



**THE OHIO STATE
UNIVERSITY**

CSE 5525: Foundations of Speech and Language Processing

Dialogue

Huan Sun (CSE@OSU)

