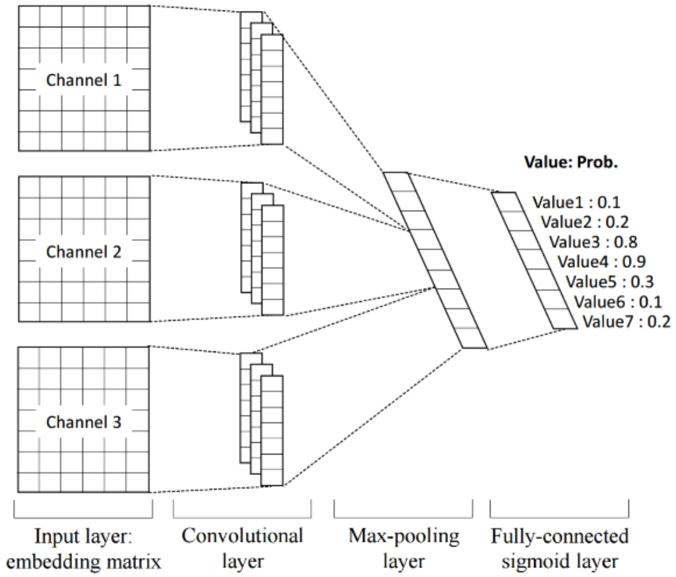

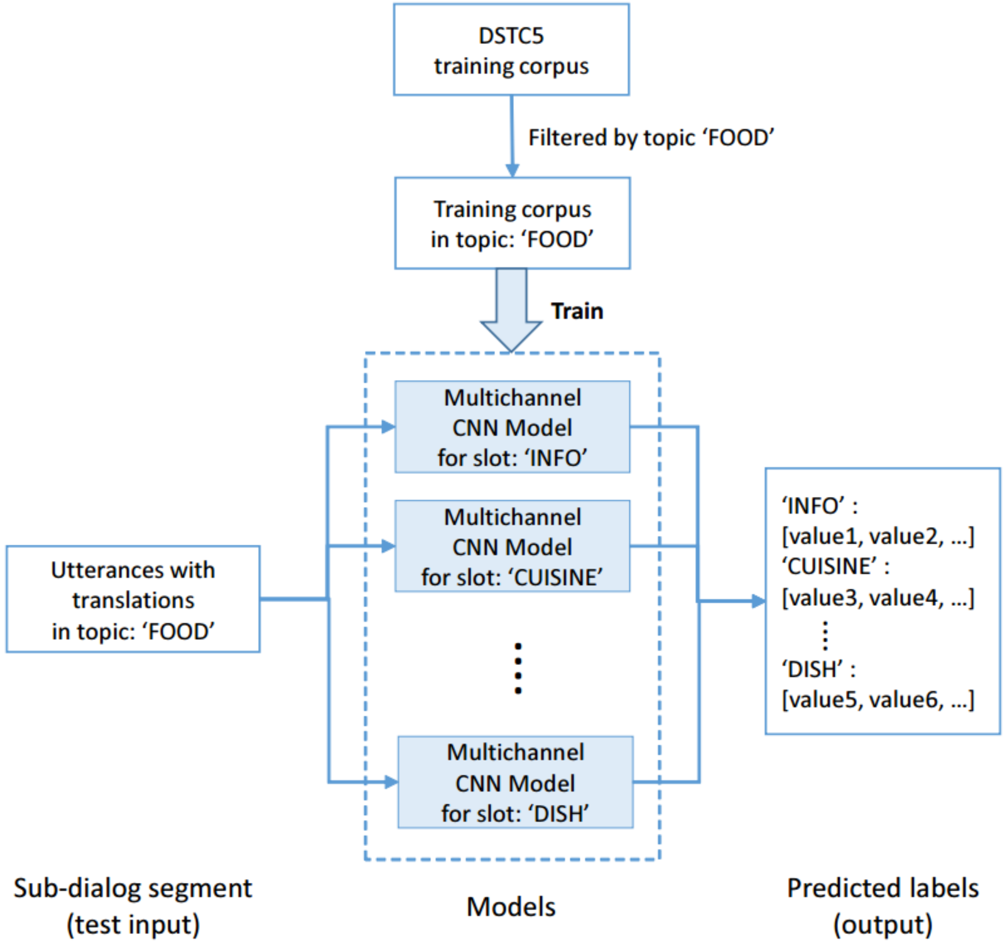

Material: http://deepdialogue.miulab.tw

DST – Task Lineages (Lee & Stent, 2016)

- Slot values shared across tasks
- Utterances with complex constraints on user goals
- Interleaved multiple task discussions

Task Frame:

Connection to Manhattan and find me a Thai restaurant, not Italian

$$\begin{bmatrix} \begin{bmatrix} \textbf{Task} & \text{Transit} \\ \textbf{DAIs} & (0.8, \text{inform}(\text{dest=MH})^{0.1}_{0.7}) \end{bmatrix} \\ \begin{bmatrix} \textbf{Task} & \text{Restaurant} \\ \textbf{DAIs} & (0.7, \text{inform}(\text{food=thai})^{0.9}_{1.2}) \\ & (0.6, \text{deny}(\text{food=italian})^{1.4}_{1.7}) \end{bmatrix} \end{bmatrix}$$

(confidence, dialog act item $^{\text{Start\_time}}_{\text{End\_time}}$)

Task State:

Thai restaurant, not Italian

$$\begin{bmatrix} \textbf{Task} & \text{Restaurant} \\ \textbf{Constraints} & (0.7, \text{food} = \text{thai}) \\ & (0.6, \text{food} \neq \text{italian}) \\ \textbf{DB} & [\text{“Thai To Go”, “Pa de Thai”}] \\ \textbf{Timestamps} & 01/01/2016 : 12\text{-}00\text{-}00 \\ \dots & \dots \end{bmatrix}$$

NTU MIULAB

Material: http://deepdialogue.miulab.tw

DST – Task Lineages (Lee & Stent, 2016)

Turn	User Input
0	Weather in New York. Connection to Manhattan
1	Want to go to Thai
2	I want to travel to Thai

Task Lineages

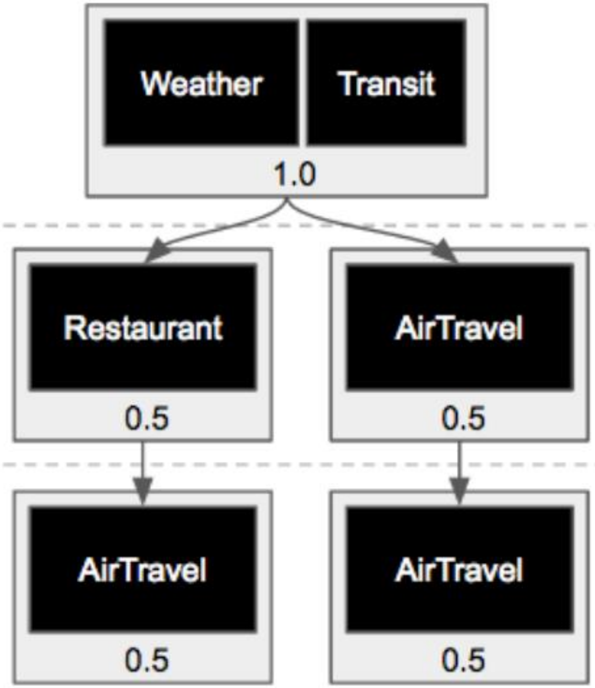

DST – Scalability (Rastogi+, 2017)

- Focus only on the relevant slots
- Better generalization to ASR lattices, visual context, etc.

S> How about 6 pm?
U> I am busy then, book it for 7 pm instead.

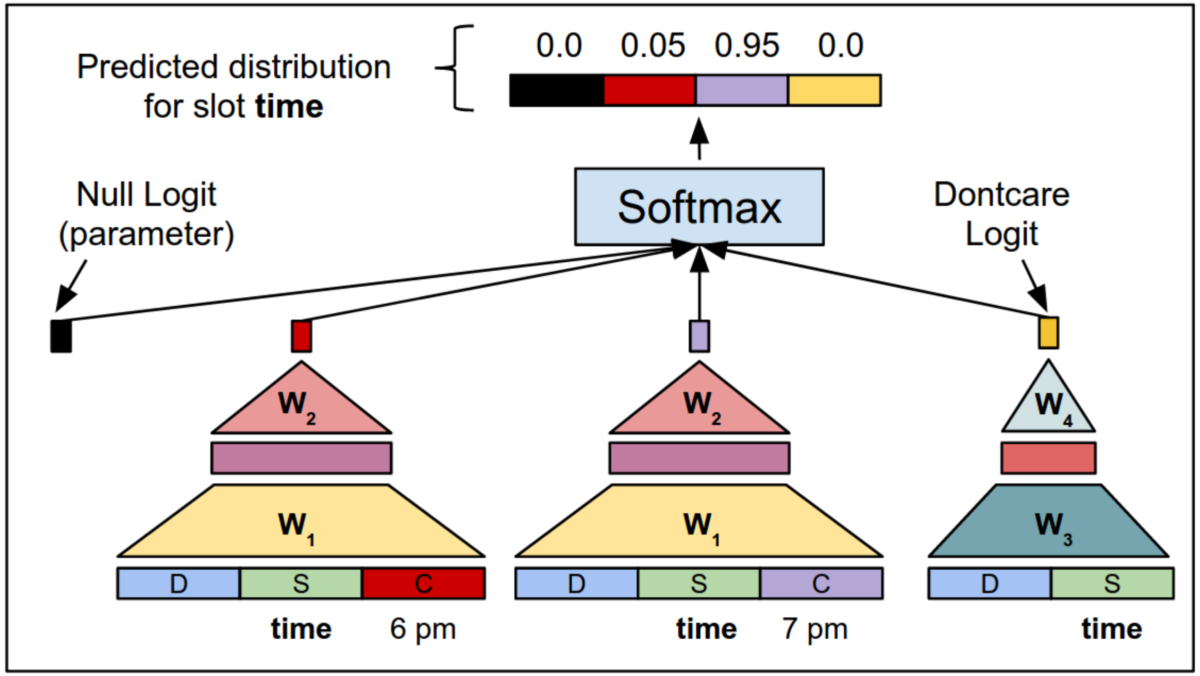

DST – Handling Unknown Values (Xu & Hu, 2018)

- Issue: fixed value sets in DST

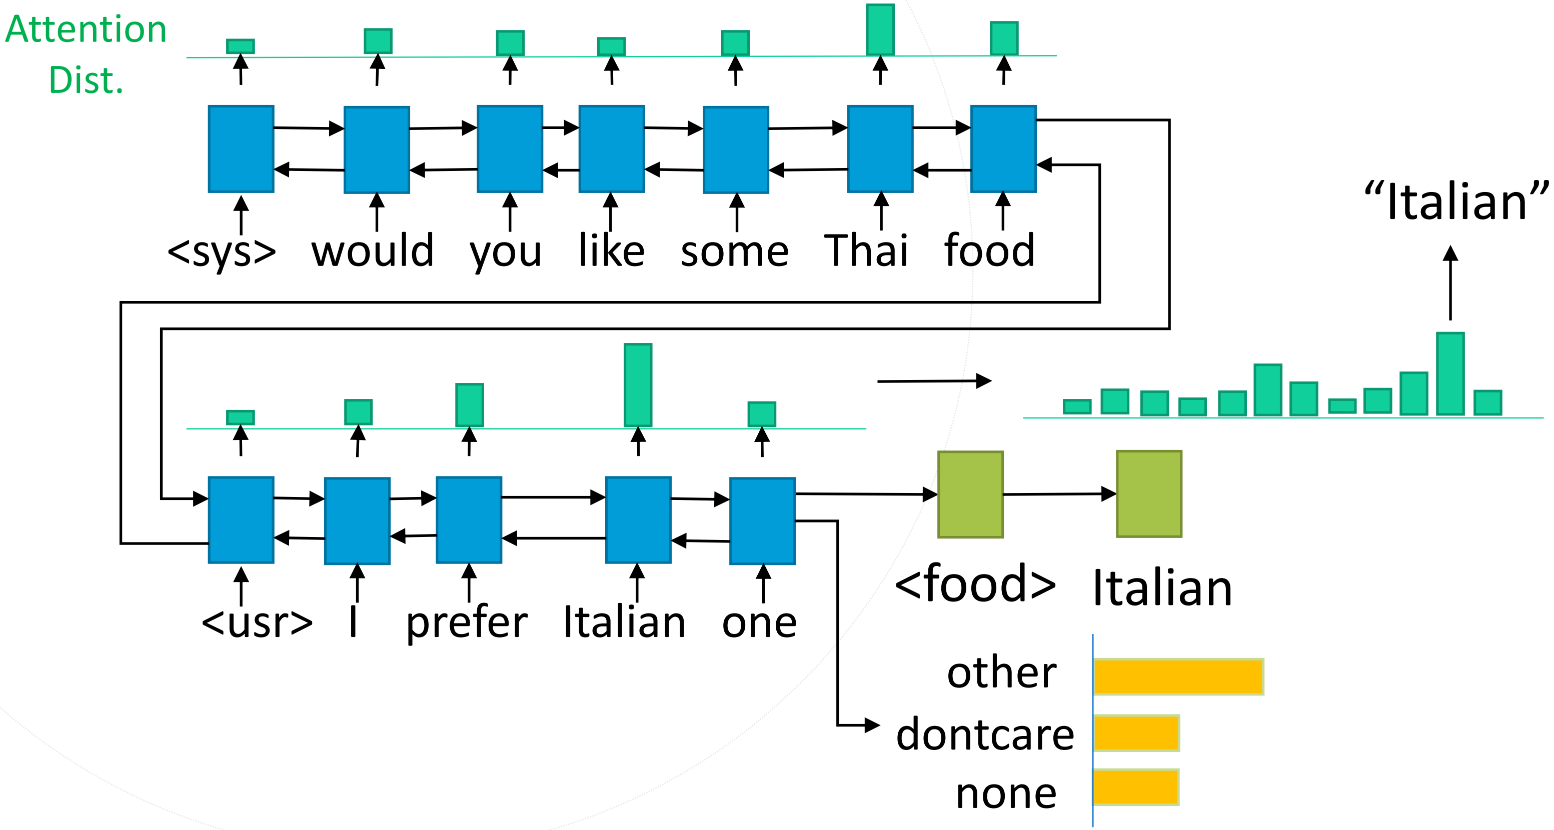

Pointer networks for generating unknown values

Joint NLU and DST (Gupta+, 2018)

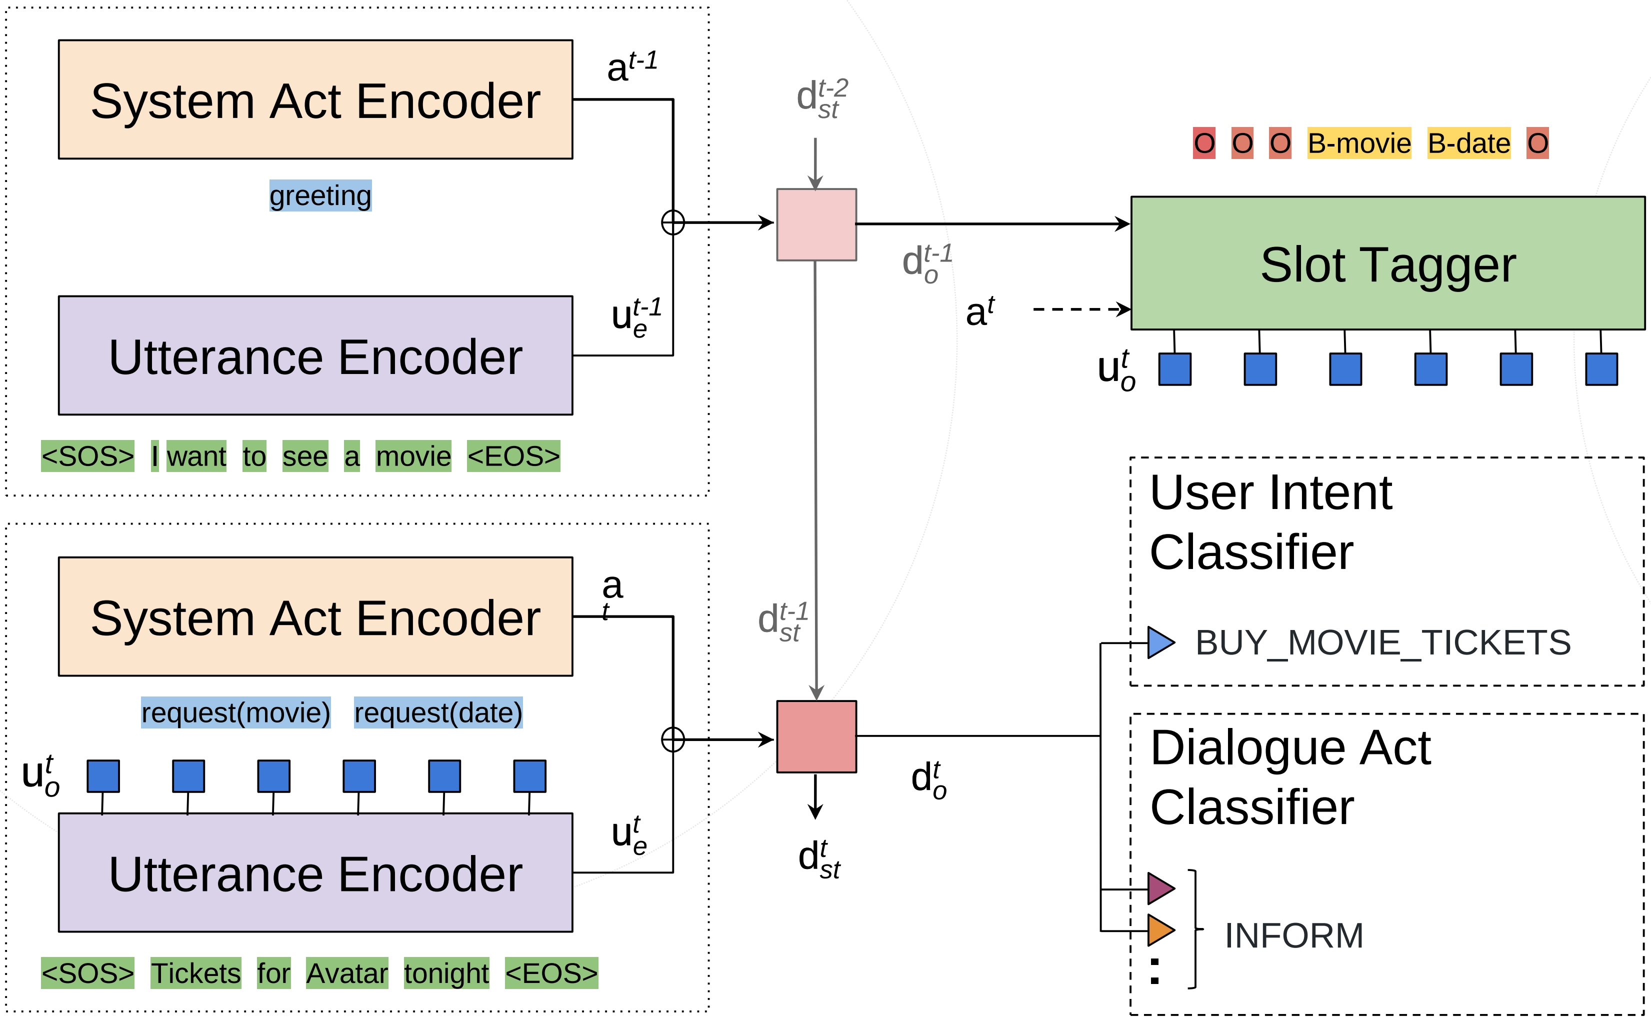

Joint NLU and DST (Gupta+, 2018)

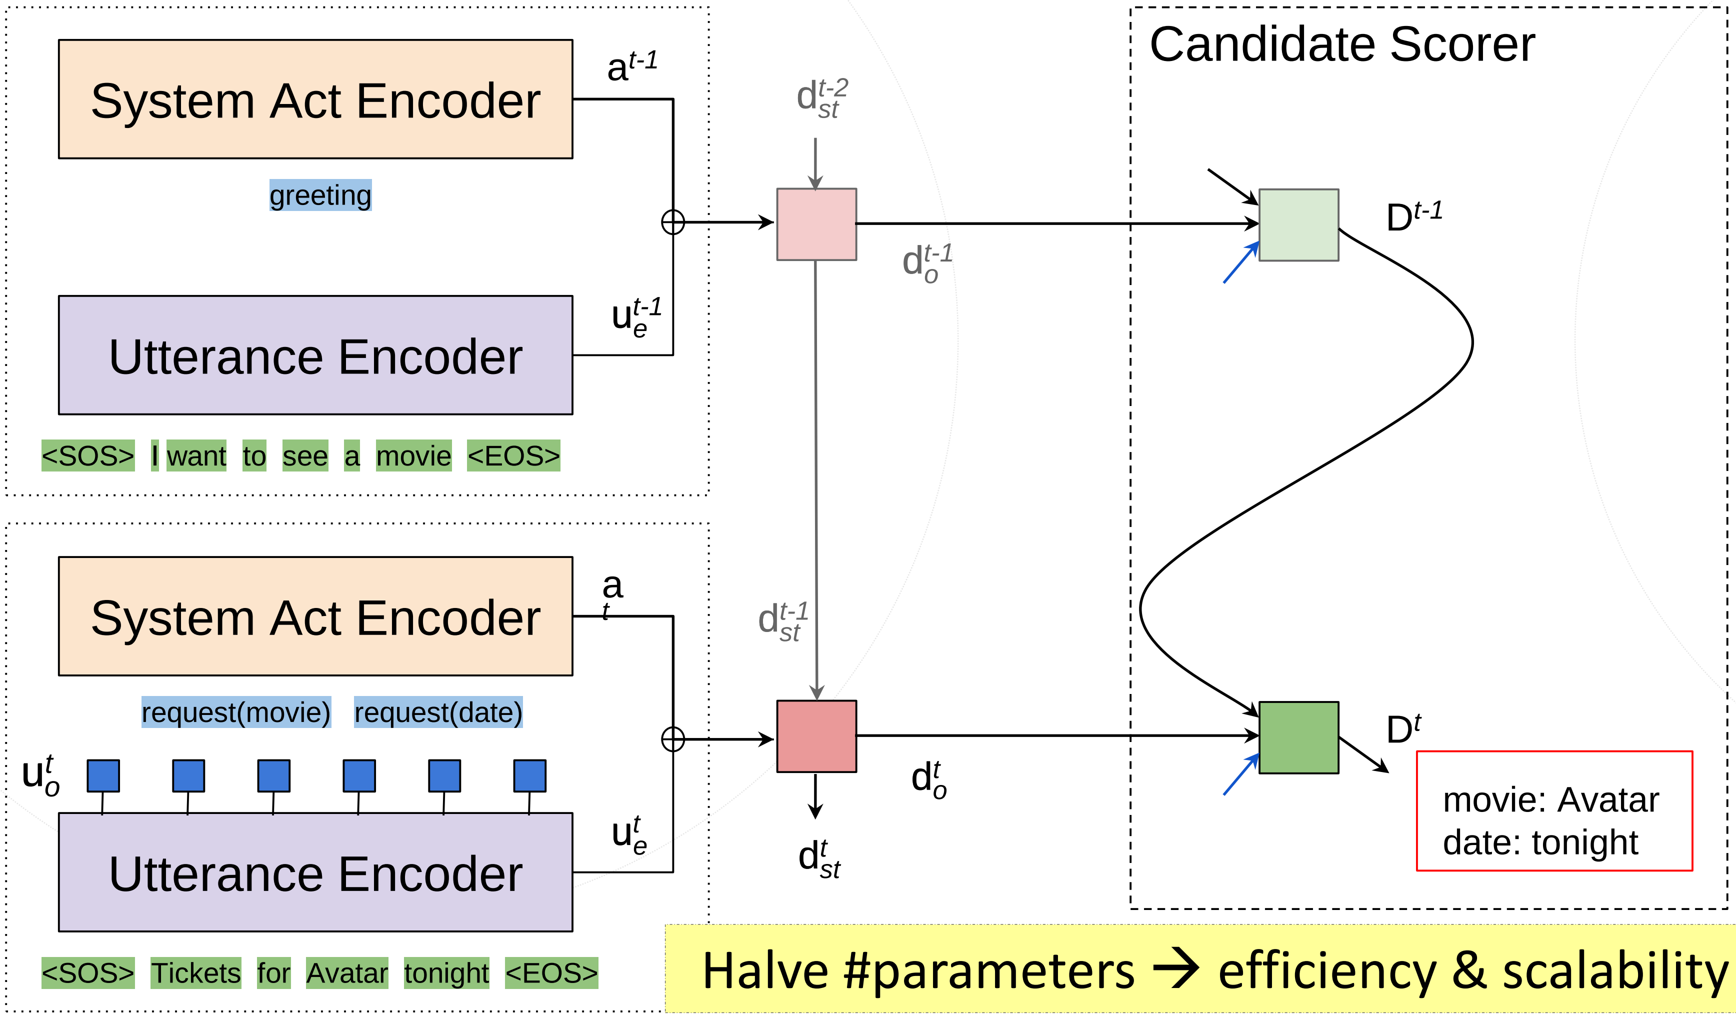

Halve #parameters → efficiency & scalability

Outline

- Introduction
- Background Knowledge
- **Modular Dialogue System**
  - Spoken/Natural Language Understanding (SLU/NLU)
  - **Dialogue Management**
    - Dialogue State Tracking (DST)
    - **Dialogue Policy Optimization**
  - Natural Language Generation (NLG)
- System Evaluation
- Recent Trends of Learning Dialogues

Material: http://deepdialogue.miulab.tw

Elements of Dialogue Management

Dialogue Policy Optimization

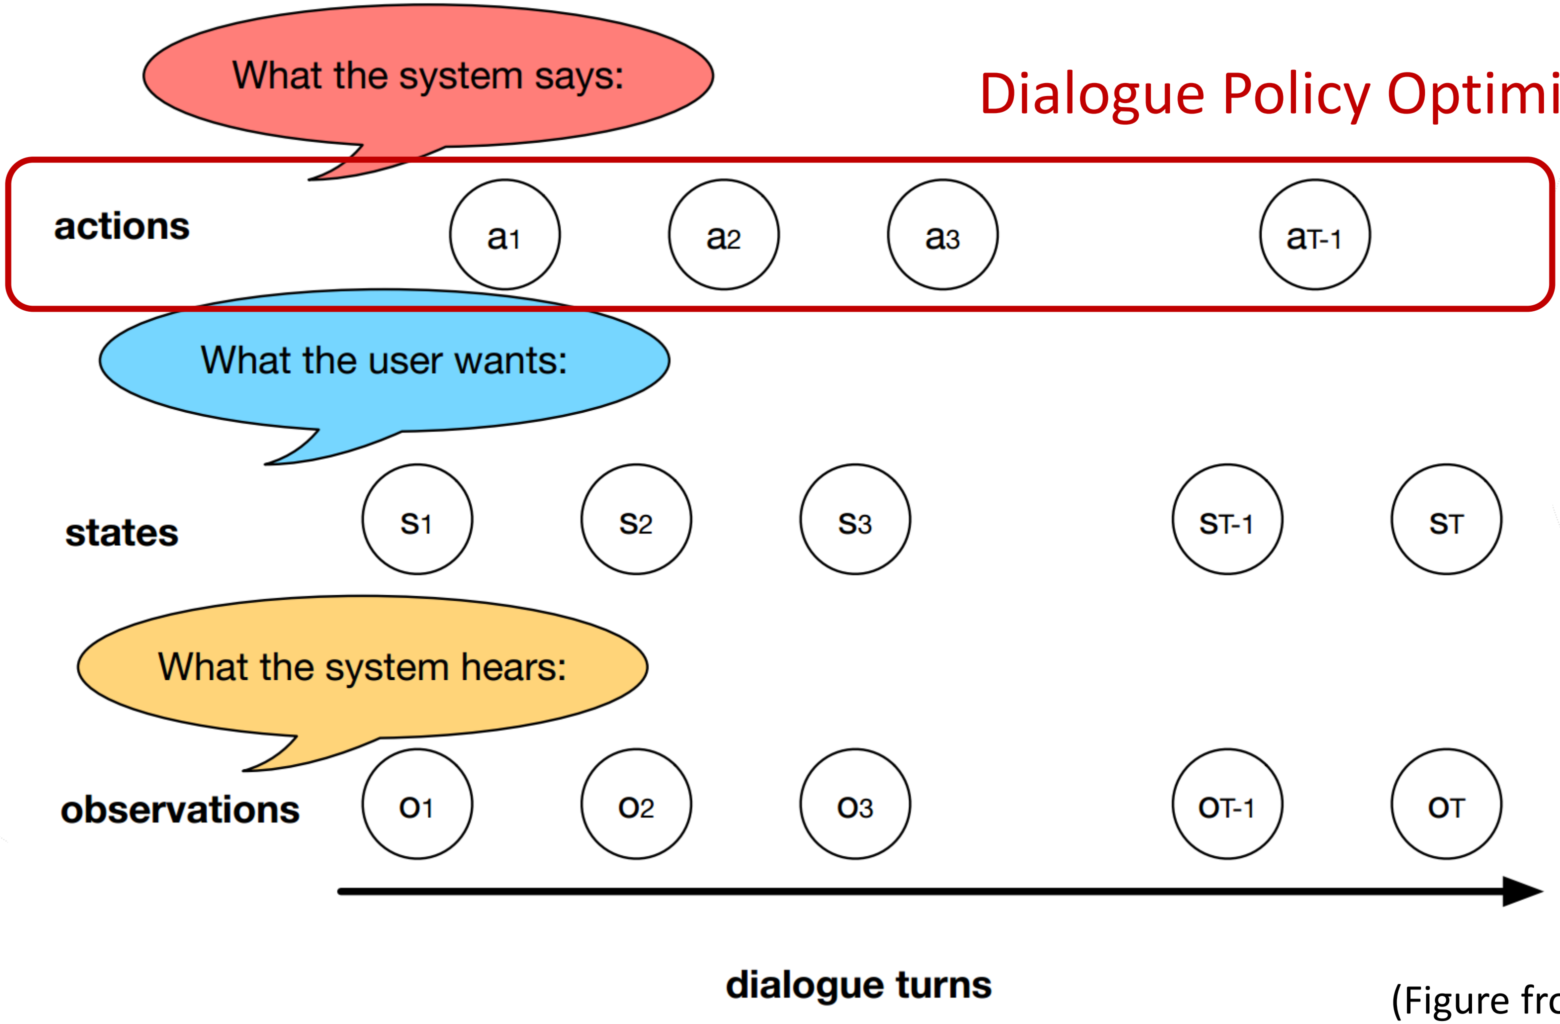

(Figure from Gašić)

Dialogue Policy Optimization

- Dialogue management in a RL framework

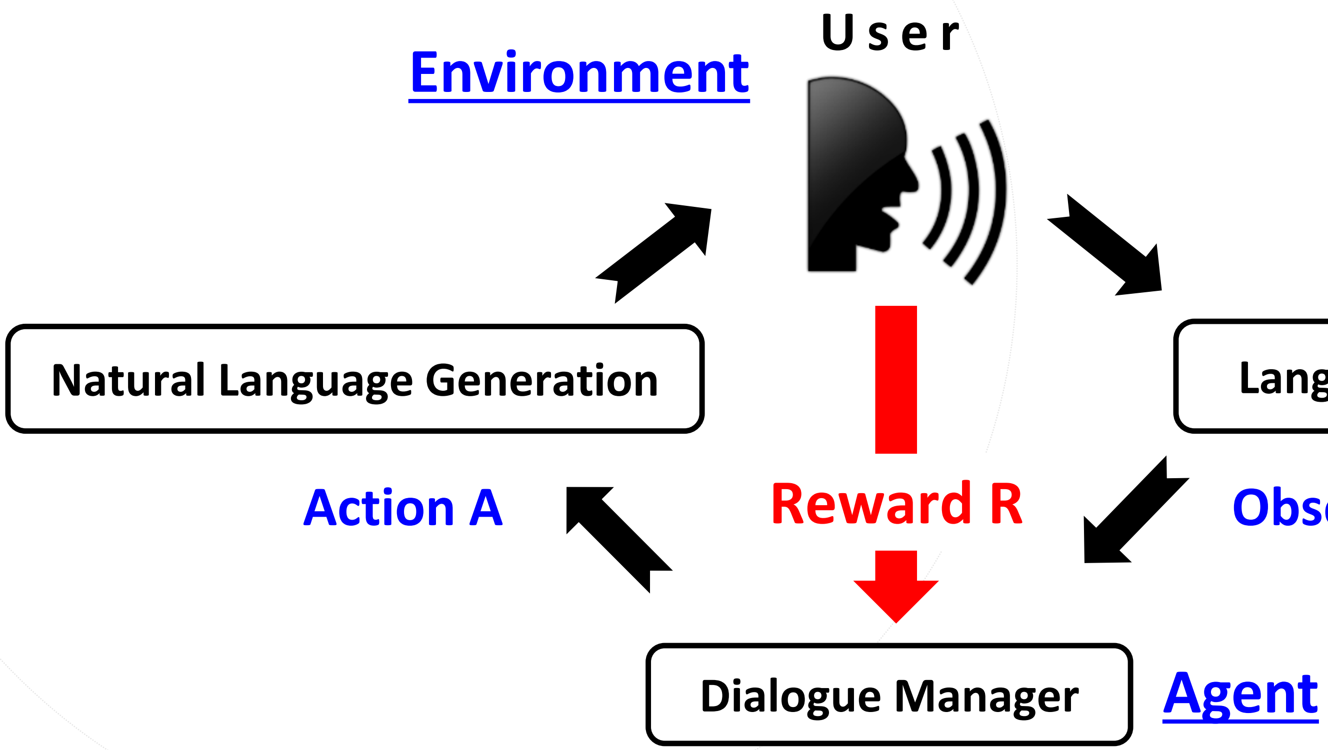

Optimized dialogue policy selects the best action that can maximize the future reward. Correct rewards are a crucial factor in dialogue policy training

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

RL for Dialogue Policy Optimization

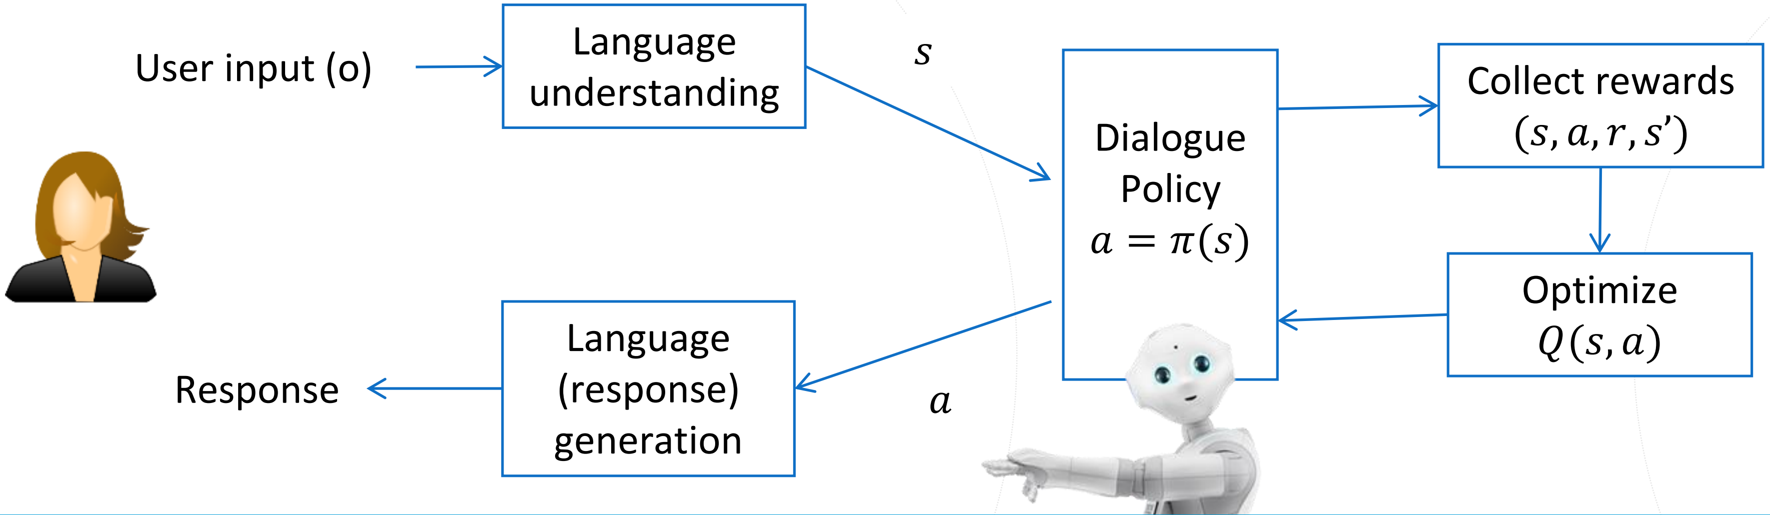

Type of Bots	State	Action	Reward
Social ChatBots	Chat history	System Response	# of turns maximized; Intrinsically motivated reward
InfoBots (interactive Q/A)	User current question + Context	Answers to current question	Relevance of answer; # of turns minimized
Task-Completion Bots	User current input +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

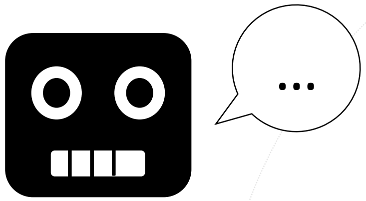

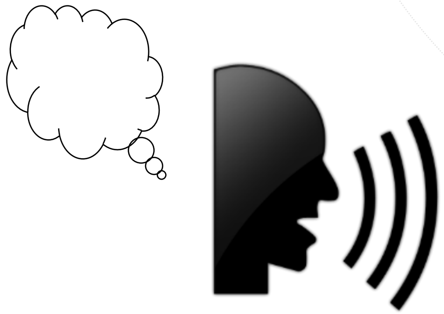

The user simulator is usually required for dialogue system training before deployment

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

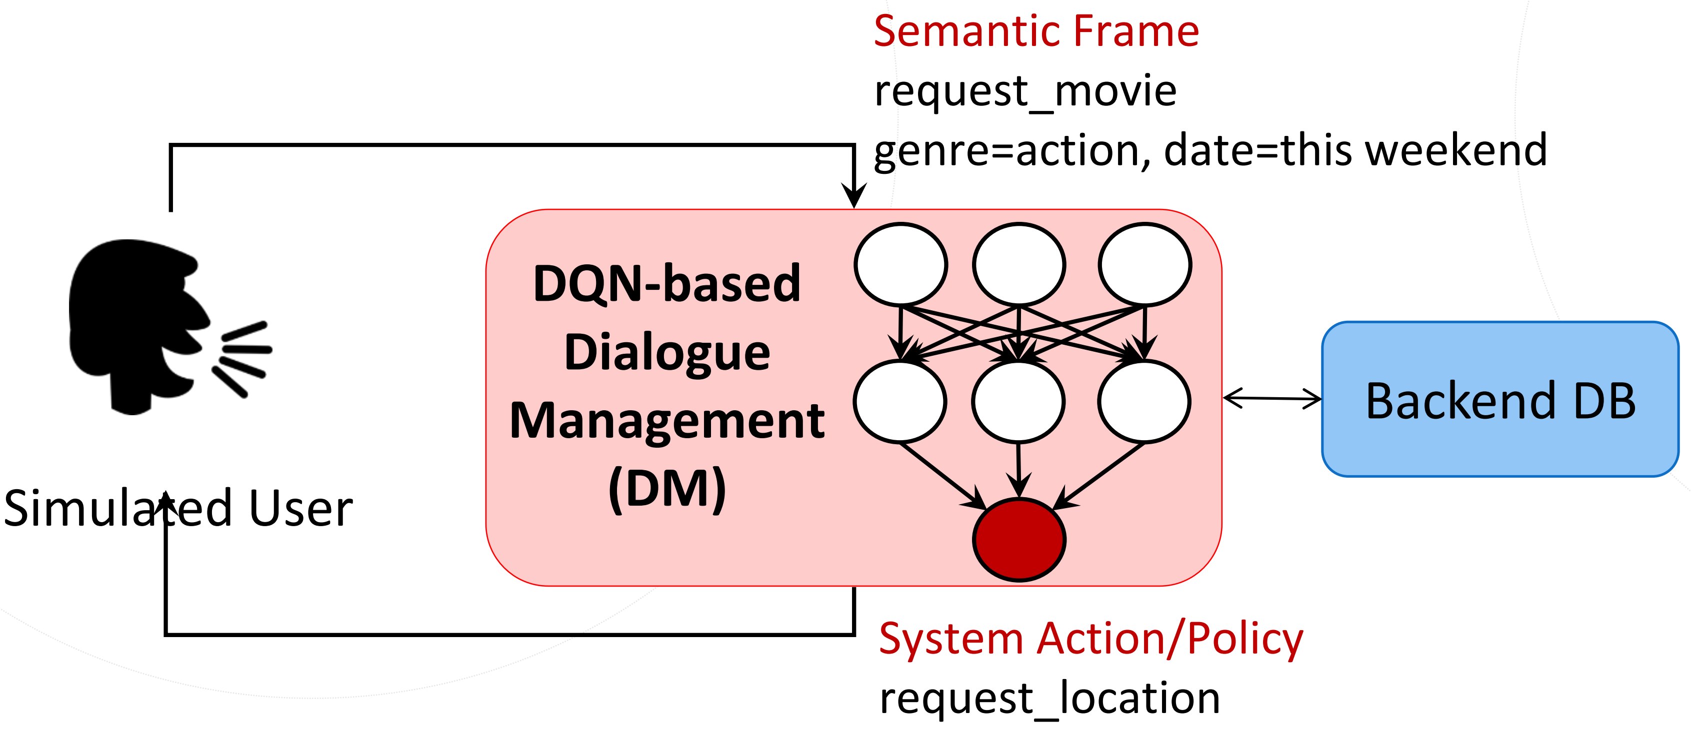

E2E Task-Completion Bot (TC-Bot) (Li+, 2017)

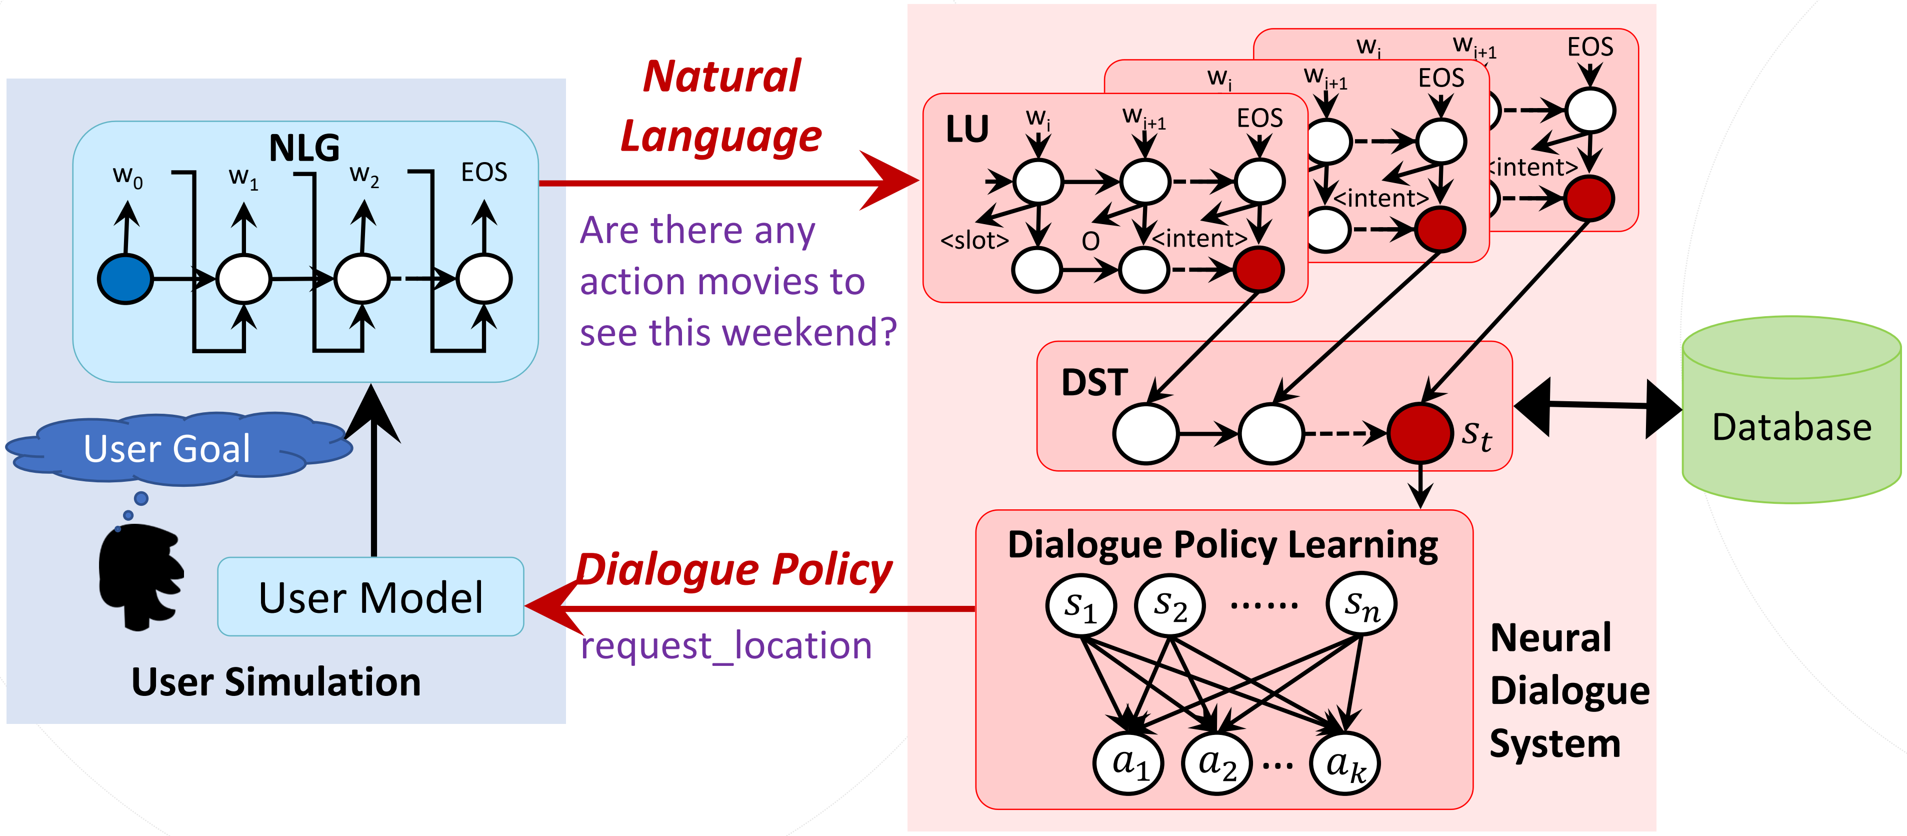

Idea: SL for each component and RL for end-to-end training

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

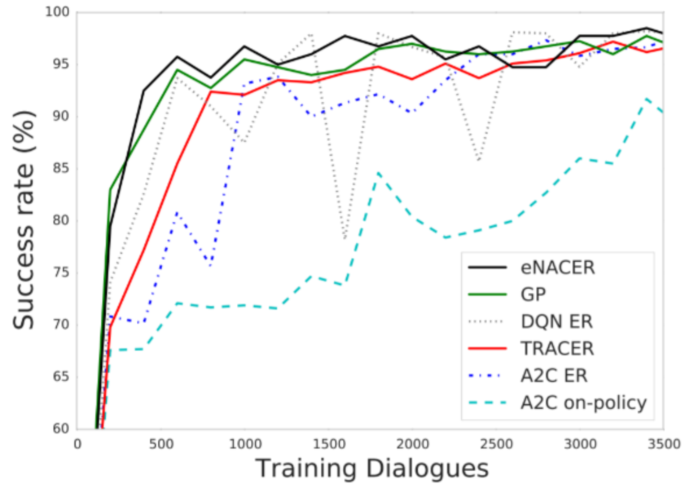

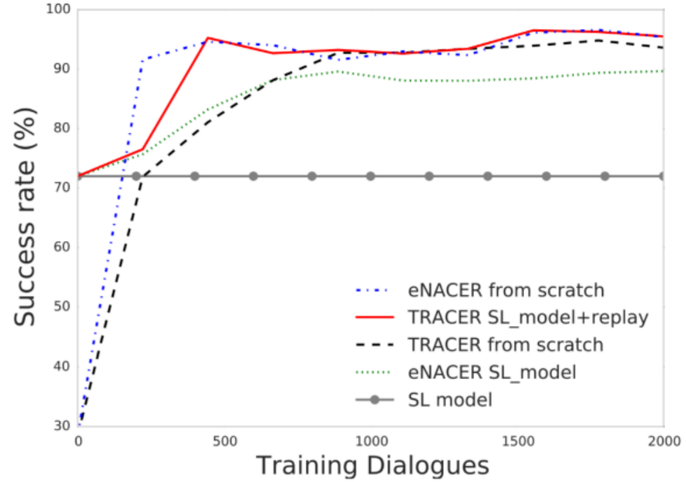

Material: http://deepdialogue.miulab.tw

Learning to Negotiate (Lewis+, 2017)

- Task: multi-issue bargaining
  - Each agent has its own value function

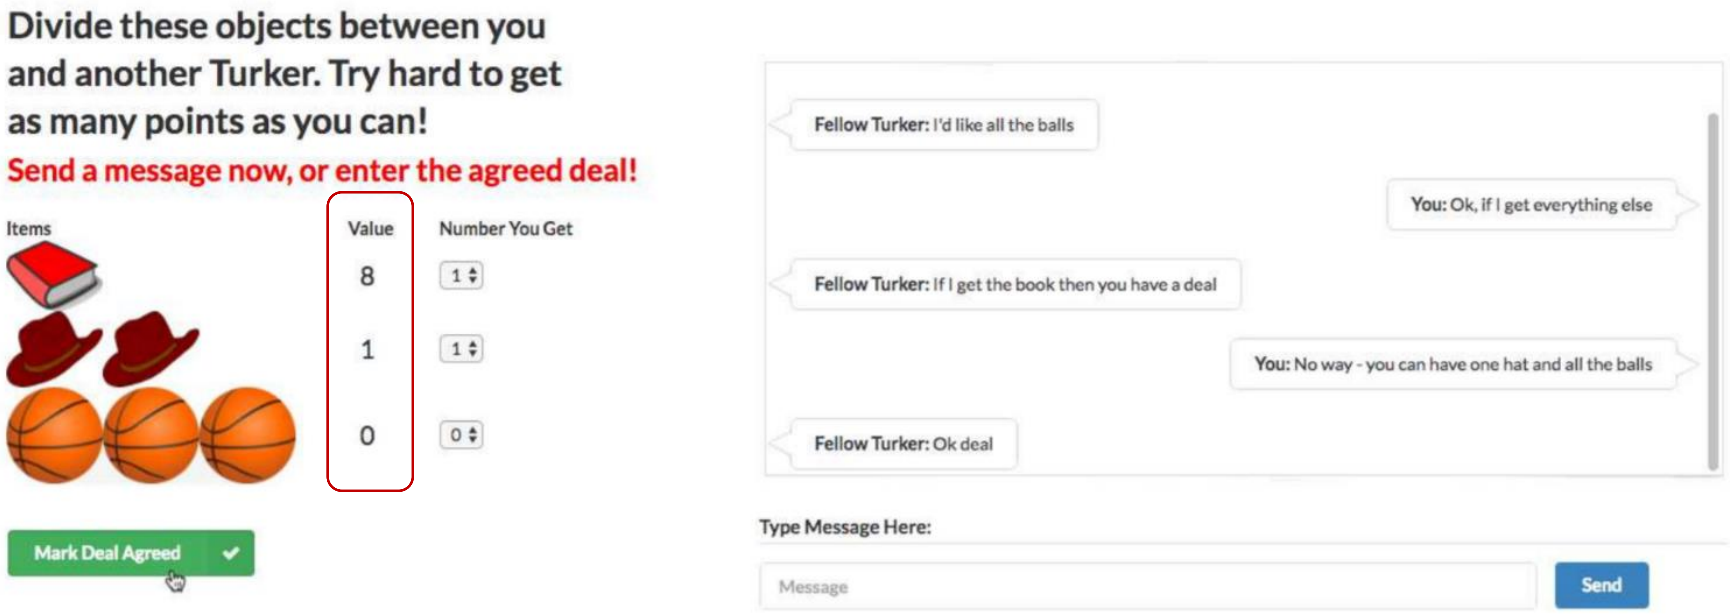

Learning to Negotiate (Lewis+, 2017)

- Dialogue rollouts to simulate a future conversation
- SL + RL
  - SL aims to imitate human users' actions
  - RL tries to make agents focus on the goal

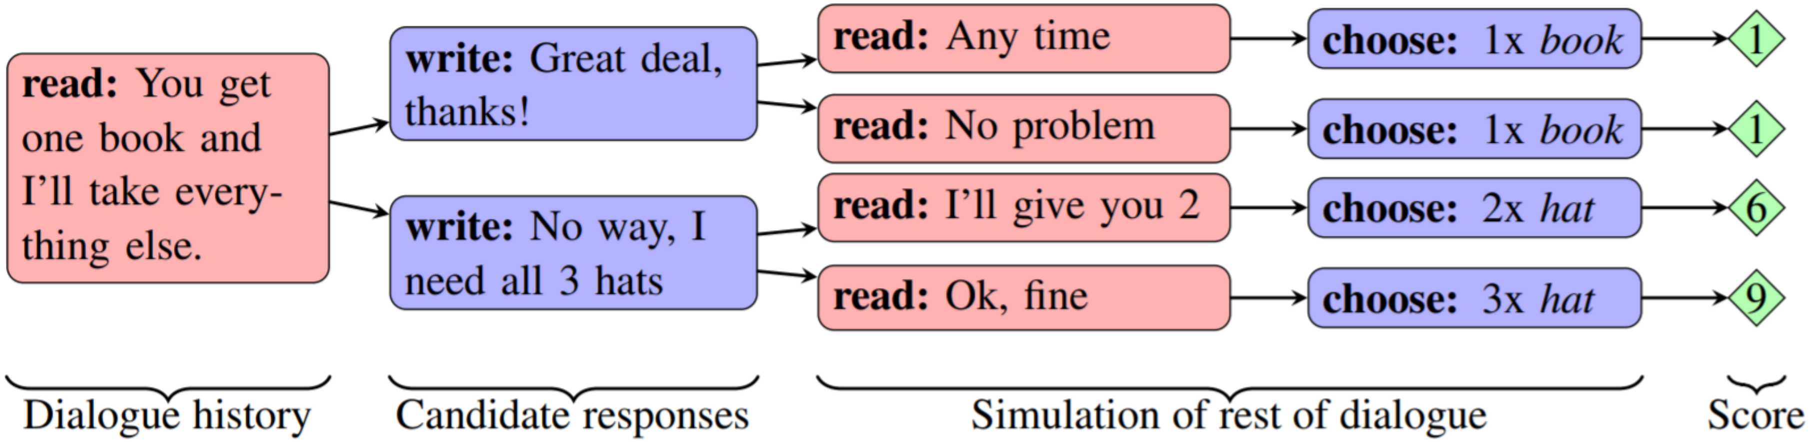

Online Training (Su+, 2015; Su+, 2016)

- 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

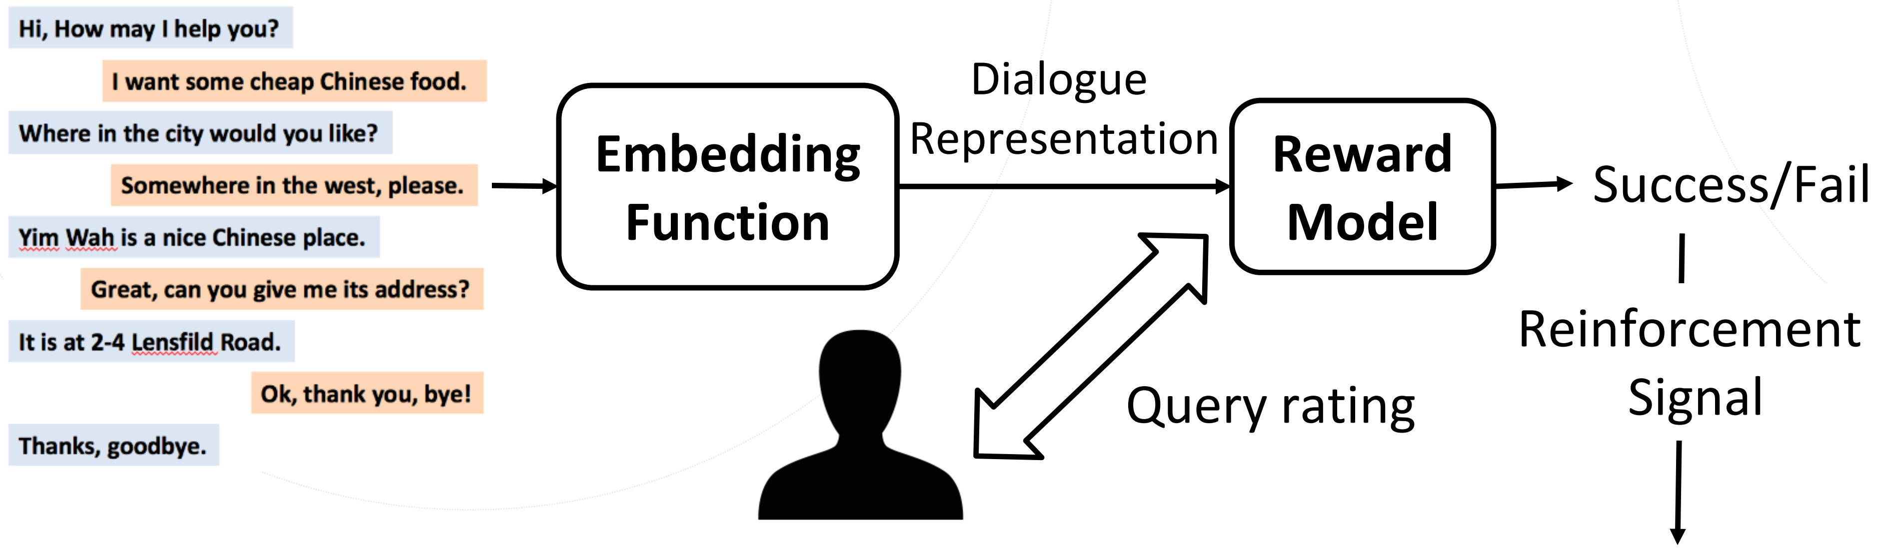

Material: http://deepdialogue.miulab.tw

Interactive RL for DM (Shah+, 2016)

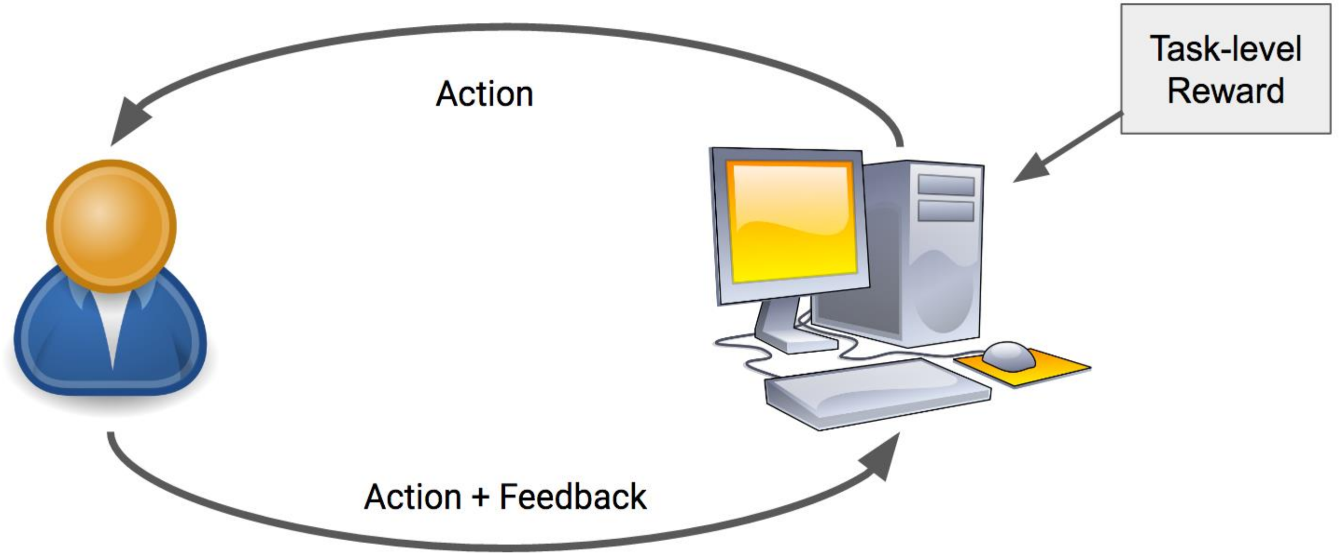

Immediate Feedback

Explicit

Implicit

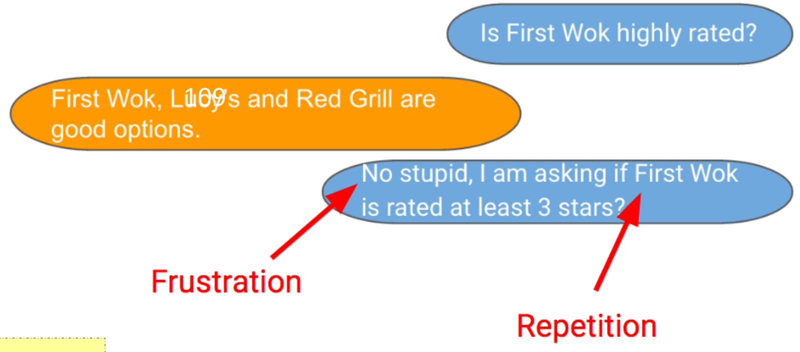

Use a third agent for providing interactive feedback to the DM

NTU MIULAB

Multi-Domain – Hierarchical RL (Peng+, 2017)

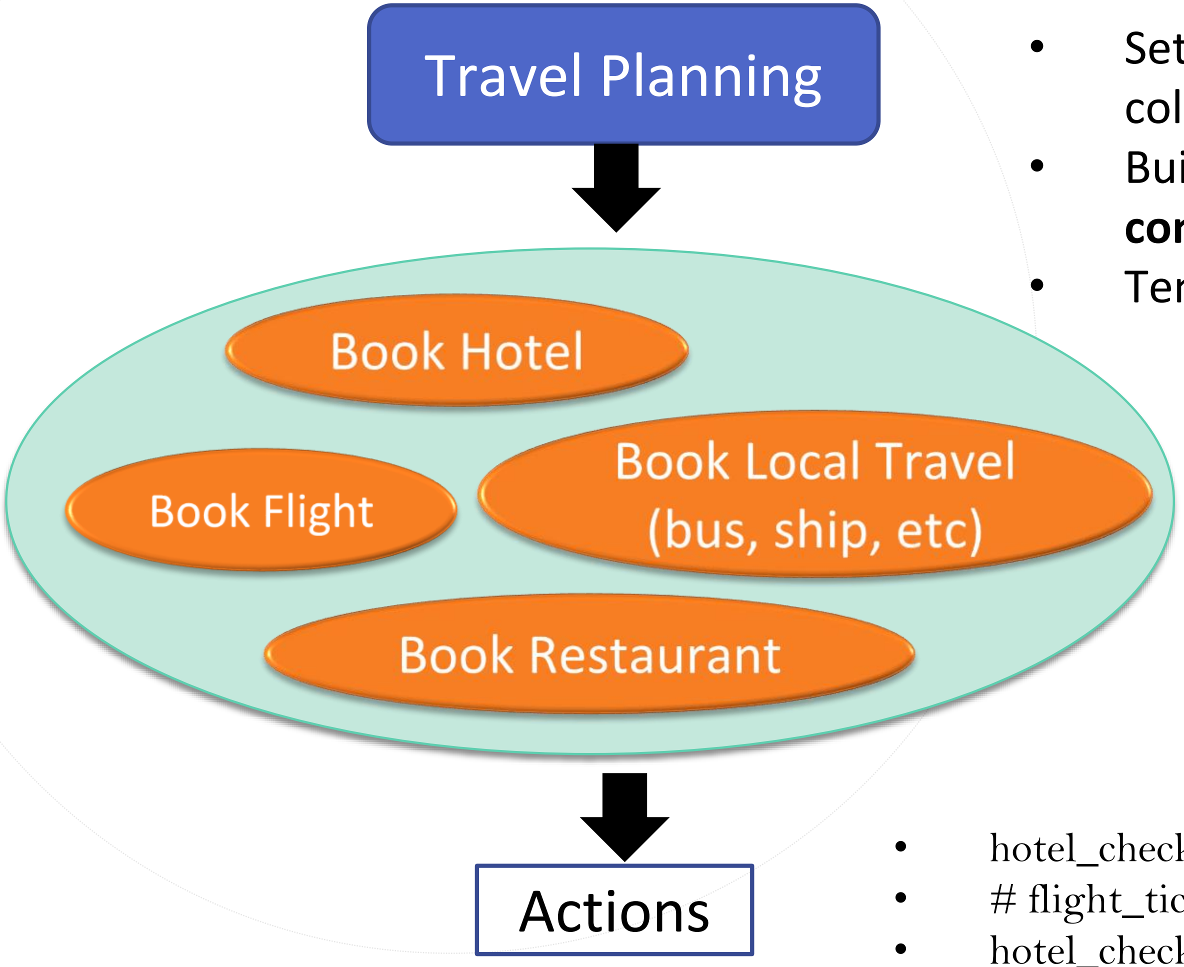

- Set of tasks that need to be fulfilled collectively!
- Build a DM for **cross-subtask constraints** (**slot constraints**)
- Temporally constructed goals

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

Multi-Domain – Hierarchical RL (Peng+, 2017)

- Model makes decisions over two levels: *meta-controller & 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

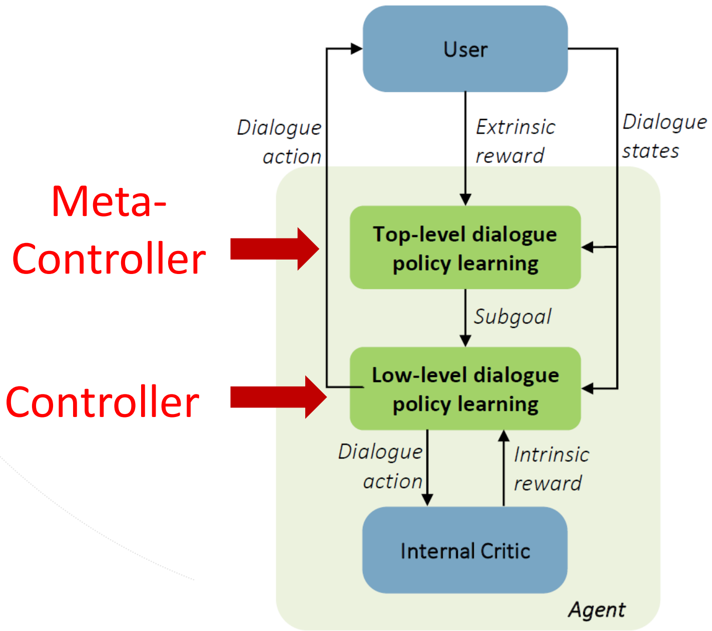

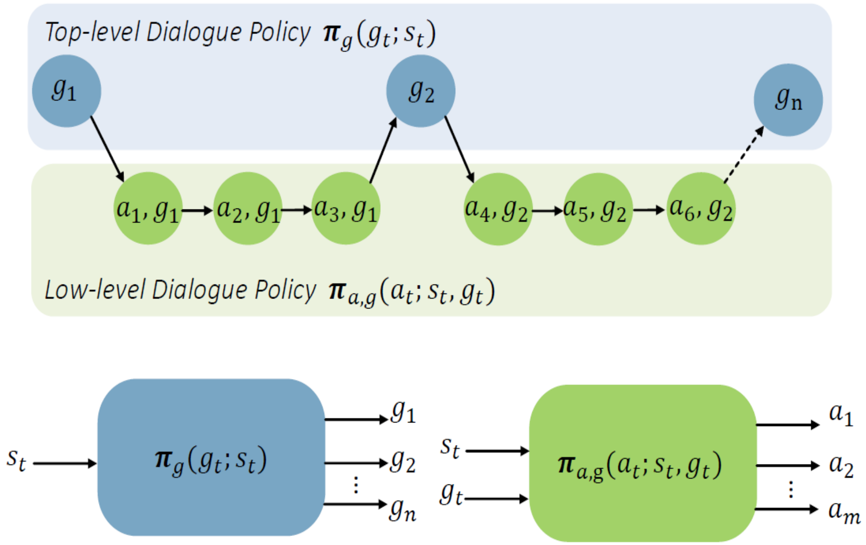

(mitigate reward sparsity issues)

Planning – Deep Dyna-Q (Peng+, 2018)

- Idea: learning with real users with planning

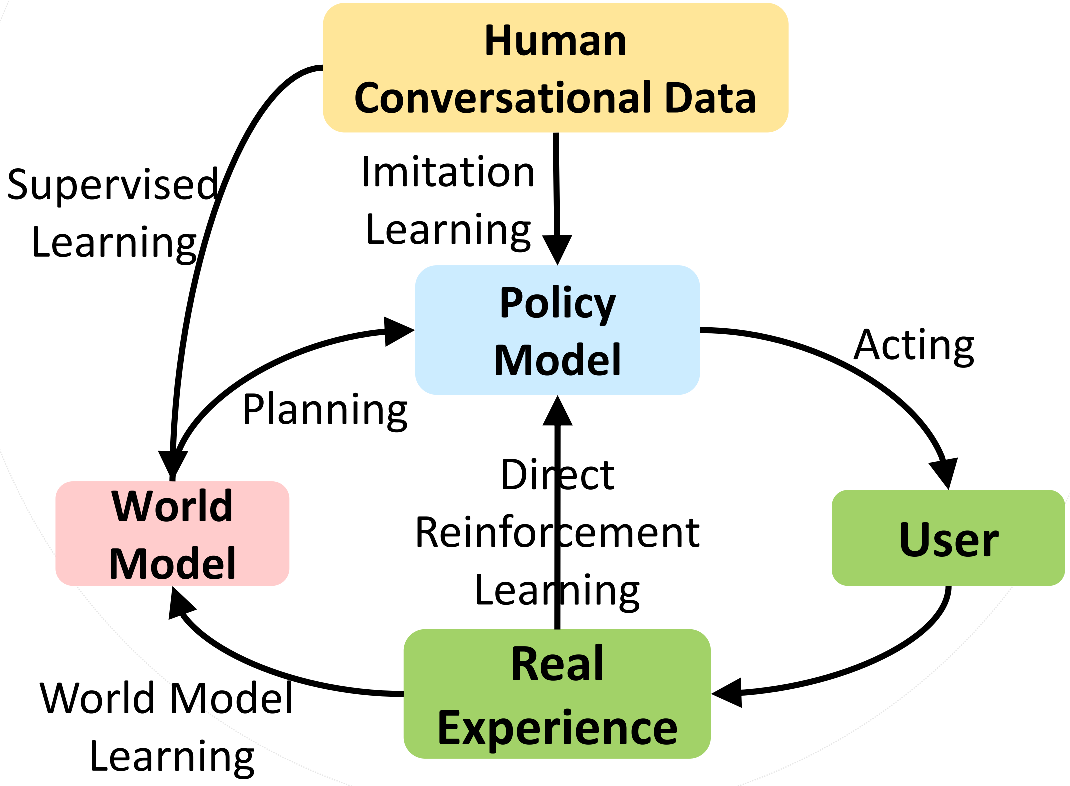

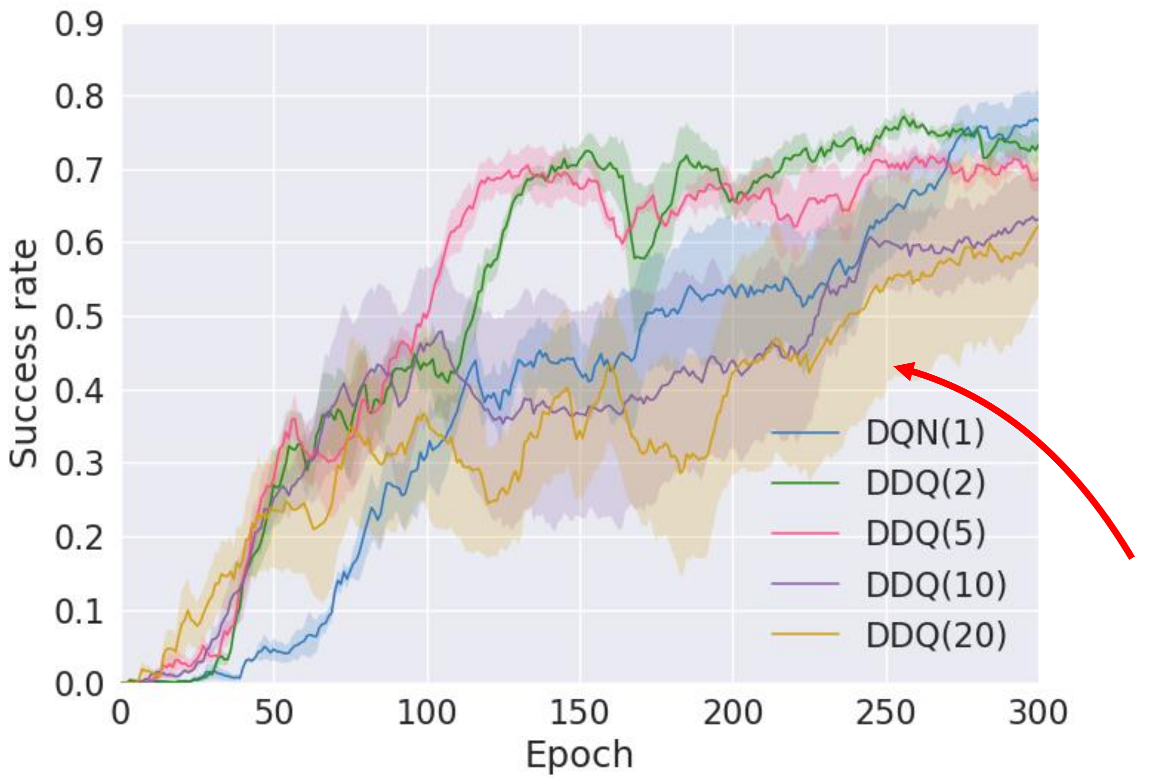

Policy learning suffers from the poor quality of fake experiences

Robust Planning – Discriminative Deep Dyna-Q (Su+, 2018)

- Idea: add a *discriminator* to filter out the bad experiences

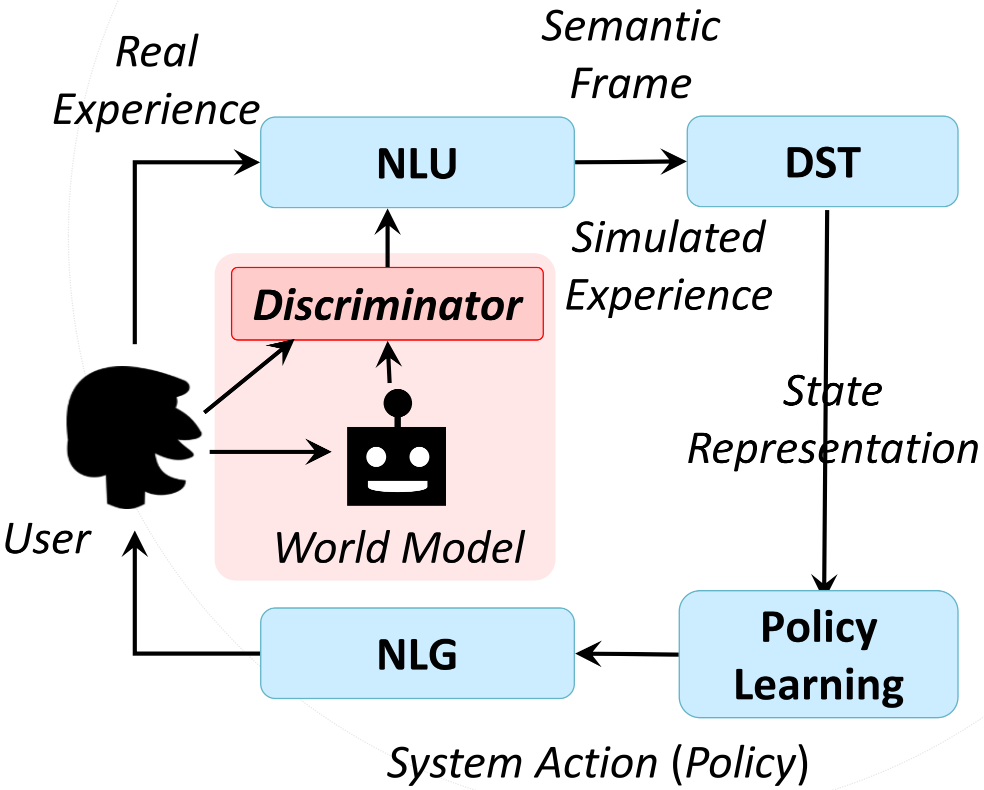

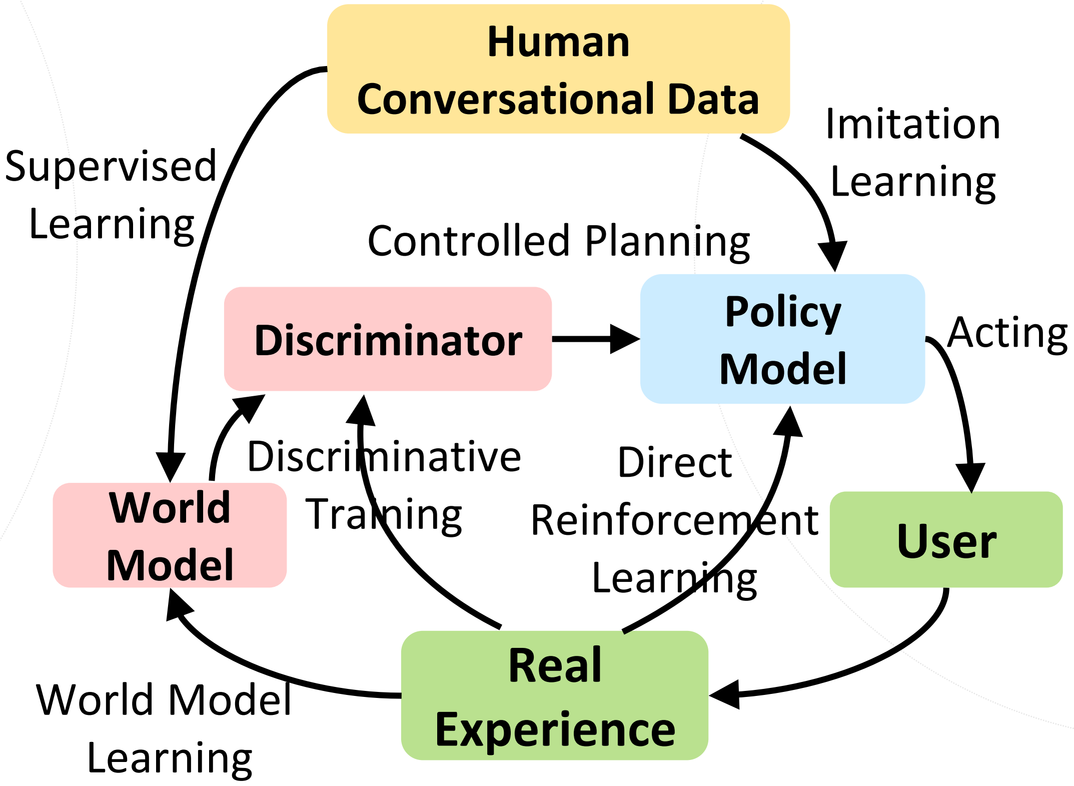

Robust Planning – Discriminative Deep Dyna-Q (Su+, 2018)

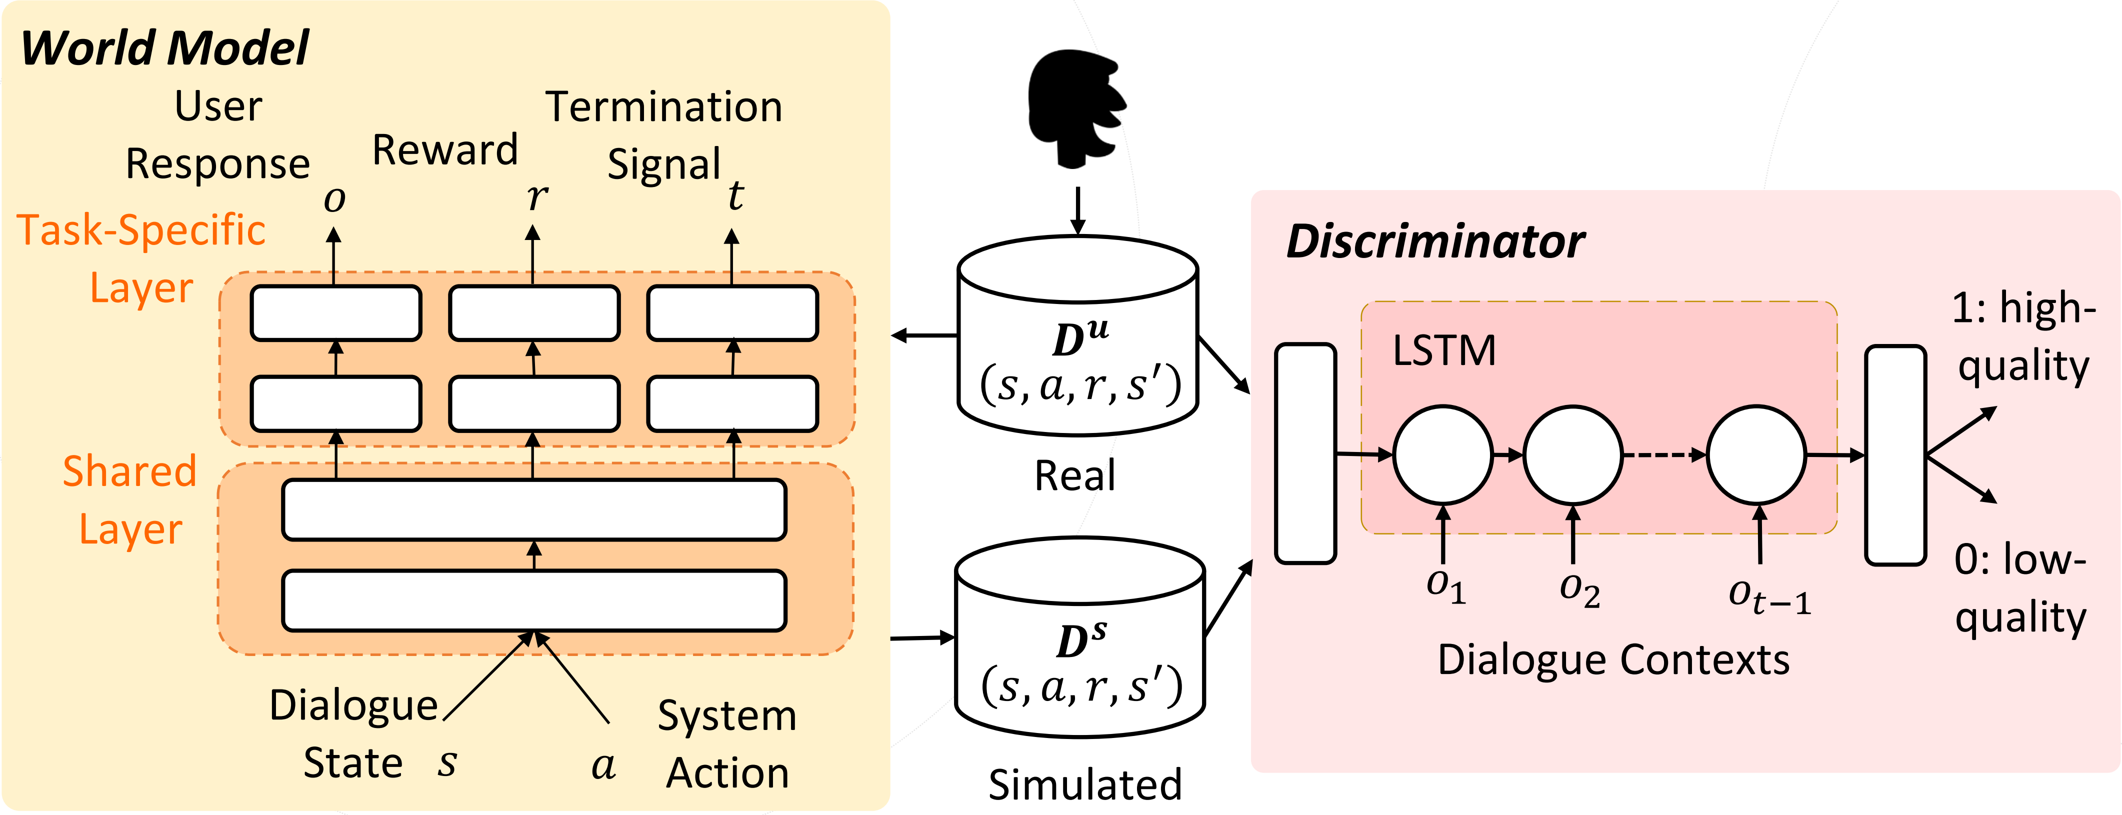

Robust Planning – Discriminative Deep Dyna-Q (Su+, 2018)

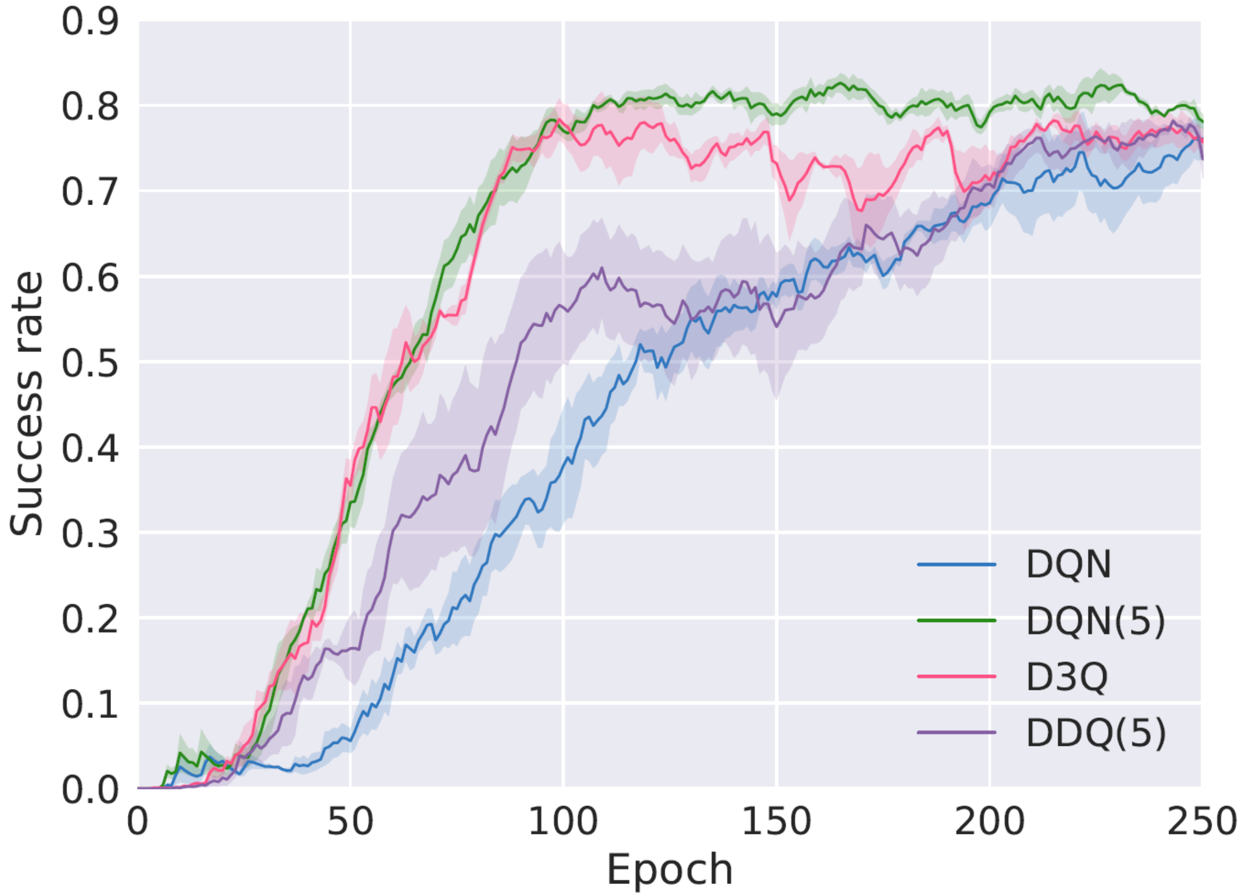

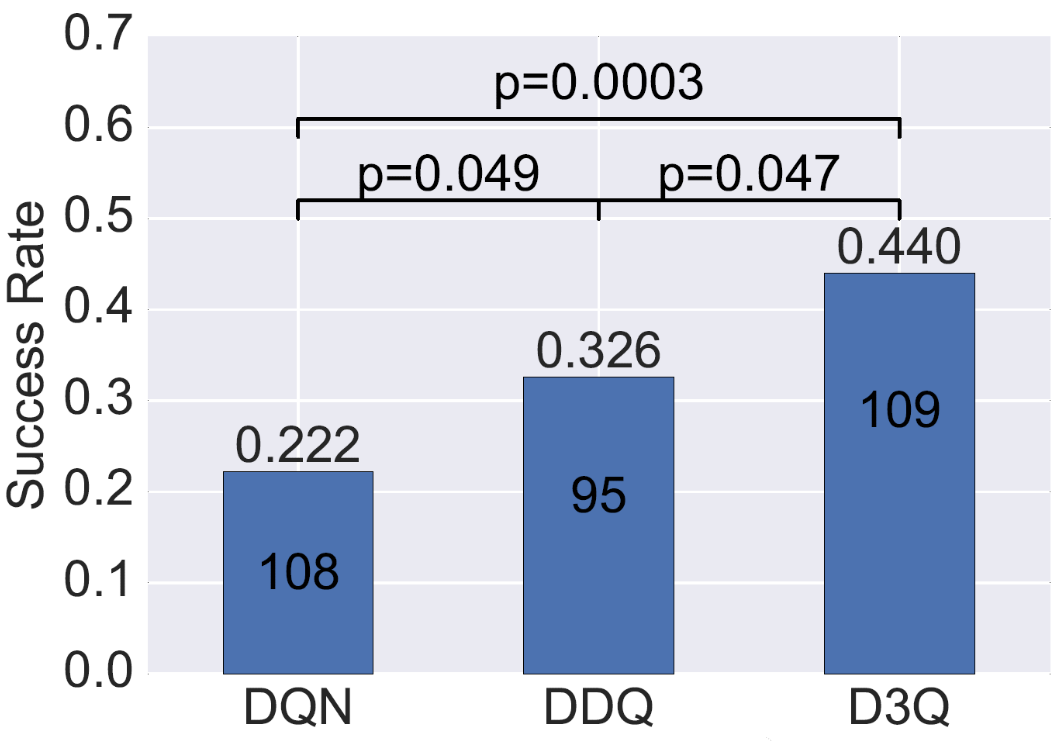

The policy learning is more robust and shows the improvement in human evaluation

Dialogue Management Evaluation

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

RL-Based DM Challenge

- SLT 2018 Microsoft Dialogue Challenge: End-to-End Task-Completion Dialogue Systems
  - Domain 1: Movie-ticket booking
  - Domain 2: Restaurant reservation
  - Domain 3: Taxi ordering

Outline

- Introduction
- Background Knowledge
- **Modular Dialogue System**
  - Spoken/Natural Language Understanding (SLU/NLU)
  - Dialogue Management
    - Dialogue State Tracking (DST)
    - Dialogue Policy Optimization
  - **Natural Language Generation (NLG)**
- System Evaluation
- Recent Trends of Learning Dialogues

Natural Language Generation (NLG)

- Mapping dialogue acts into natural language

inform(name=Seven_Days, foodtype=Chinese)

↓

Seven Days is a nice Chinese restaurant

Template-Based NLG

- Define a set of rules to map frames to NL

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

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

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

- Divide the problem into pipeline

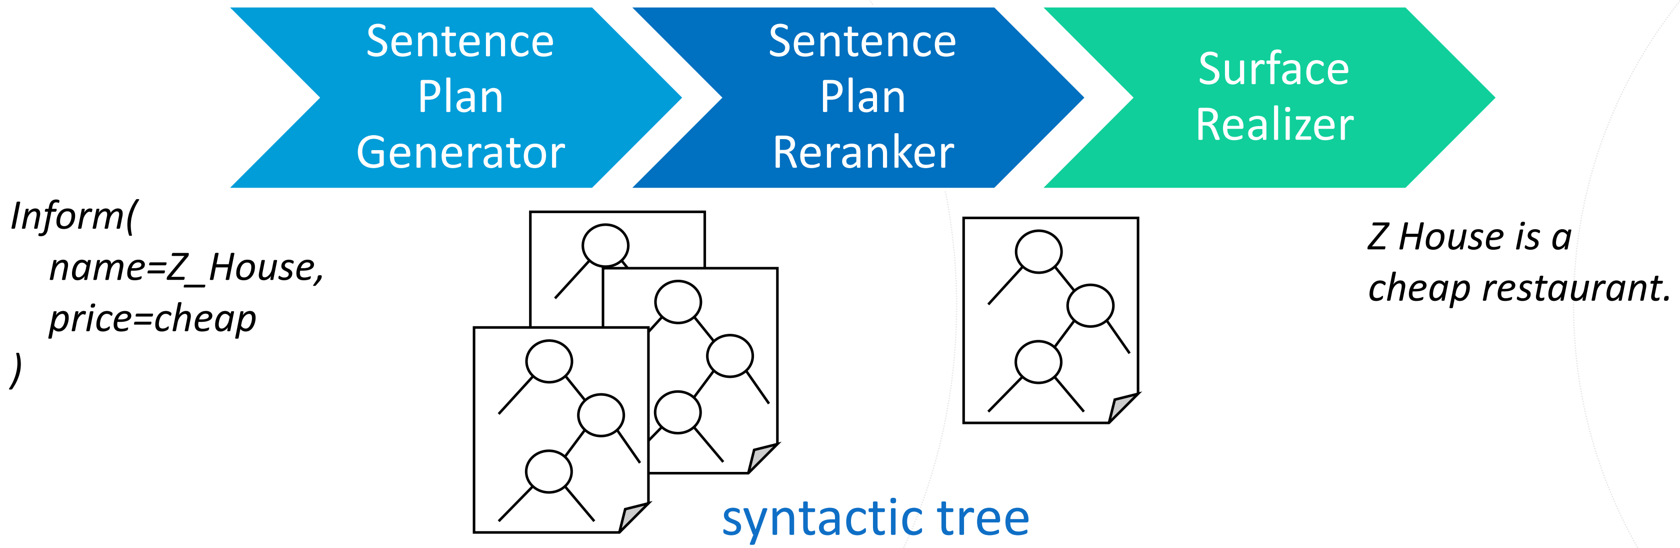

- Statistical sentence plan generator (Stent et al., 2009)
- Statistical surface realizer (Dethlefs et al., 2013; Cuayáhuitl et al., 2014; …)

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

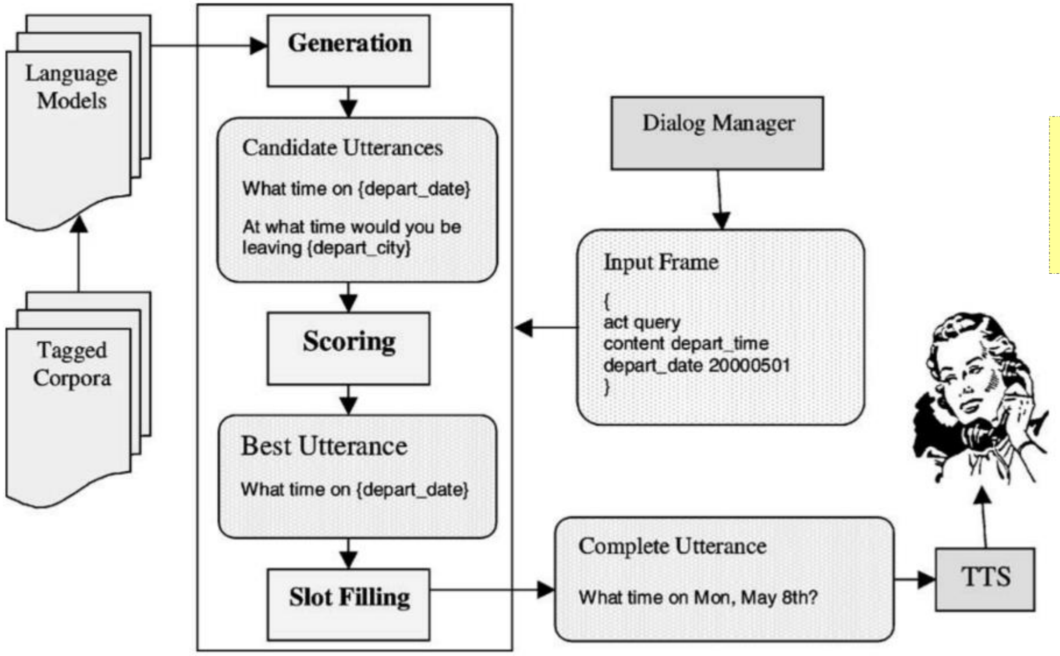

Pros: easy to implement/ understand, simple rules
Cons: computationally inefficient

Material: http://deepdialogue.miulab.tw

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

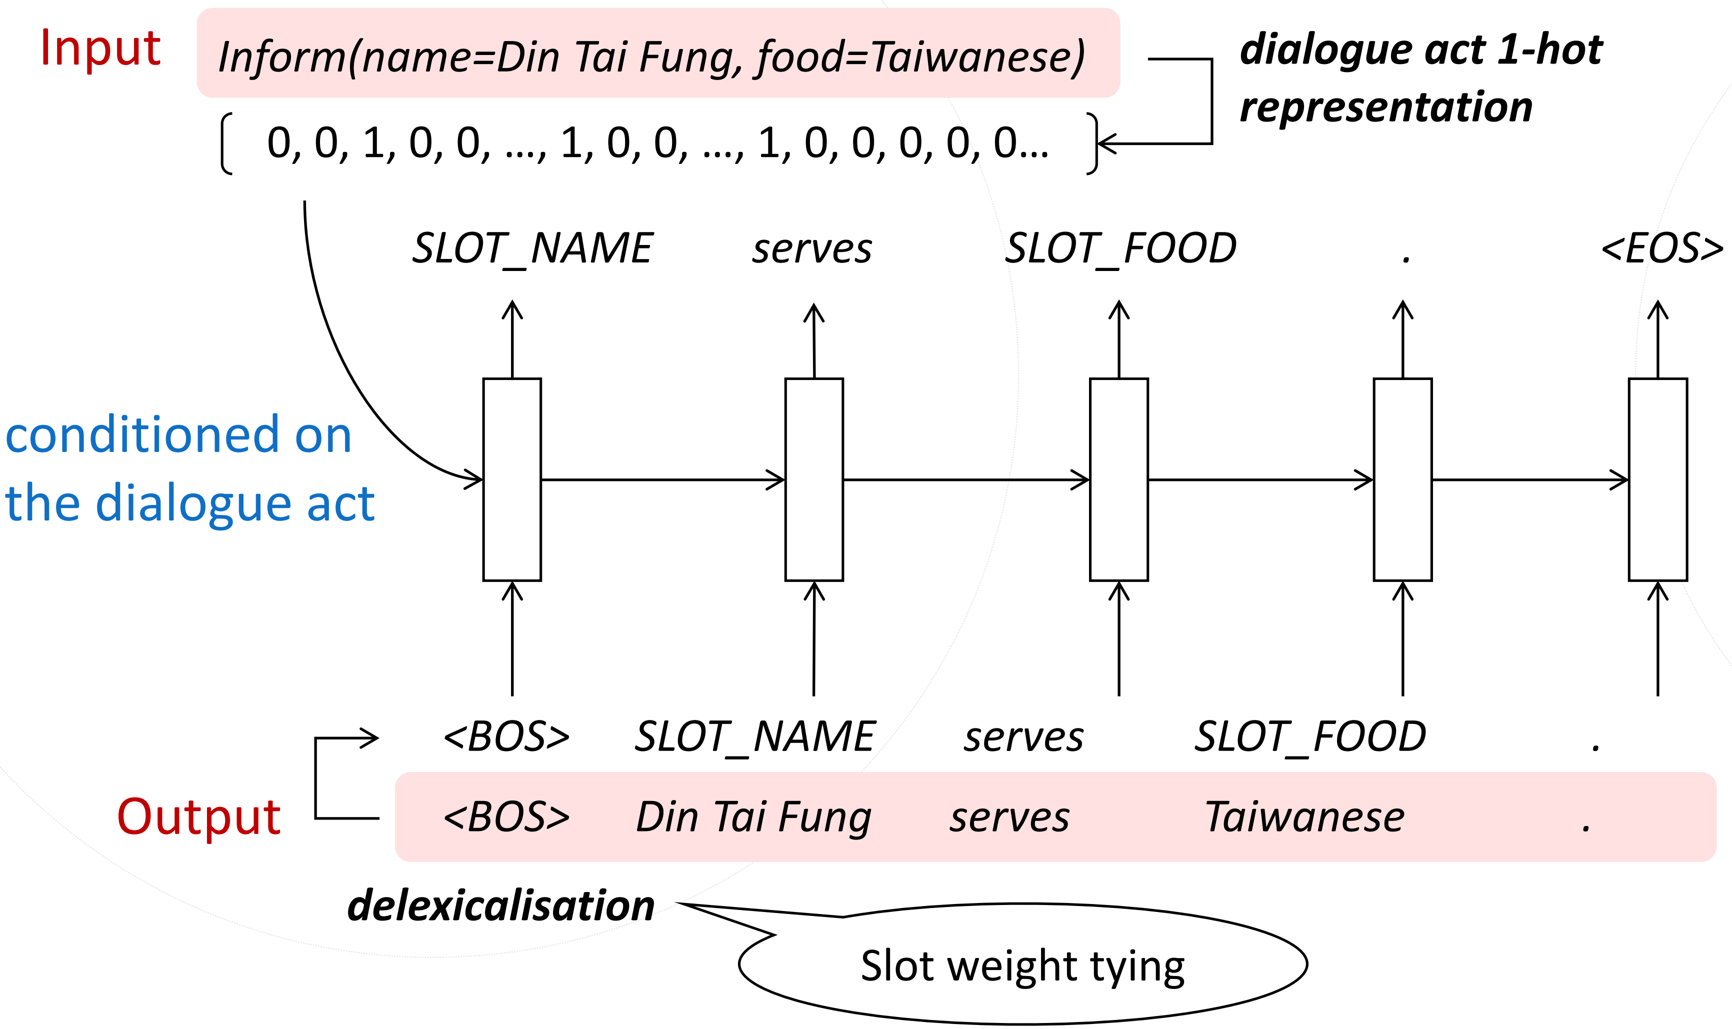

Handling Semantic Repetition

- 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

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

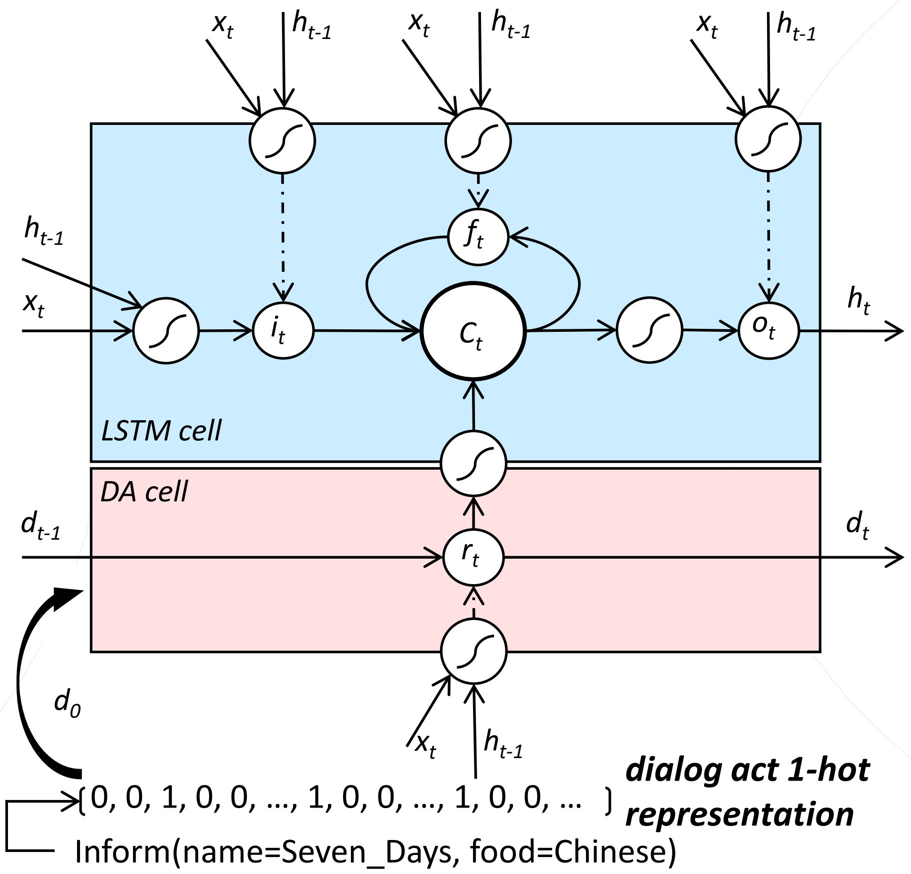

Idea: using gate mechanism to control the generated semantics (dialogue act/slots)

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)

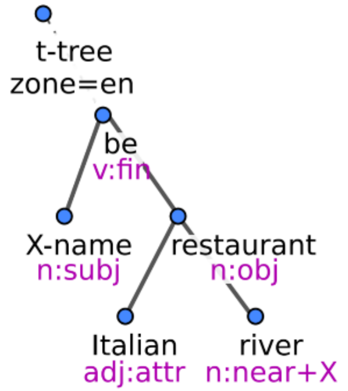

(<root> <root> ((X-name n:subj) be v:fin ((Italian adj:attr) restaurant n:obj (river n:near+X))))

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.

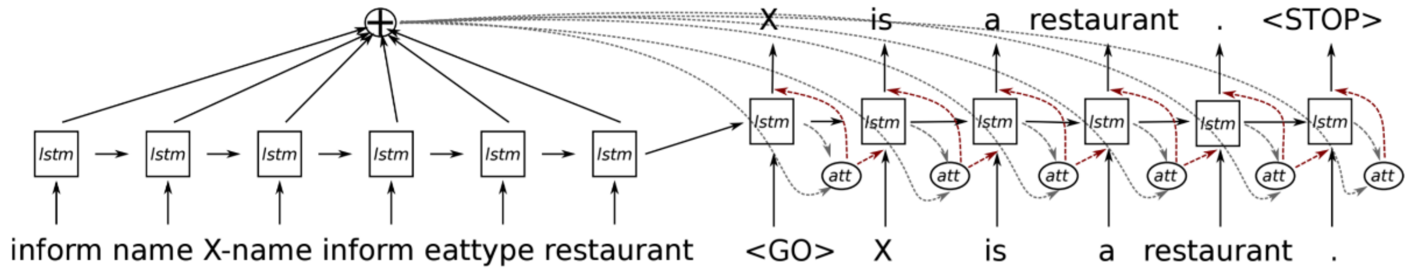

Structural NLG (Sharma+, 2017; Nayak+, 2017)

- Delexicalized slots do not consider the word level information

Generated output:	There are no restaurants around which serve INFORM-FOOD food.			
Delexicalized slot input:	INFORM-FOOD	✓	INFORM-FOOD	✗
Lexicalized value input:	chinese		pizza	

- Slot value-informed sequence to sequence models

Mention rep.	Input sequence					
SEQ	x_i	x_{i+1}	x_{i+2}	x_{i+3}	x_{i+4}	...
	decor	*decent*	service	good	cuisine	...
JOINT	x_i		x_{i+1}		x_{i+2}	
	⟨ decor, decent ⟩		⟨ service, good ⟩		⟨ cuisine, null ⟩	
CONCAT	$x_{i,1}$	$x_{i,2}$	$x_{i+1,1}$	$x_{i+1,2}$	$x_{i+2,1}$	$x_{i+2,2}$
	decor	decent	service	good	cuisine	null

Structural NLG (Nayak+, 2017)

- Sentence plans as part of the input sequence

Plan sup.	Input tokens					
NONE	decor	decent	service	decent	quality	good
FLAT	decor	decent	service	decent		
	quality	good				
POSITIONAL	<B>	decor	decent	service	decent	
	<I>	quality	good			

Contextual NLG (Dušek and Jurčíček, 2016)

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

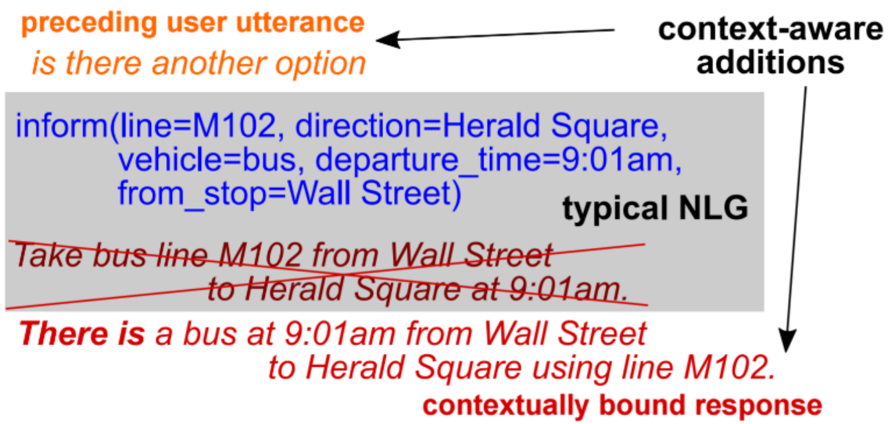

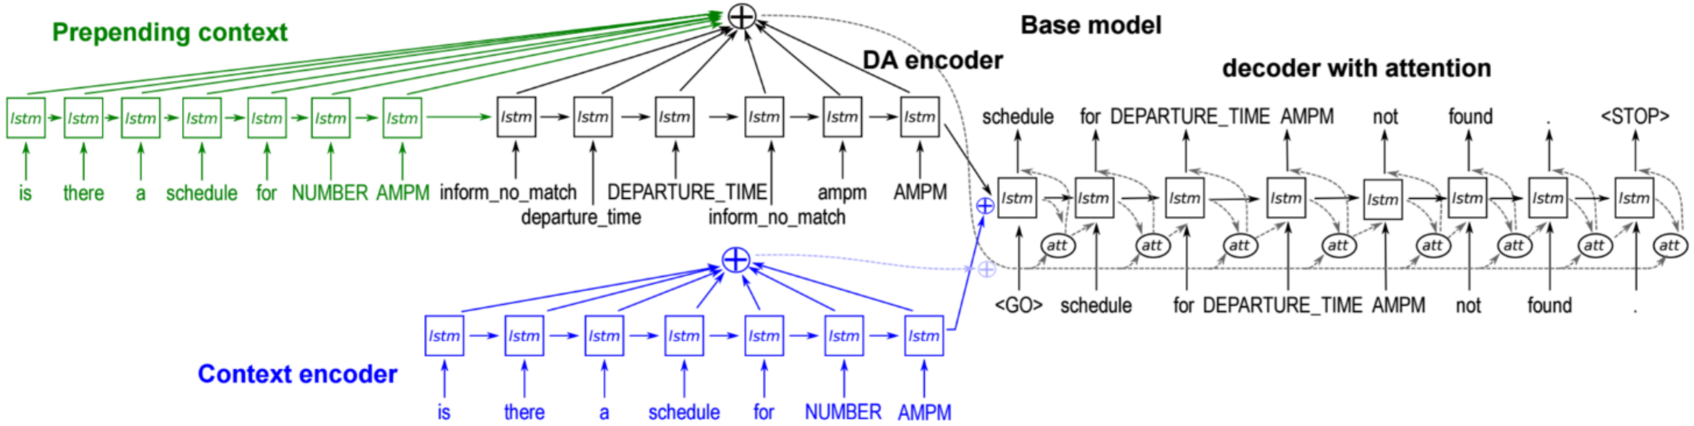

Controlled Text Generation (Hu et al., 2017)

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

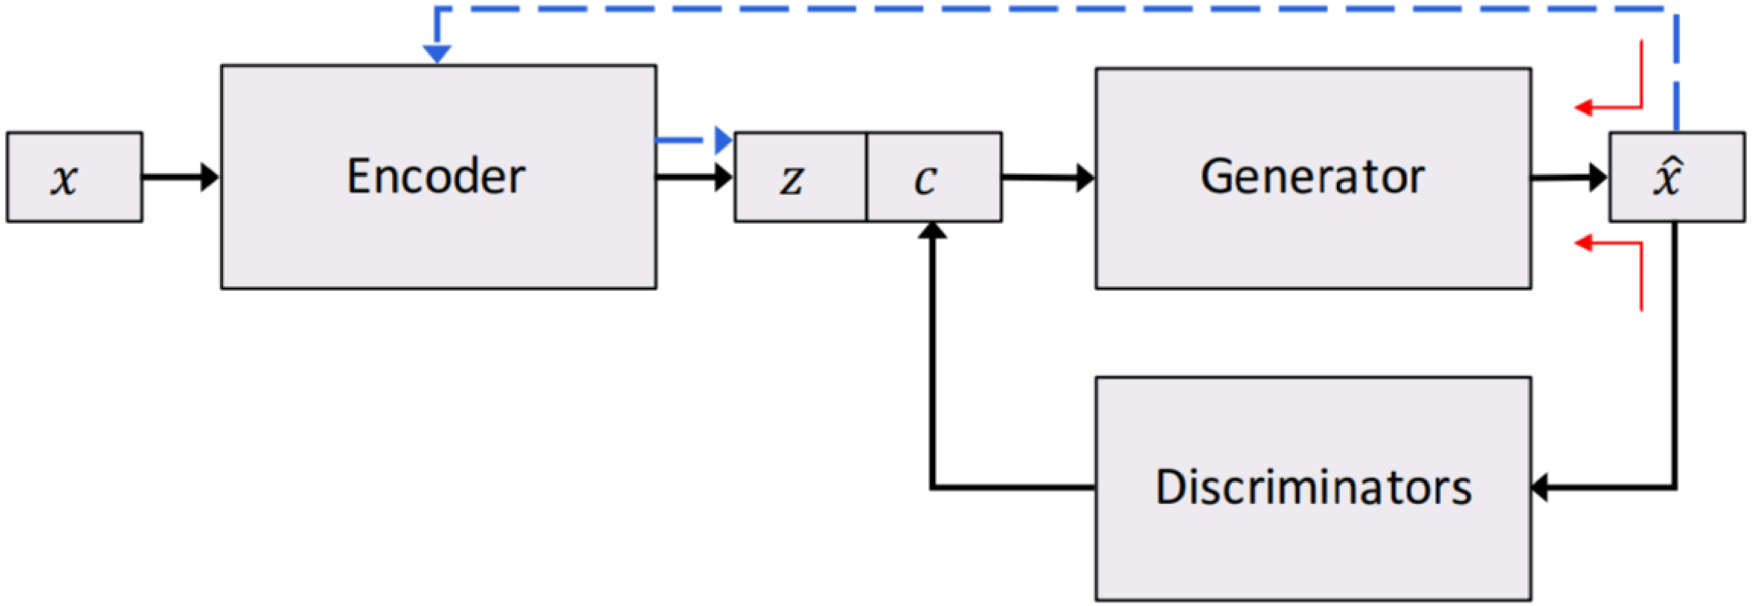

Knowledge-Grounded Conversations (Ghazvininejad+, 2017)

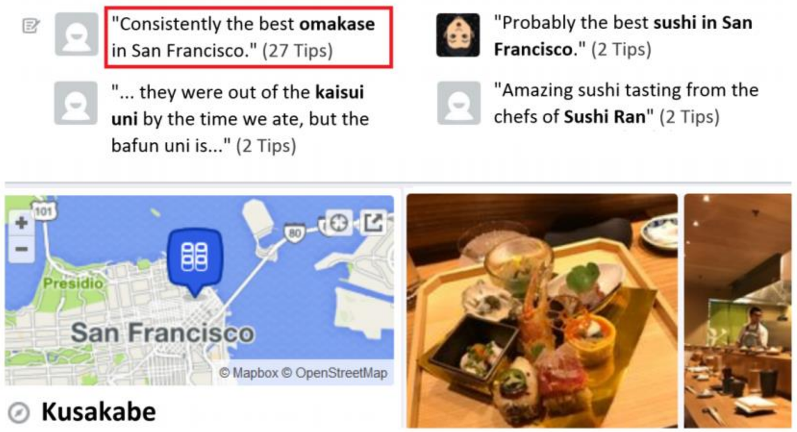

User input: *Going to Kusakabe tonight.*
Neural model: Have a great time!
Human: You'll love it! Try omasake, the best in town.

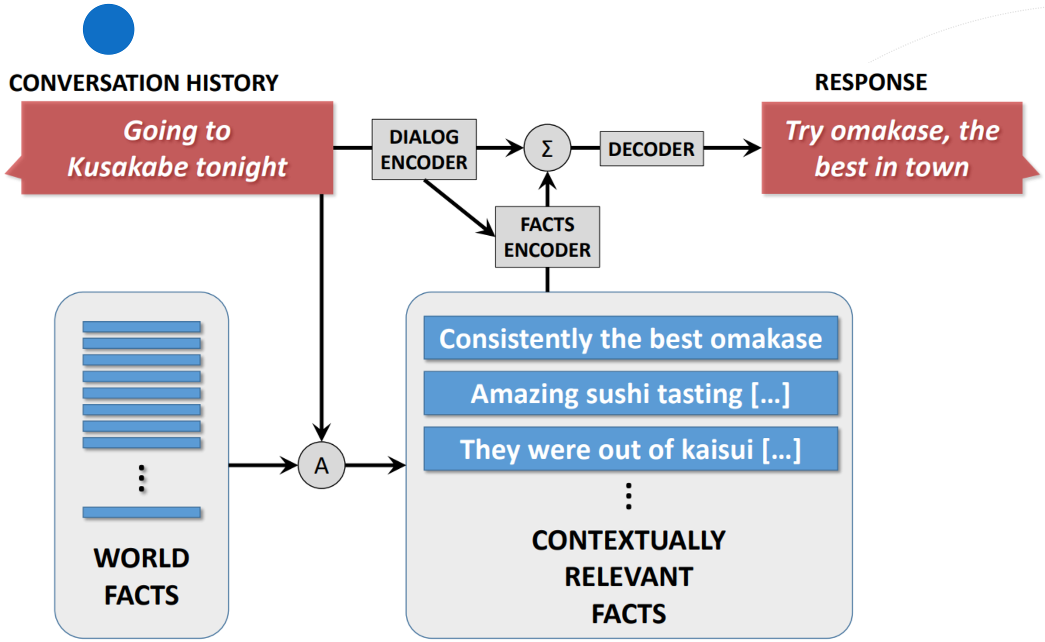

A: **Looking forward to trying** *@pizzalibretto* **tonight! my expectations are high**.
B: **Get** *the rocco salad*. **Can you eat** *calamari*?

A: **Anyone in** *Chi* **have a** *dentist office* **they recommend? I'm never going back to** *[...]* **and would love a reco!**
B: **Really looved** *Ora* **in** *Wicker Park*.

A: **I'm at** *California Academy of Sciences*
B: **Make sure you catch** *the show* **at** *the Planetarium*. **Tickets are usually limited.**

A: **I'm at** *New Wave Cafe*.
B: **Try to get to** *Dmitri's* **for dinner**. **Their** *pan fried scallops* **and** *shrimp scampi* **are to die for.**

A: **I just bought:** *[...] 4.3-inch portable GPS navigator* **for my wife, shh, don't tell her**.
B: **I heard this brand** *loses battery power*.

NTU MIULAB

Material: http://deepdialogue.miulab.tw

Hierarchical NLG w/ Linguistic Patterns (Su+, 2018)

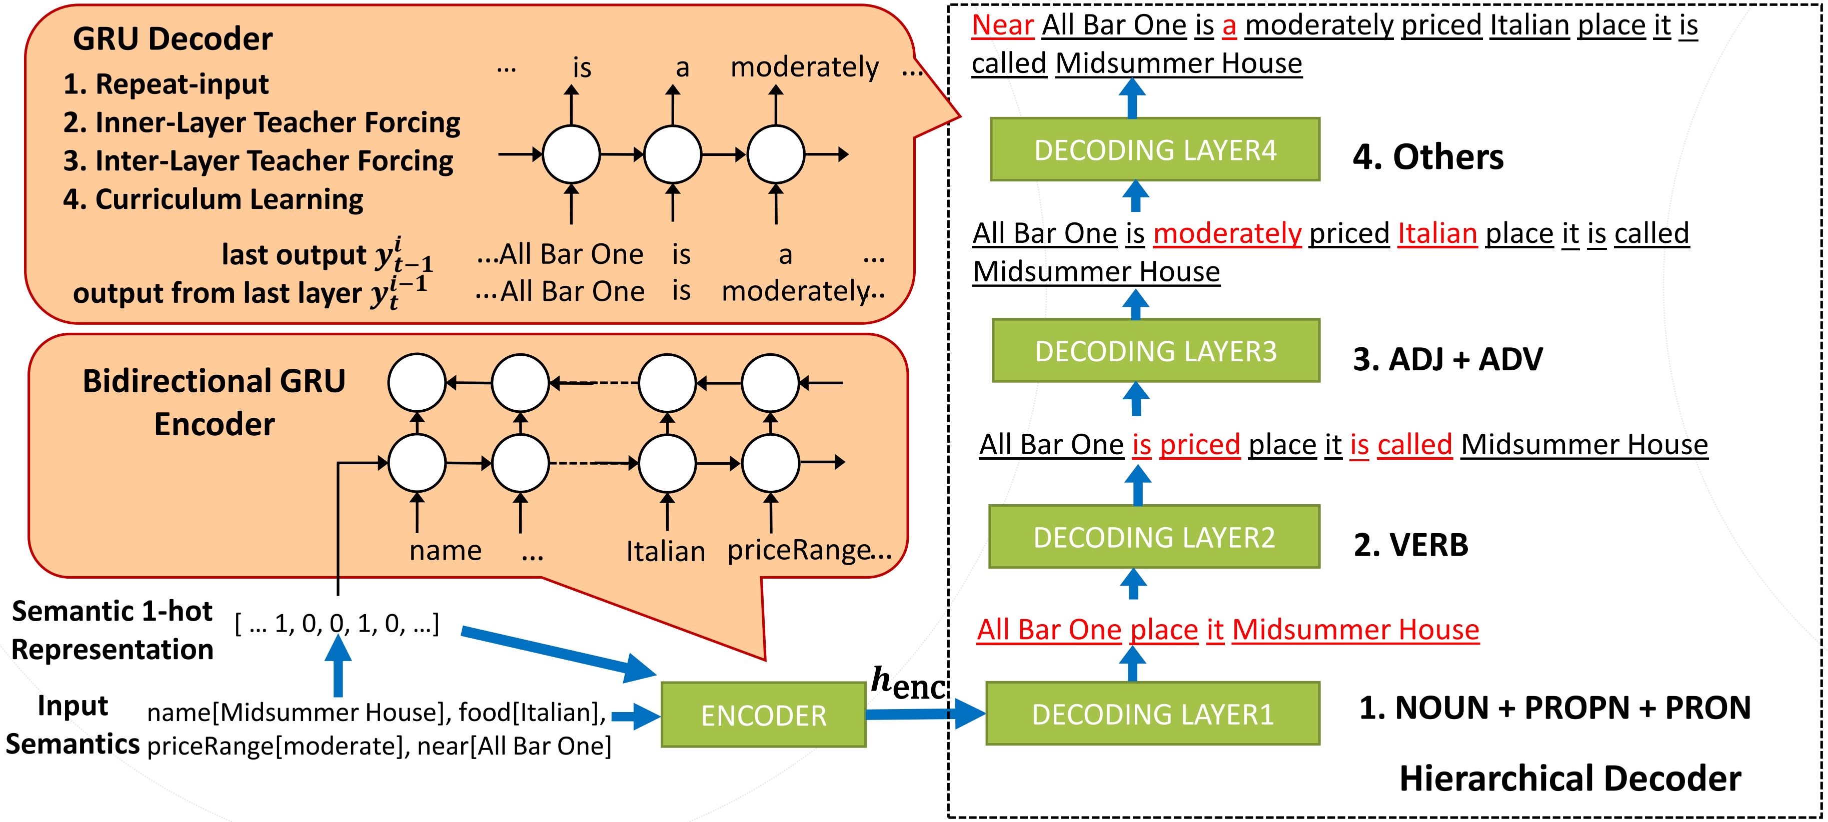

Idea: gradually generate words based on the linguistic knowledge

Material: http://deepdialogue.miulab.tw

Learning Discourse-Level Diversity (Zhao+, 2017)

- Conditional VAE
- Improves diversity of responses

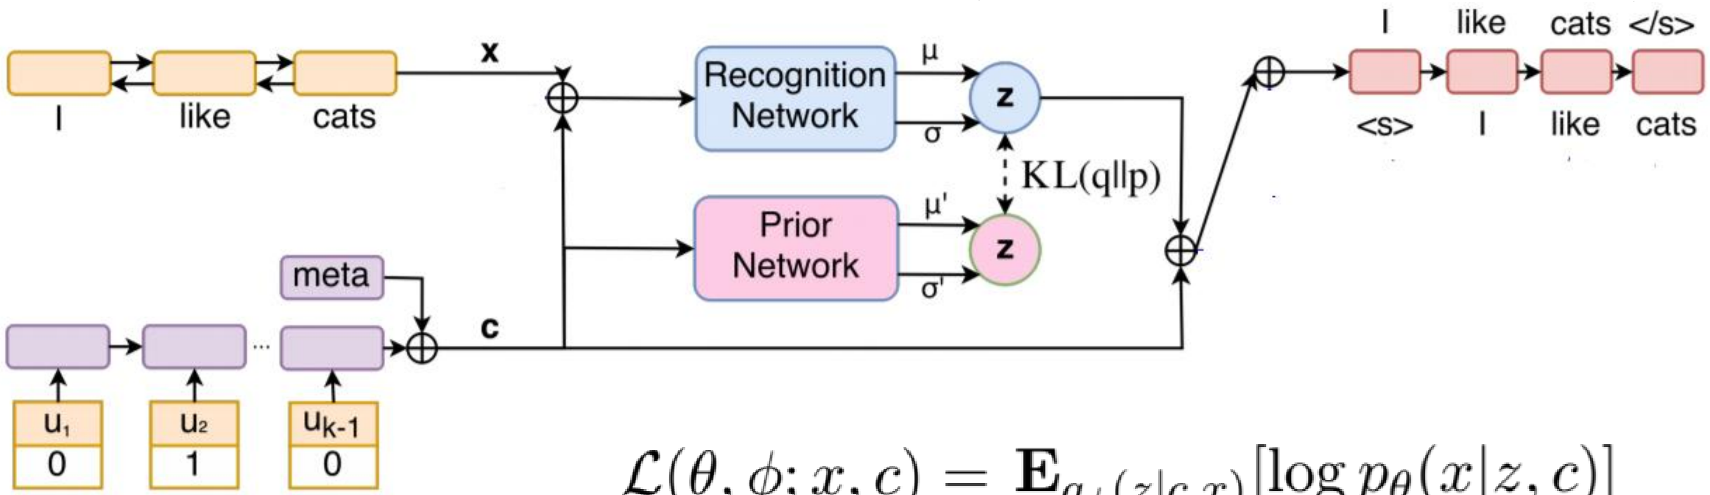

$$\mathcal{L}(\theta, \phi; x, c) = \mathbf{E}_{q_\phi(z|c,x)}[\log p_\theta(x|z, c)] - KL(q_\phi(z|x, c)\|p_\theta(z|c))$$

Learning Discourse-Level Diversity (Zhao+, 2017)

- Conditional VAE
- Improves diversity of responses with dialogue acts

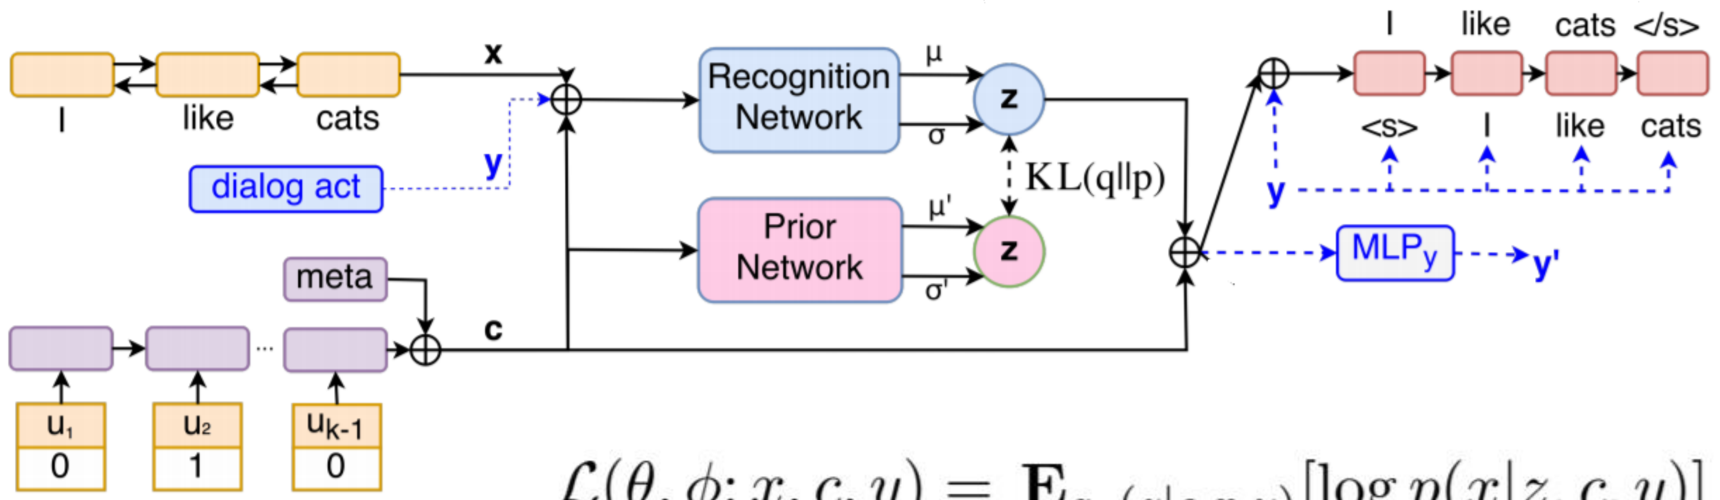

$$\mathcal{L}(\theta, \phi; x, c, y) = \mathbf{E}_{q_\phi(z|c,x,y)}[\log p(x|z, c, y)] + \mathbf{E}_{q_\phi(z|c,x,y)}[\log p(y|z, c)] - KL(q_\phi(z|x, c, y) \| P_\theta(z|c))$$

Learning Discourse-Level Diversity (Zhao+, 2017)

- Knowledge guided conditional VAE
- Improves diversity of responses with dialogue acts

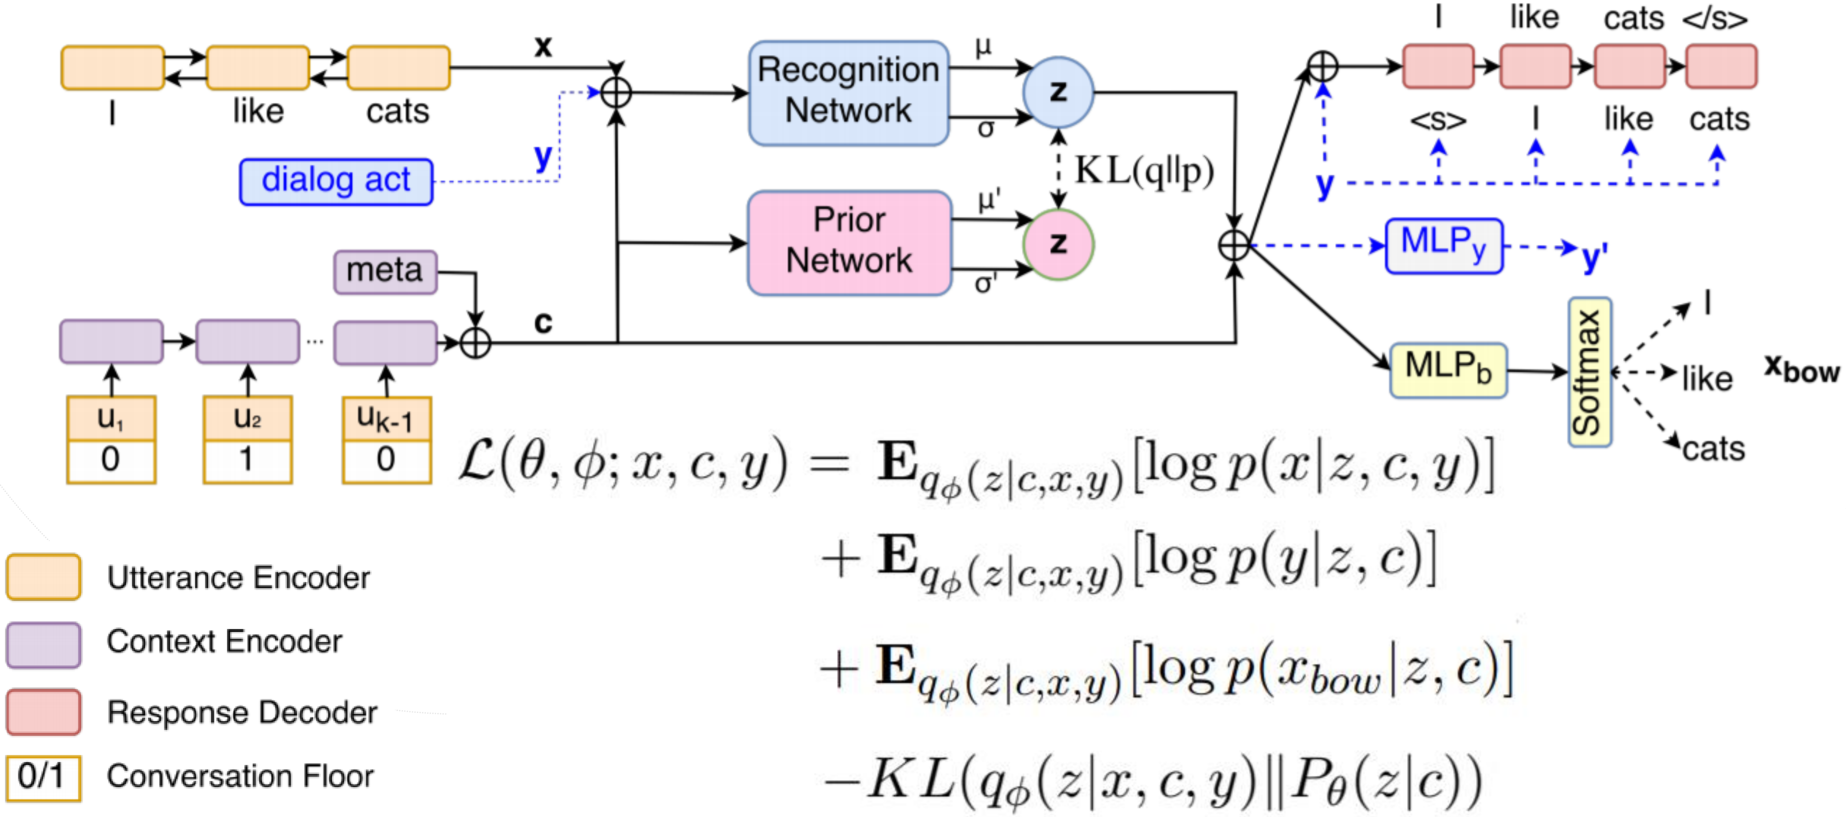

$$\mathcal{L}(\theta,\phi;x,c,y) = \mathbf{E}_{q_\phi(z|c,x,y)}[\log p(x|z,c,y)]$$
$$+\mathbf{E}_{q_\phi(z|c,x,y)}[\log p(y|z,c)]$$
$$+\mathbf{E}_{q_\phi(z|c,x,y)}[\log p(x_{bow}|z,c)]$$
$$-KL(q_\phi(z|x,c,y)\|P_\theta(z|c))$$

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

- Introduction
- Background Knowledge
- Modular Dialogue System
- **System Evaluation**
- Recent Trends of Learning Dialogues

Dialogue System Evaluation

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

Material: http://deepdialogue.miulab.tw

Crowdsourcing for System Evaluation (Yang+, 2012)

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

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.*

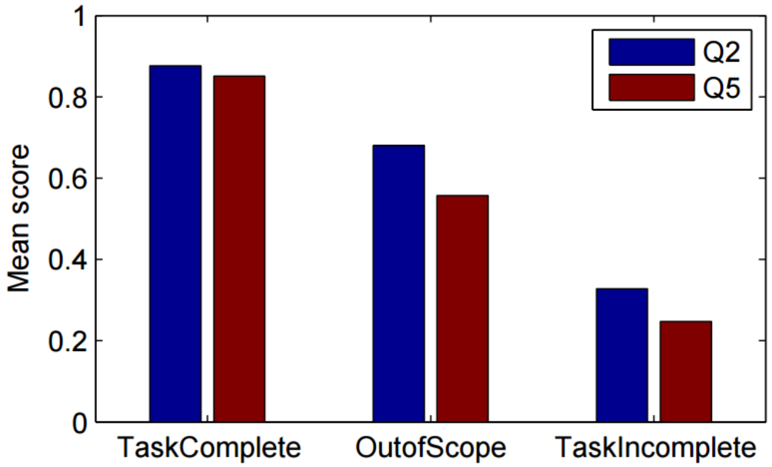

User Simulation

keeps a list of its goals and actions → randomly generates an agenda → updates its list of goals and adds new ones

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

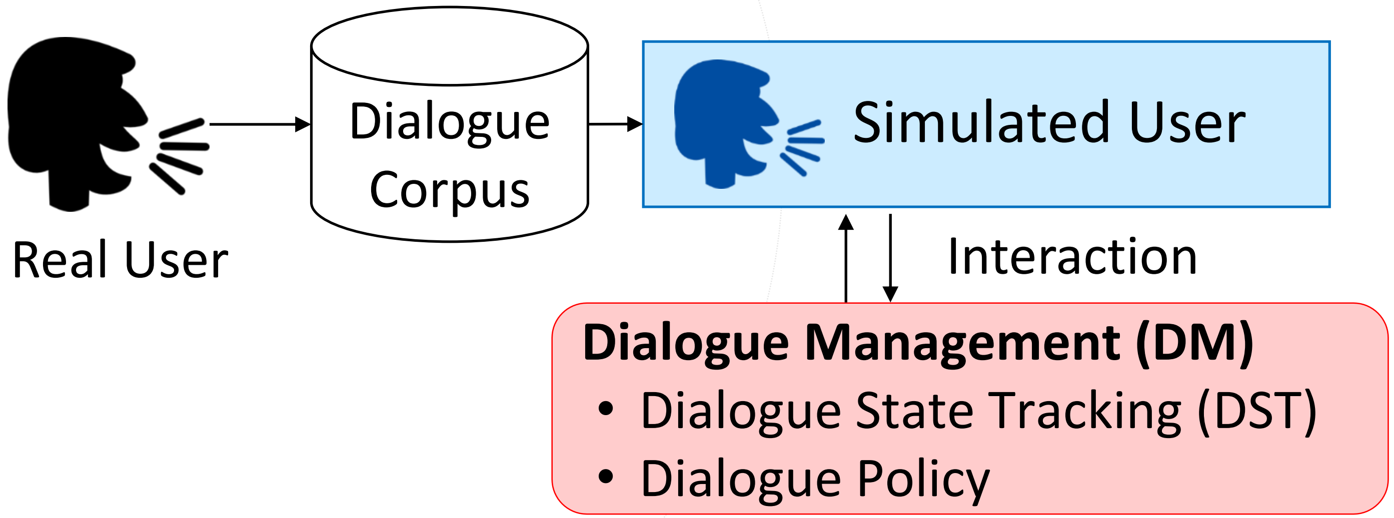

- Approach
  - Rule-based crafted by experts (Li et al., 2016)
  - Learning-based (Schatzmann et al., 2006; El Asri et al., 2016, Crook and Marin, 2017)

User Simulation

keeps a list of its goals and actions → randomly generates an agenda → updates its list of goals and adds new ones

- First, generate a user goal.
- The user goal contains:
  - Dialog act
  - Inform slots
  - Request slots


```
{
  "request_slots": {
    "ticket": "UNK",
    "theater": "UNK"
  },
  "diaact": "request",
  "inform_slots": {
    "city": "birmingham",
    "numberofpeople": "2",
    "state": "al",
    "starttime": "4 pm",
    "date": "today",
    "moviename": "deadpool"
  }
}
```


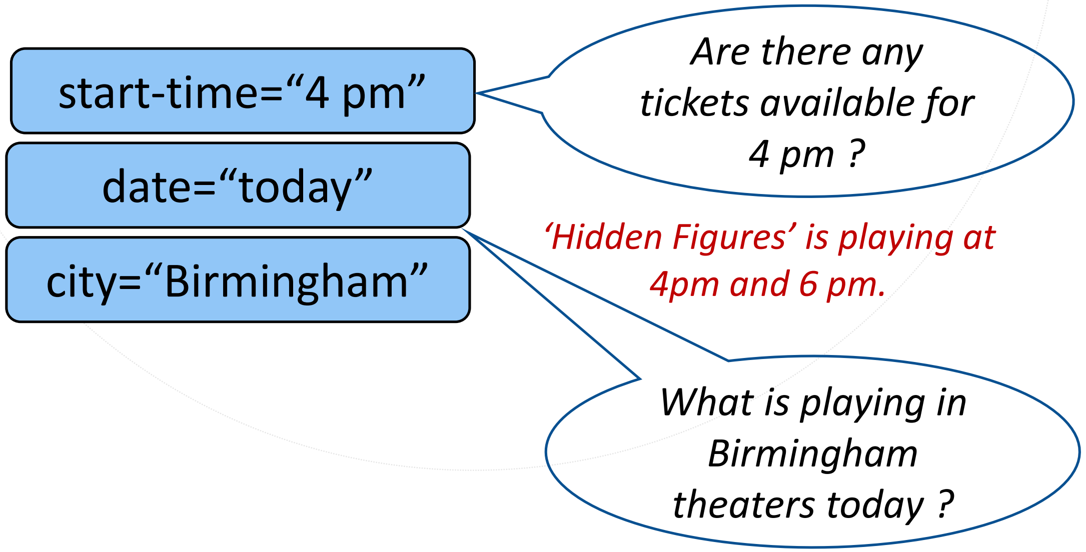

Elements of User Simulation

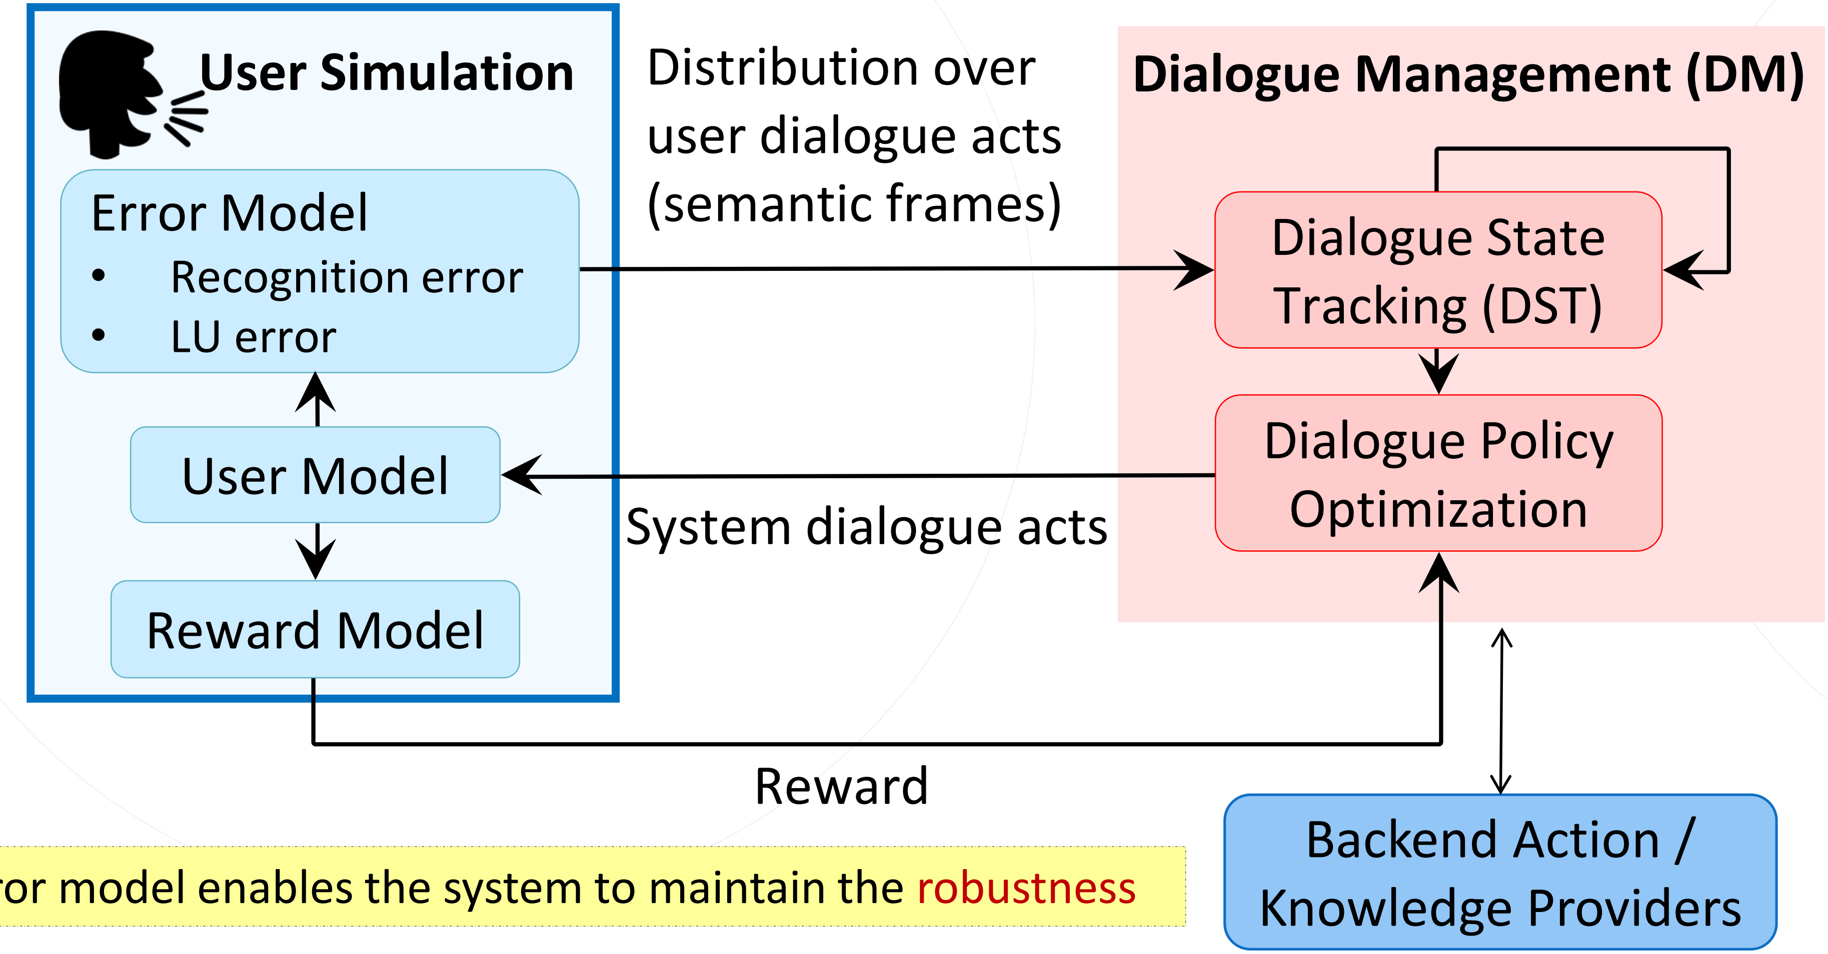

The error model enables the system to maintain the robustness

Rule-Based Simulator for RL 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

```python
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

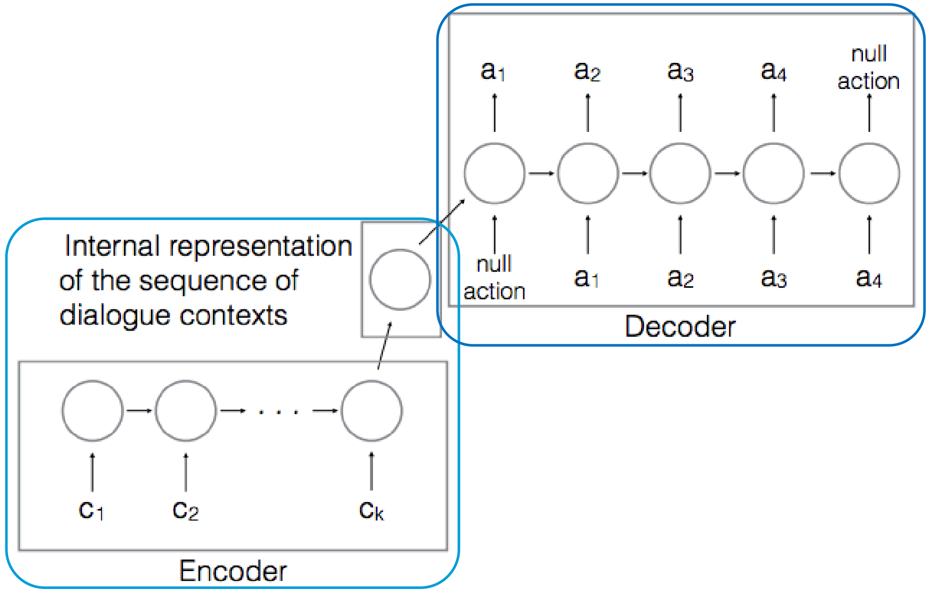

Seq2Seq User Simulation (Crook and Marin, 2017)

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

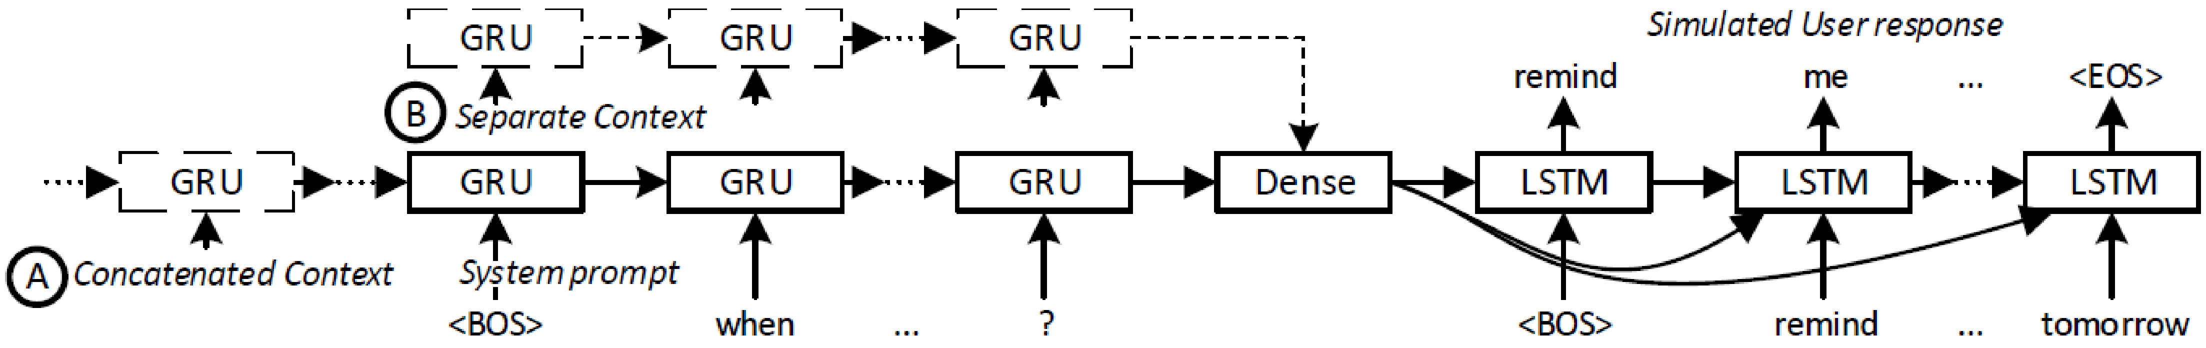

Material: http://deepdialogue.miulab.tw

User Simulator for Dialogue Evaluation Measures

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

How NOT to Evaluate Dialog System (Liu+, 2017)

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

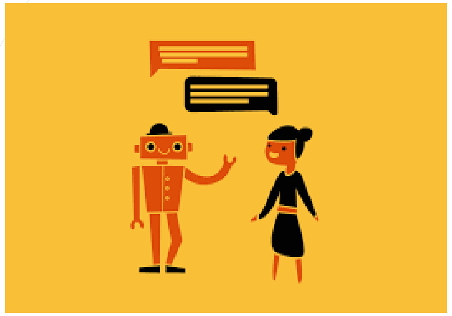

- Metrics:
  - Word overlap metrics, e.g., BLEU, METEOR, ROUGE, etc.
  - Embeddings based metrics, e.g., contextual/meaning representation between target and candidate

Dialogue Response Evaluation (Lowe+, 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!*

Outline

- Introduction
- Background Knowledge
- Modular Dialogue System
- System Evaluation
- **Recent Trends of Learning Dialogues**
  - End-to-End Neural Dialogue System
  - Multimodality
  - Dialogue Breadth & Dialogue Depth

Outline

- Introduction
- Background Knowledge
- Modular Dialogue System
- System Evaluation
- **Recent Trends of Learning Dialogues**
  - **End-to-End Neural Dialogue System**
  - Multimodality
  - Dialogue Breadth & Dialogue Depth

ChitChat Hierarchical Seq2Seq (Serban et al., 2016)

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

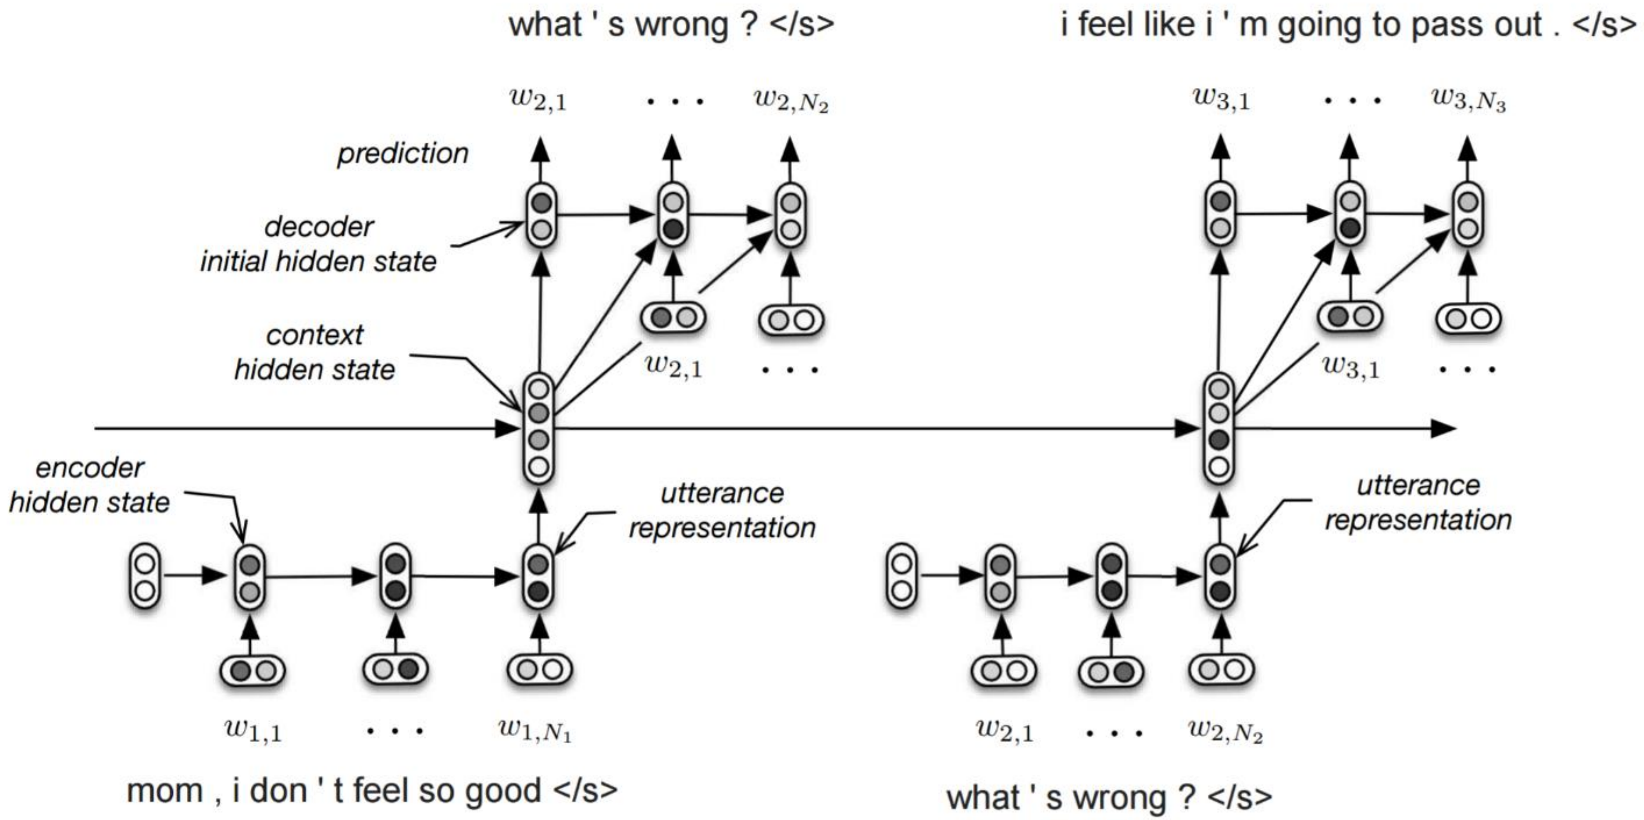

ChitChat Hierarchical Seq2Seq (Serban et.al., 2017)

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

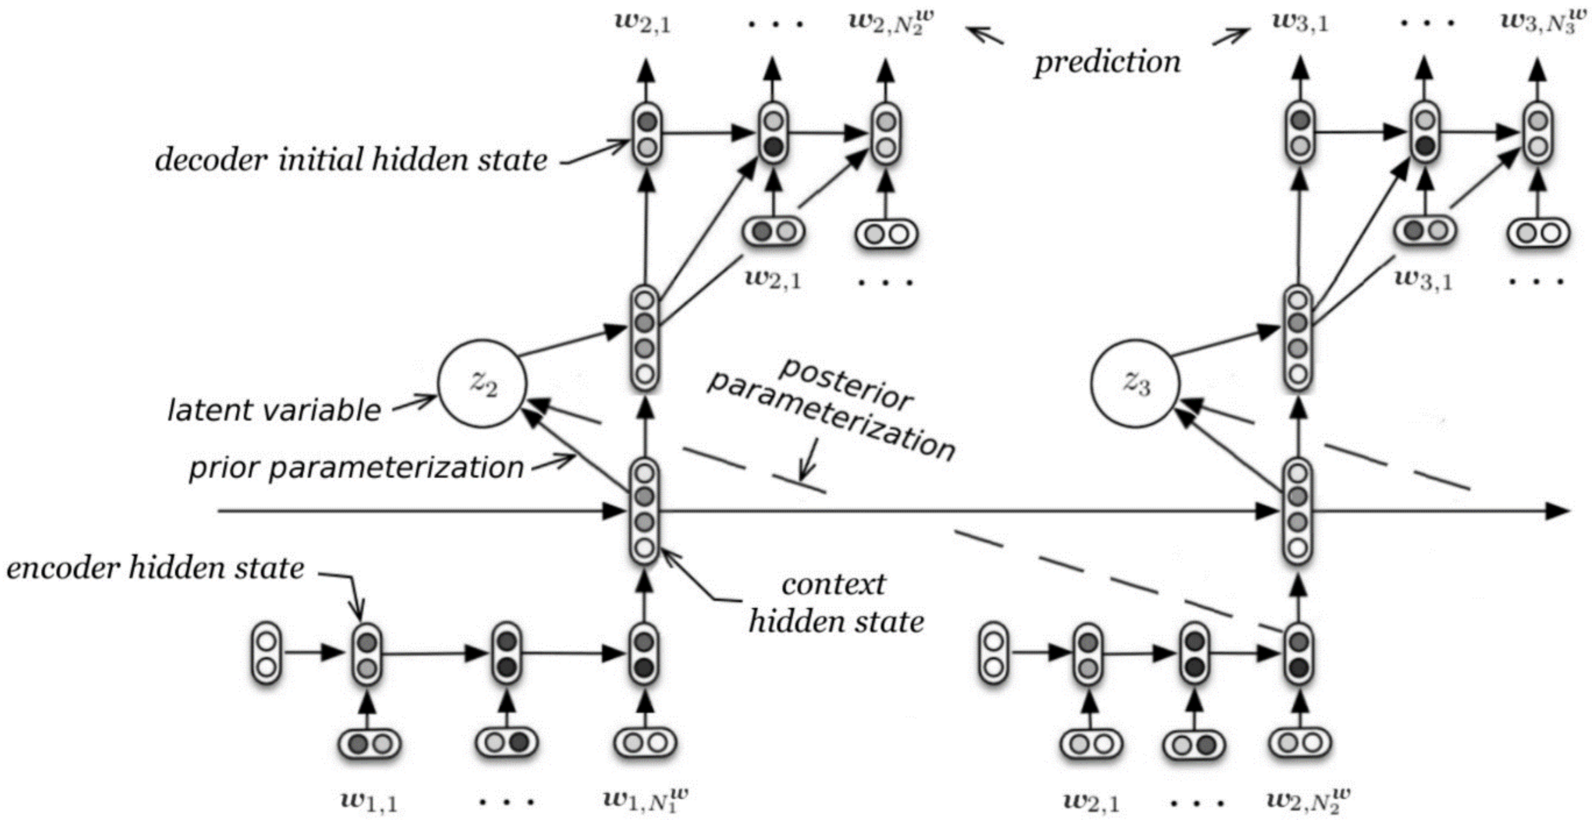

Material: http://deepdialogue.miulab.tw

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

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

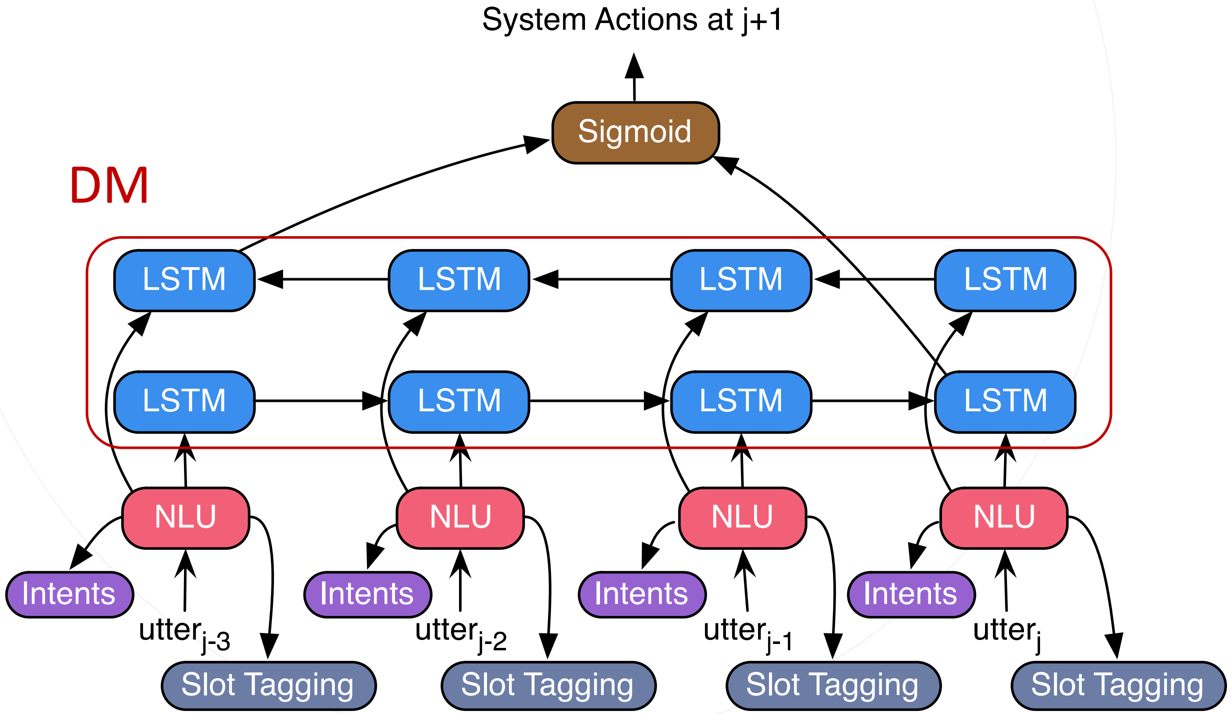

Model	DM	NLU
Baseline (CRF+SVMs)	7.7	33.1
Pipeline-BLSTM	12.0	36.4
JointModel	**22.8**	**37.4**

Both DM and NLU performance (frame accuracy) is improved

Material: http://deepdialogue.miulab.tw

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

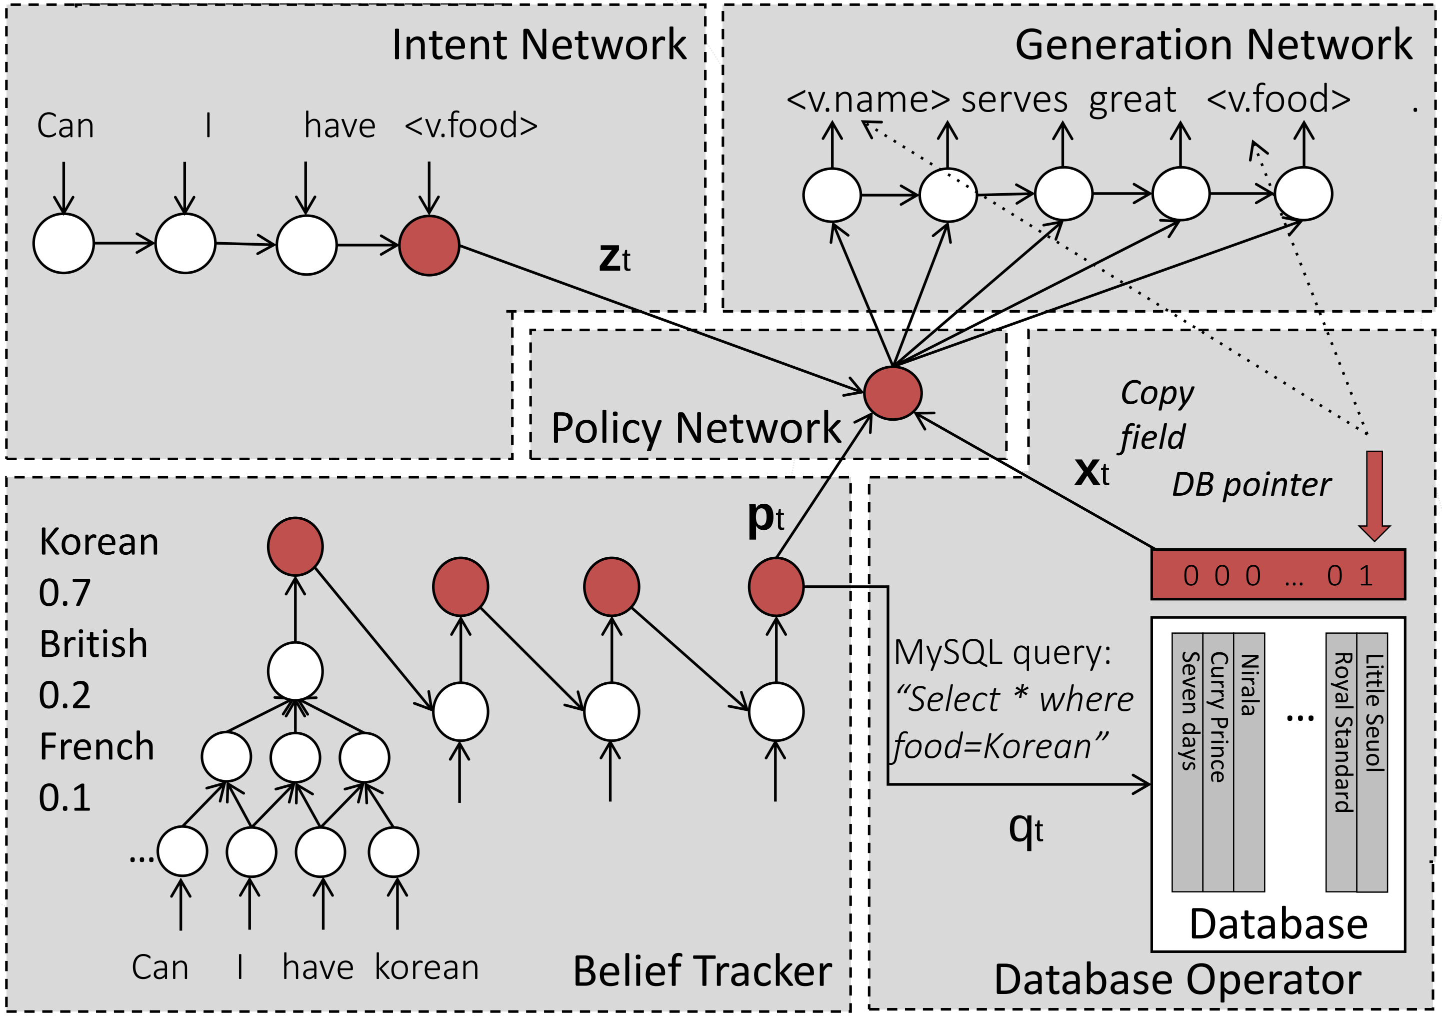

E2E MemNN for Dialogues (Bordes et al., 2016)

- 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)

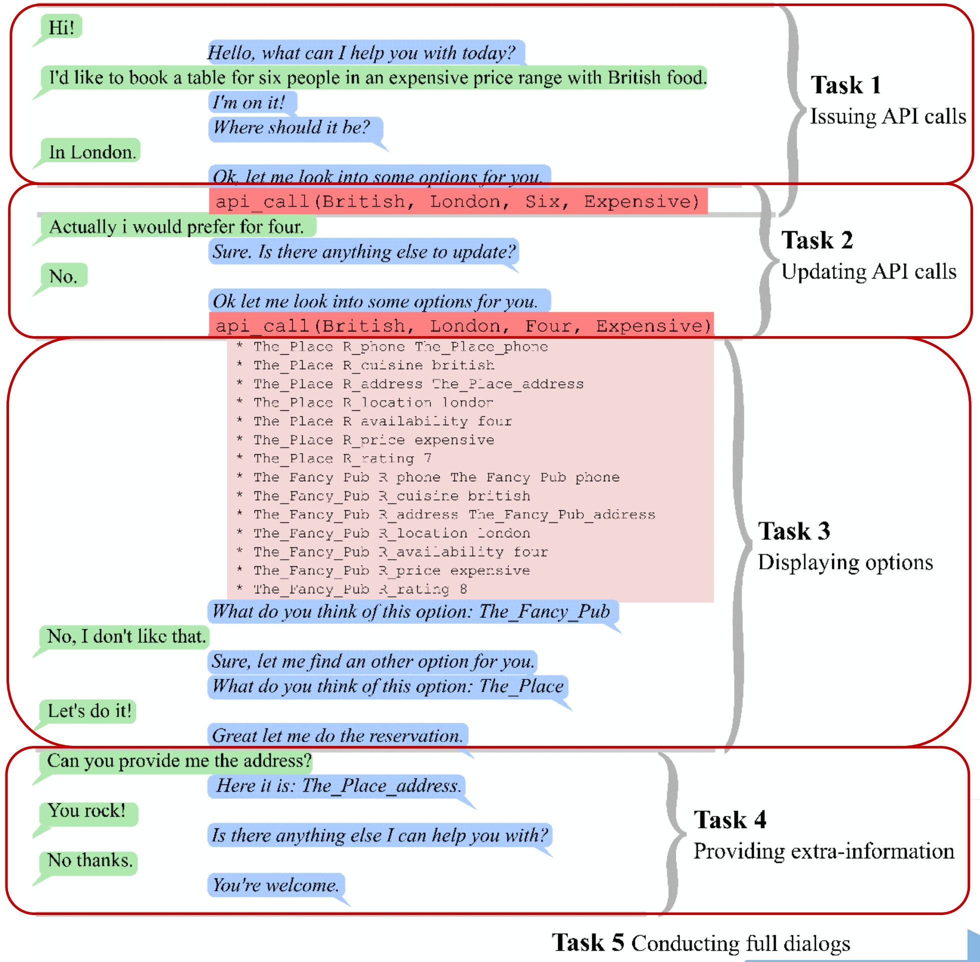

E2E RL-Based KB-InfoBot (Dhingra et al., 2017)

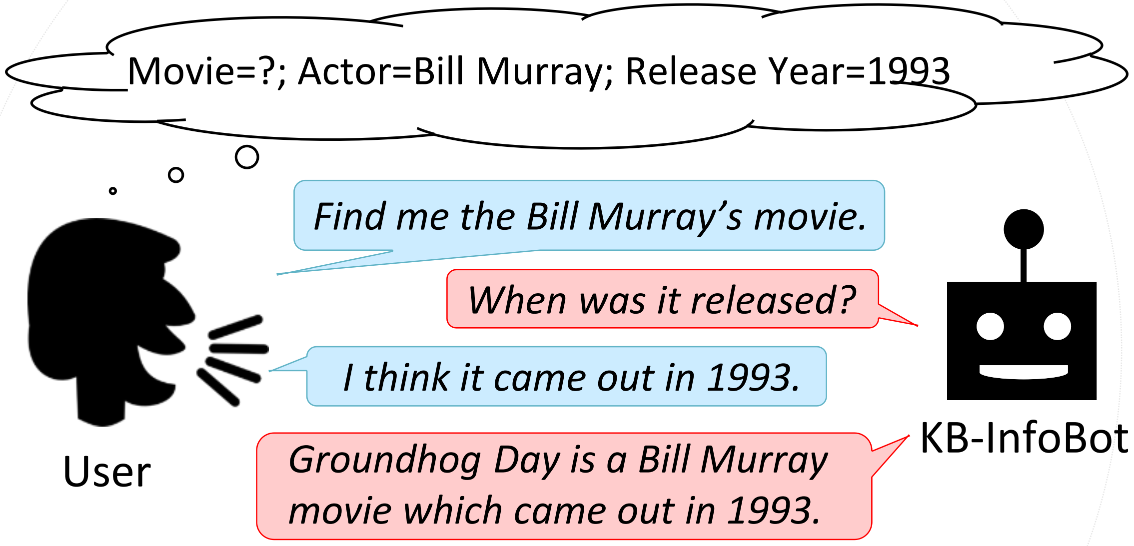

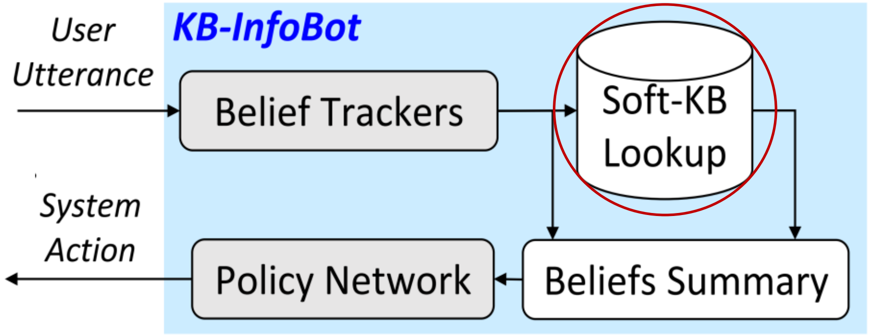

Entity-Centric Knowledge Base

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

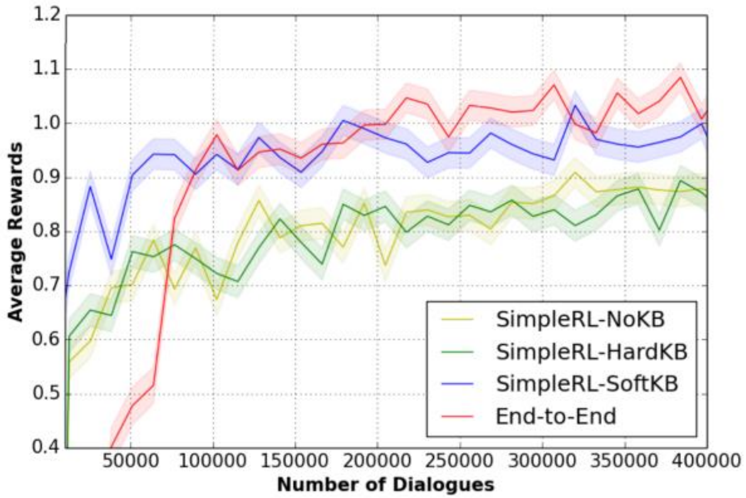

Idea: differentiable database for propagating the gradients

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

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

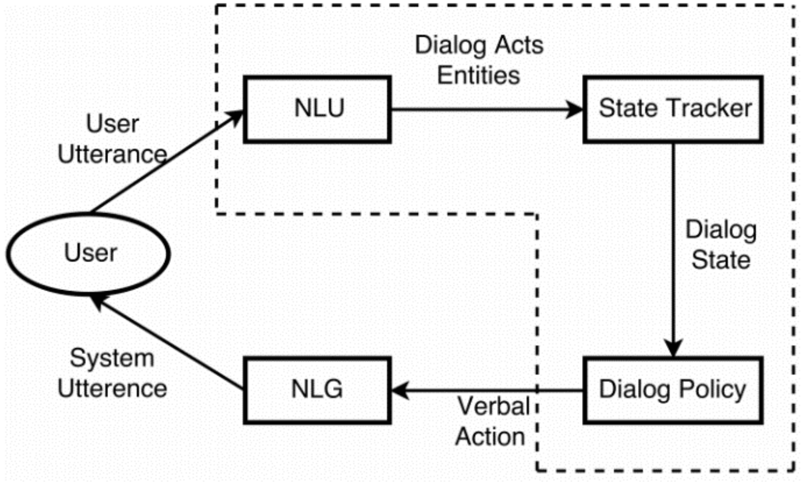

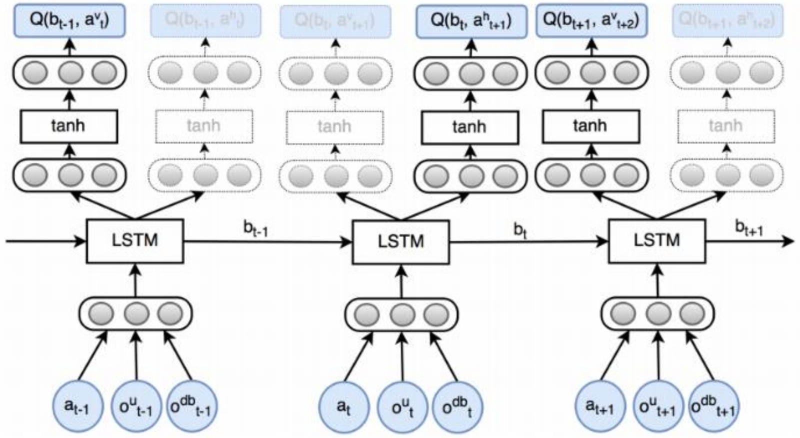

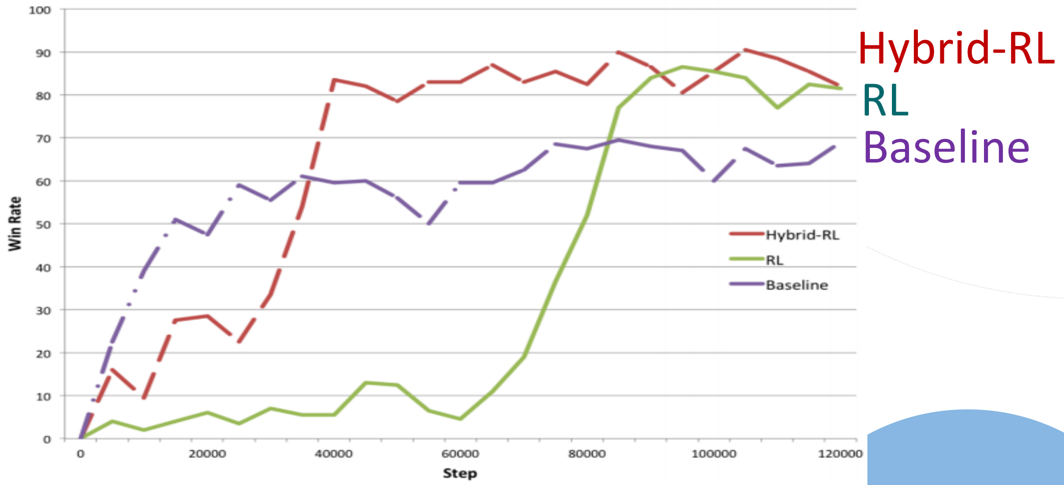

E2E LSTM-Based Dialogue Control (Williams and Zweig, 2016)

- 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

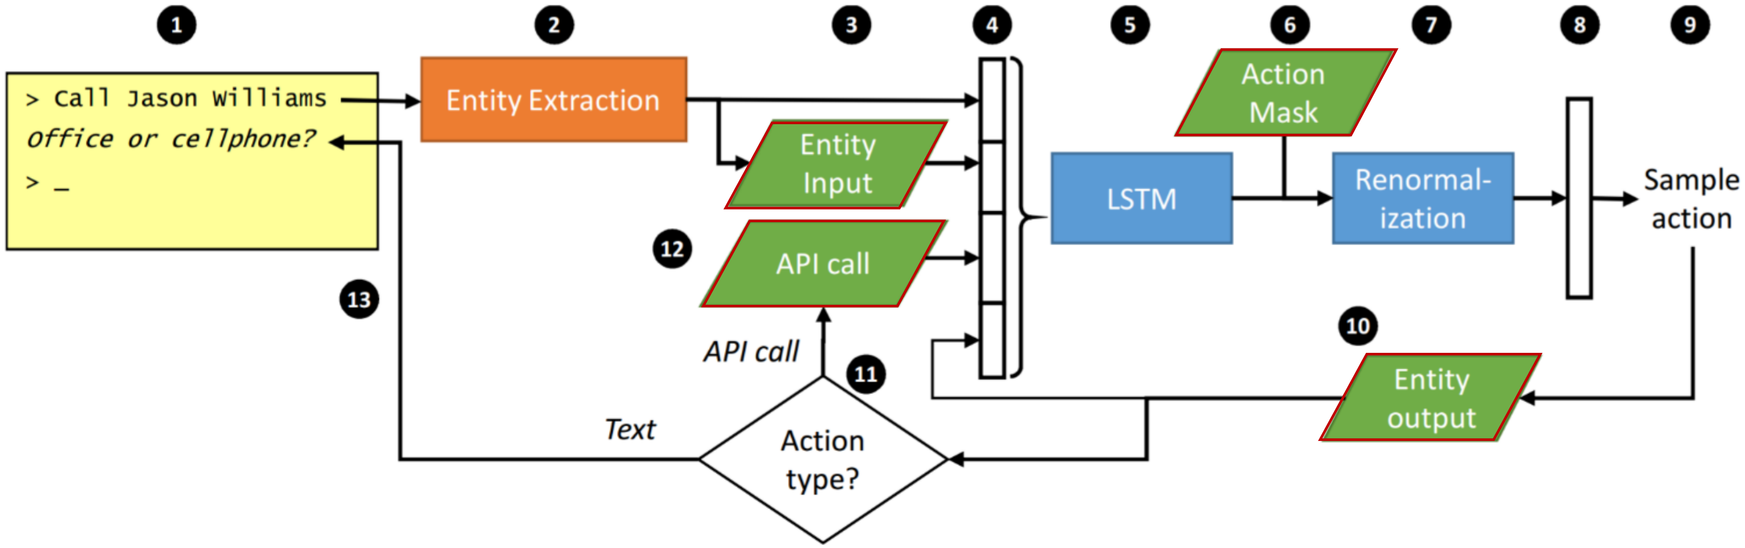

E2E Task-Completion Bot (TC-Bot) (Li et al., 2017)

- 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!

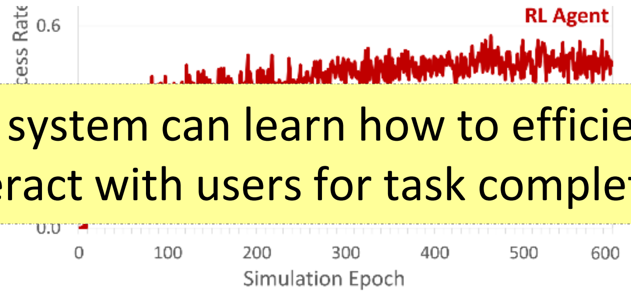

The system can learn how to efficiently interact with users for task completion

E2E Imitation and RL Agent (Liu+, 2018)

- Generate distribution over candidate slot values:

$$s_k = \mathrm{LSTM}(s_{k-1},\ [U_k,\ A_{k-1}])$$
$$P(l_k^m \mid \mathbf{U}_{\le k},\ \mathbf{A}_{<k}) = \mathrm{SlotDist}_m(s_k)$$

- Generate system action:

$$P(a_k \mid U_{\le k},\ A_{<k},\ E_{\le k})$$
$$= \mathrm{PolicyNet}(s_k, v_k, E_k)$$

- Train Supervised → REINFORCE

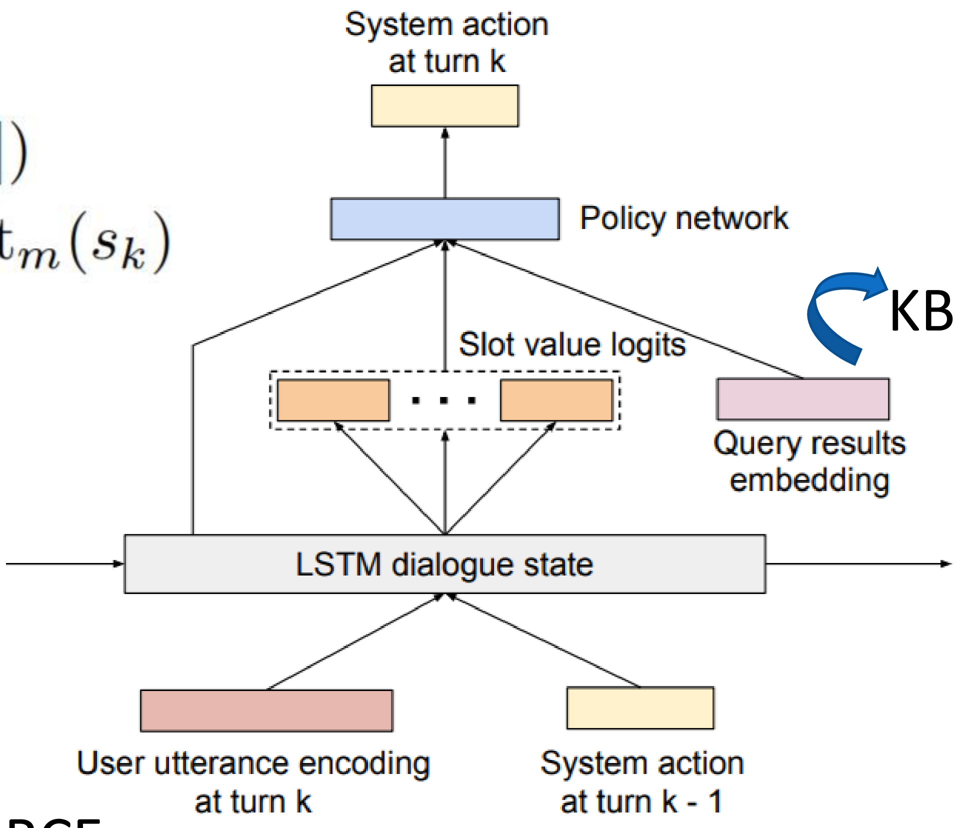

Material: http://deepdialogue.miulab.tw

Dialogue Challenge

- DSTC: Dialog System Technology Challenge

Challenge	Track	Theme
DSTC6	Track 1	End-to-End Goal-Oriented Dialog Learning
	Track 2	End-to-End Conversation Modeling
	Track 3	Dialogue Breakdown Detection
DSTC7	Track 1	Sentence Selection
	Track 2	Sentence Generation
	Track 3	AVSD: **A**udio **V**isual **S**cene-aware **D**ialog

- SLT 2018 Microsoft Dialogue Challenge: End-to-End Task-Completion Dialogue Systems
- The Conversation Intelligence Challenge: ConvAI2 - PersonaChat

Outline

- Introduction
- Background Knowledge
- Modular Dialogue System
- System Evaluation
- **Recent Trends of Learning Dialogues**
  - End-to-End Neural Dialogue System
  - **Multimodality**
  - Dialogue Breadth & Dialogue Depth

Brain Signal for Understanding

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

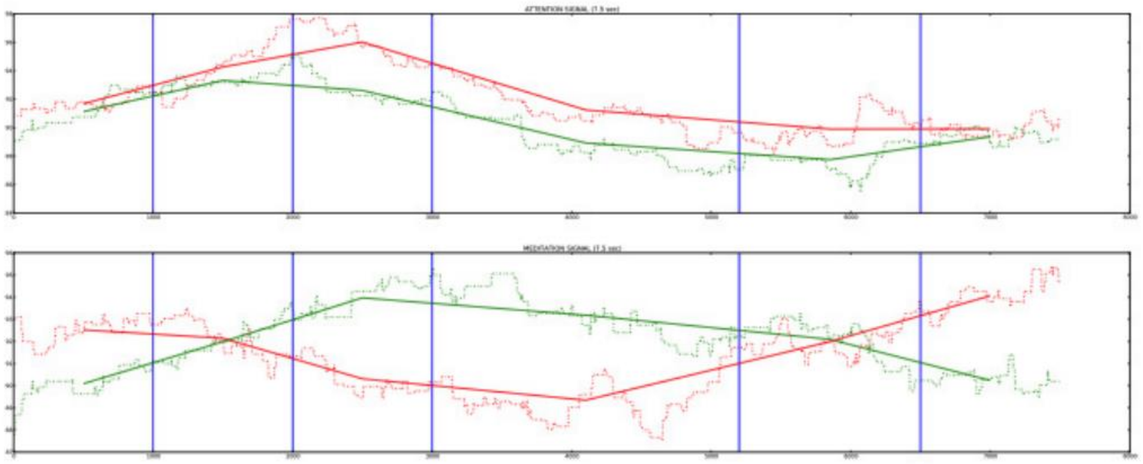

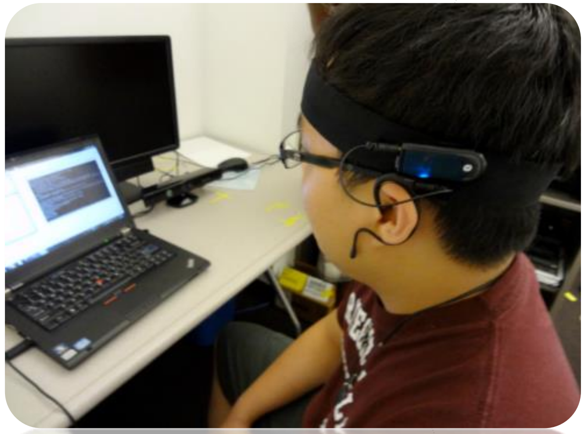

Detecting misunderstanding via brain signal in order to correct the understanding results

Video for Intent Understanding

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

- Task: user intent prediction
- Challenge: language ambiguity

send to vivian

Communication

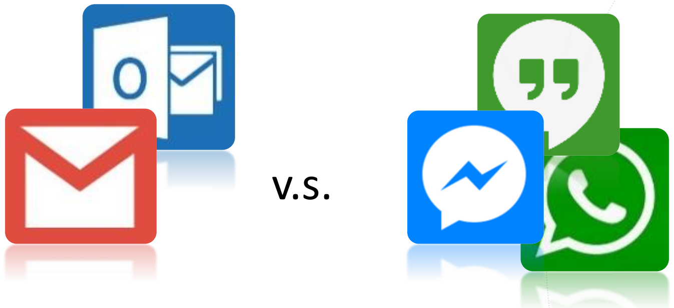

Email? **Message?**

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.

Material: http://deepdialogue.miulab.tw

Video Highlight Prediction Using Audience Chats

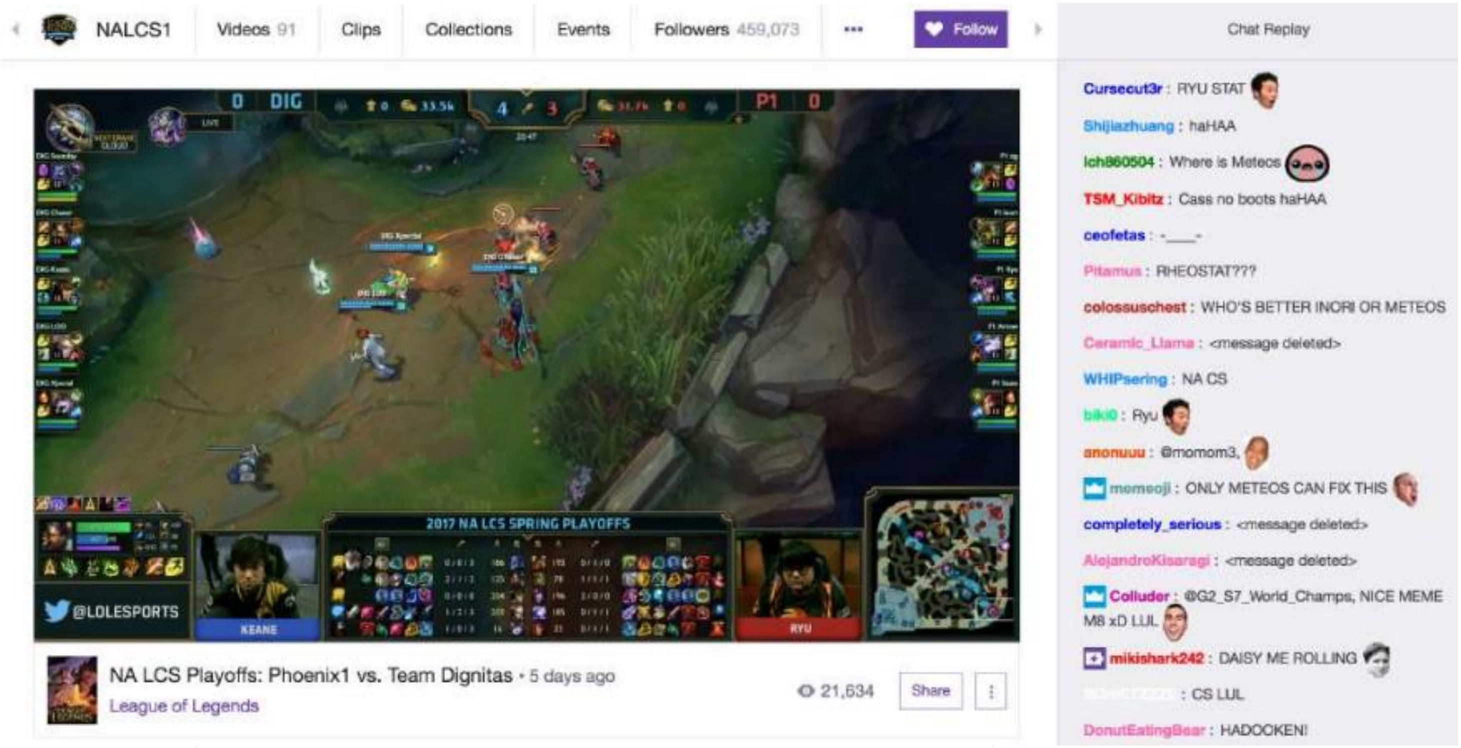

Video Highlight Prediction Using Audience Chats

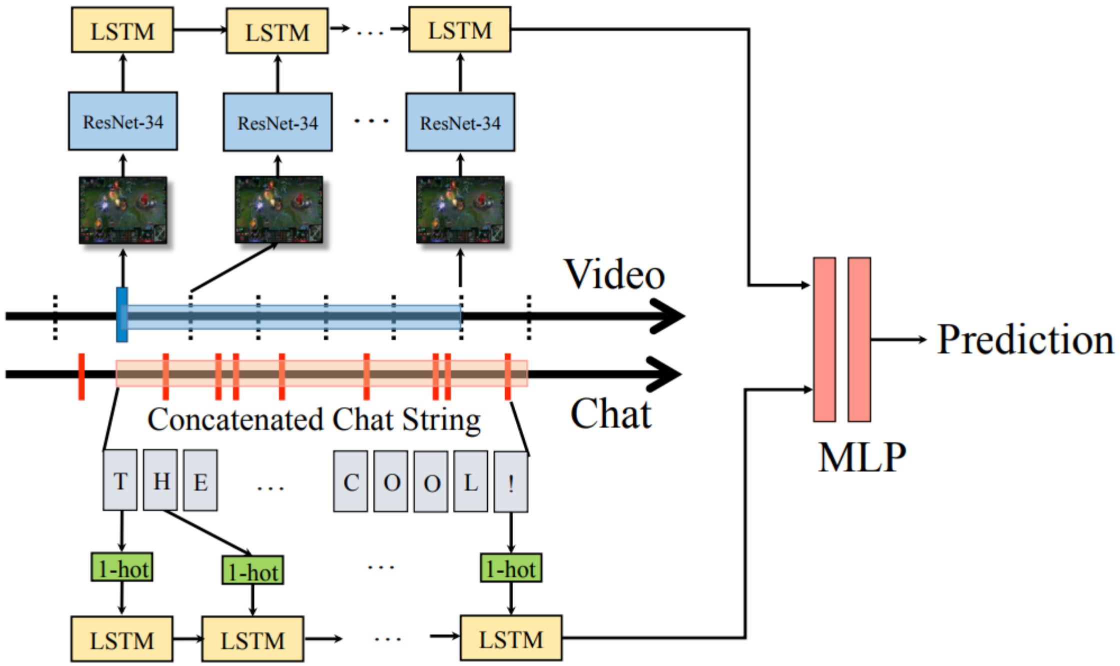

- Goal: predict highlight from the video
- Input : multi-modal and multi-lingual (real time text commentary from fans)
- Output: tag if a frame part of a highlight or not

Outline

- Introduction
- Background Knowledge
- Modular Dialogue System
- System Evaluation
- **Recent Trends of Learning Dialogues**
  - End-to-End Neural Dialogue System
  - Multimodality
  - **Dialogue Breadth & Dialogue Depth**

Material: http://deepdialogue.miulab.tw

Evolution Roadmap

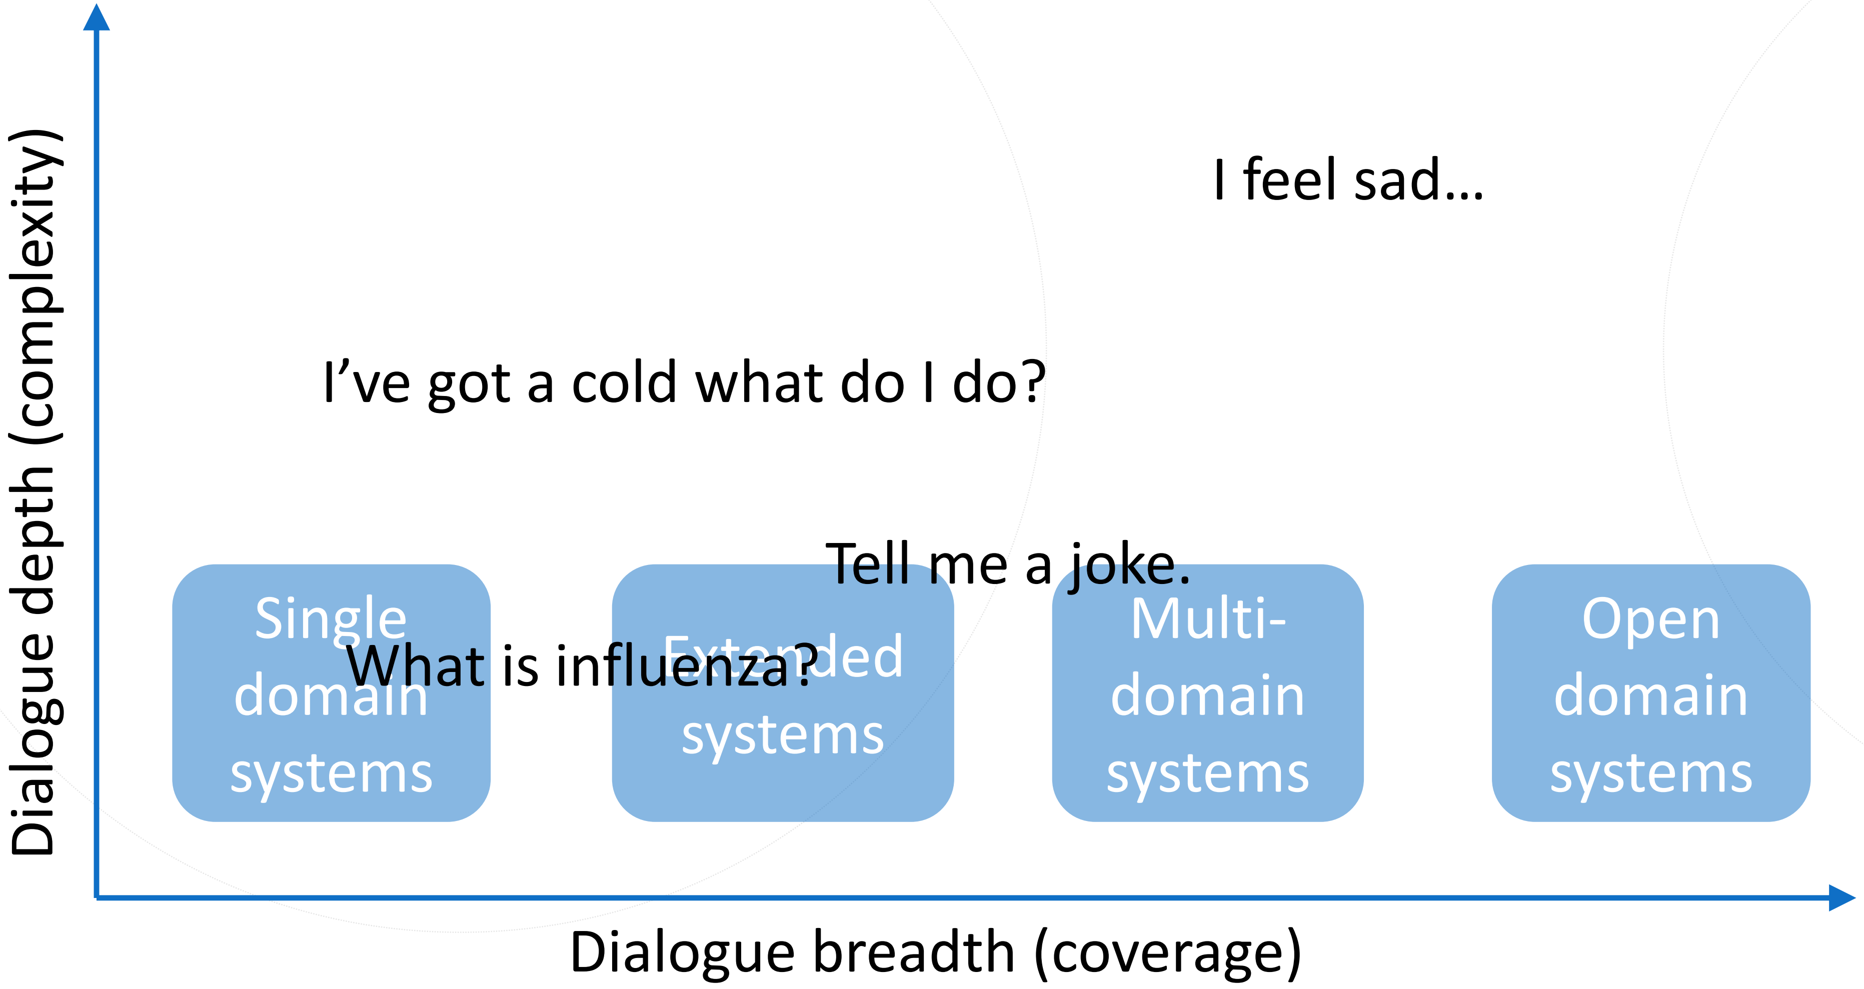

Material: http://deepdialogue.miulab.tw

Evolution Roadmap

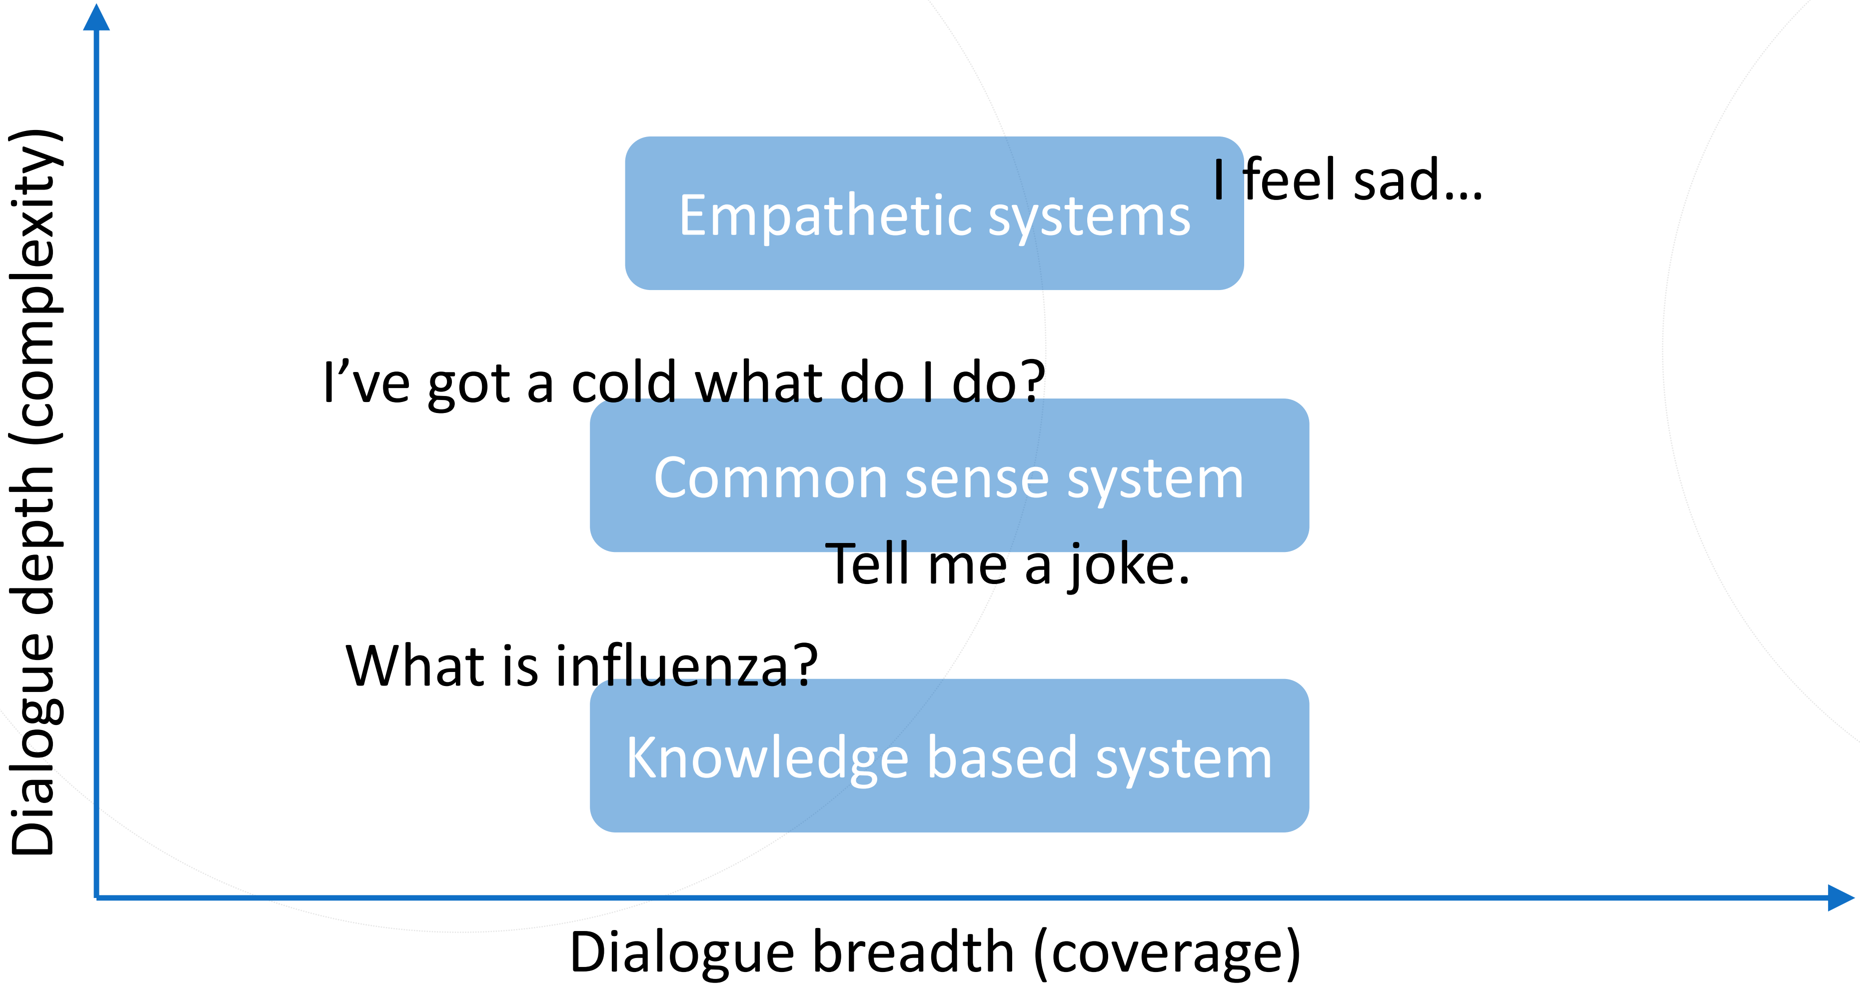

Common Sense for Dialogue Planning (Sun+, 2016)

- High-level intention may span several domains

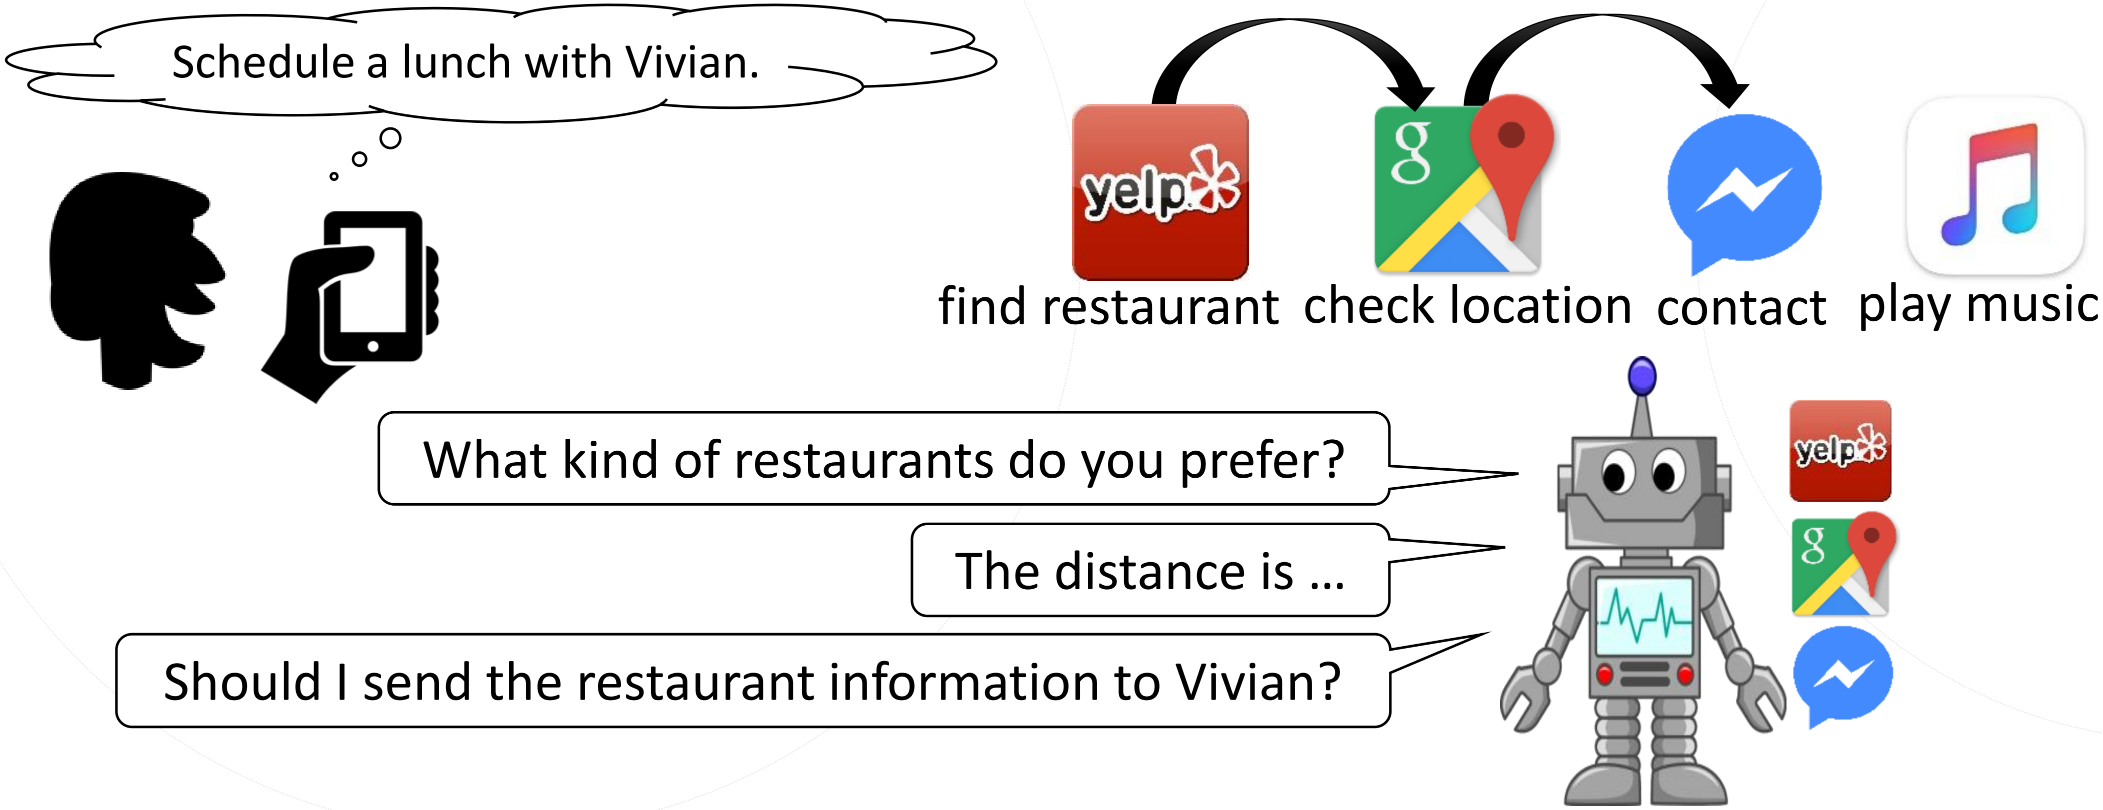

Users can interact via high-level descriptions and the system learns how to plan the dialogues

Material: http://deepdialogue.miulab.tw

Empathy in Dialogue System (Fung+, 2016)

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

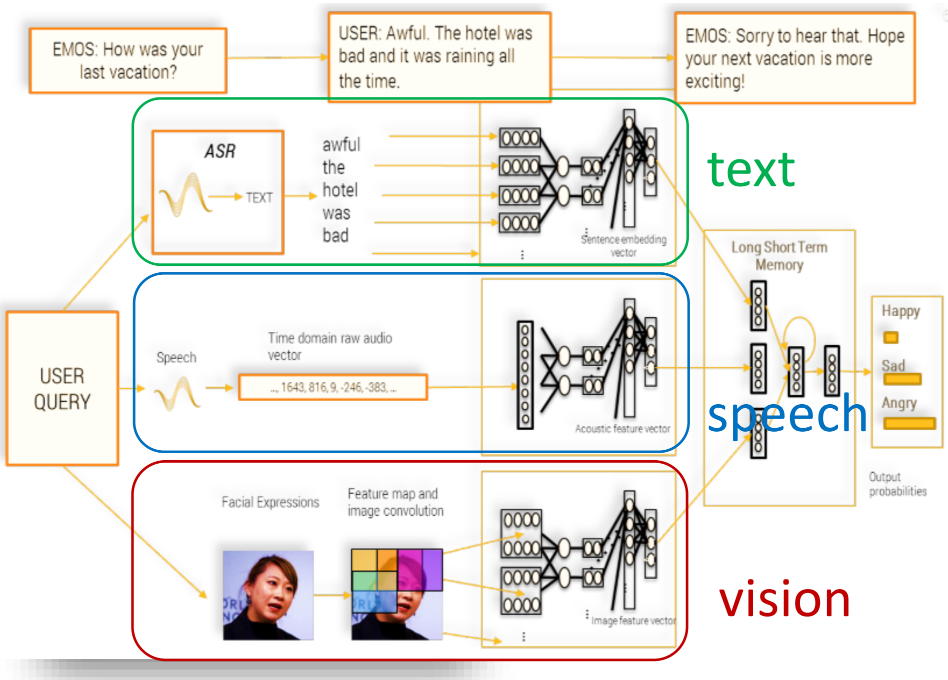

Emotion Recognizer

Zara - The Empathetic Supergirl

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

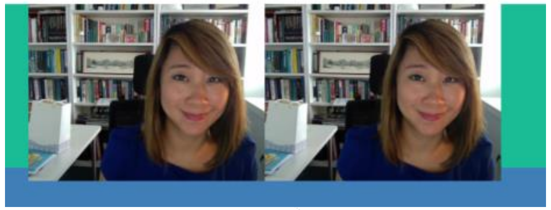

Face recognition output

(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 via 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

Conclusions

Summarized Challenges

- 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
- Data collection and analysis from un-structured data
- Complex-cascade systems requires high accuracy for working good as a whole

Brief Conclusions

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

http://deepdialogue.miulab.tw

NTU MIULAB

THANK YOU

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

Yun-Nung (Vivian) Chen

http://vivianchen.idv.tw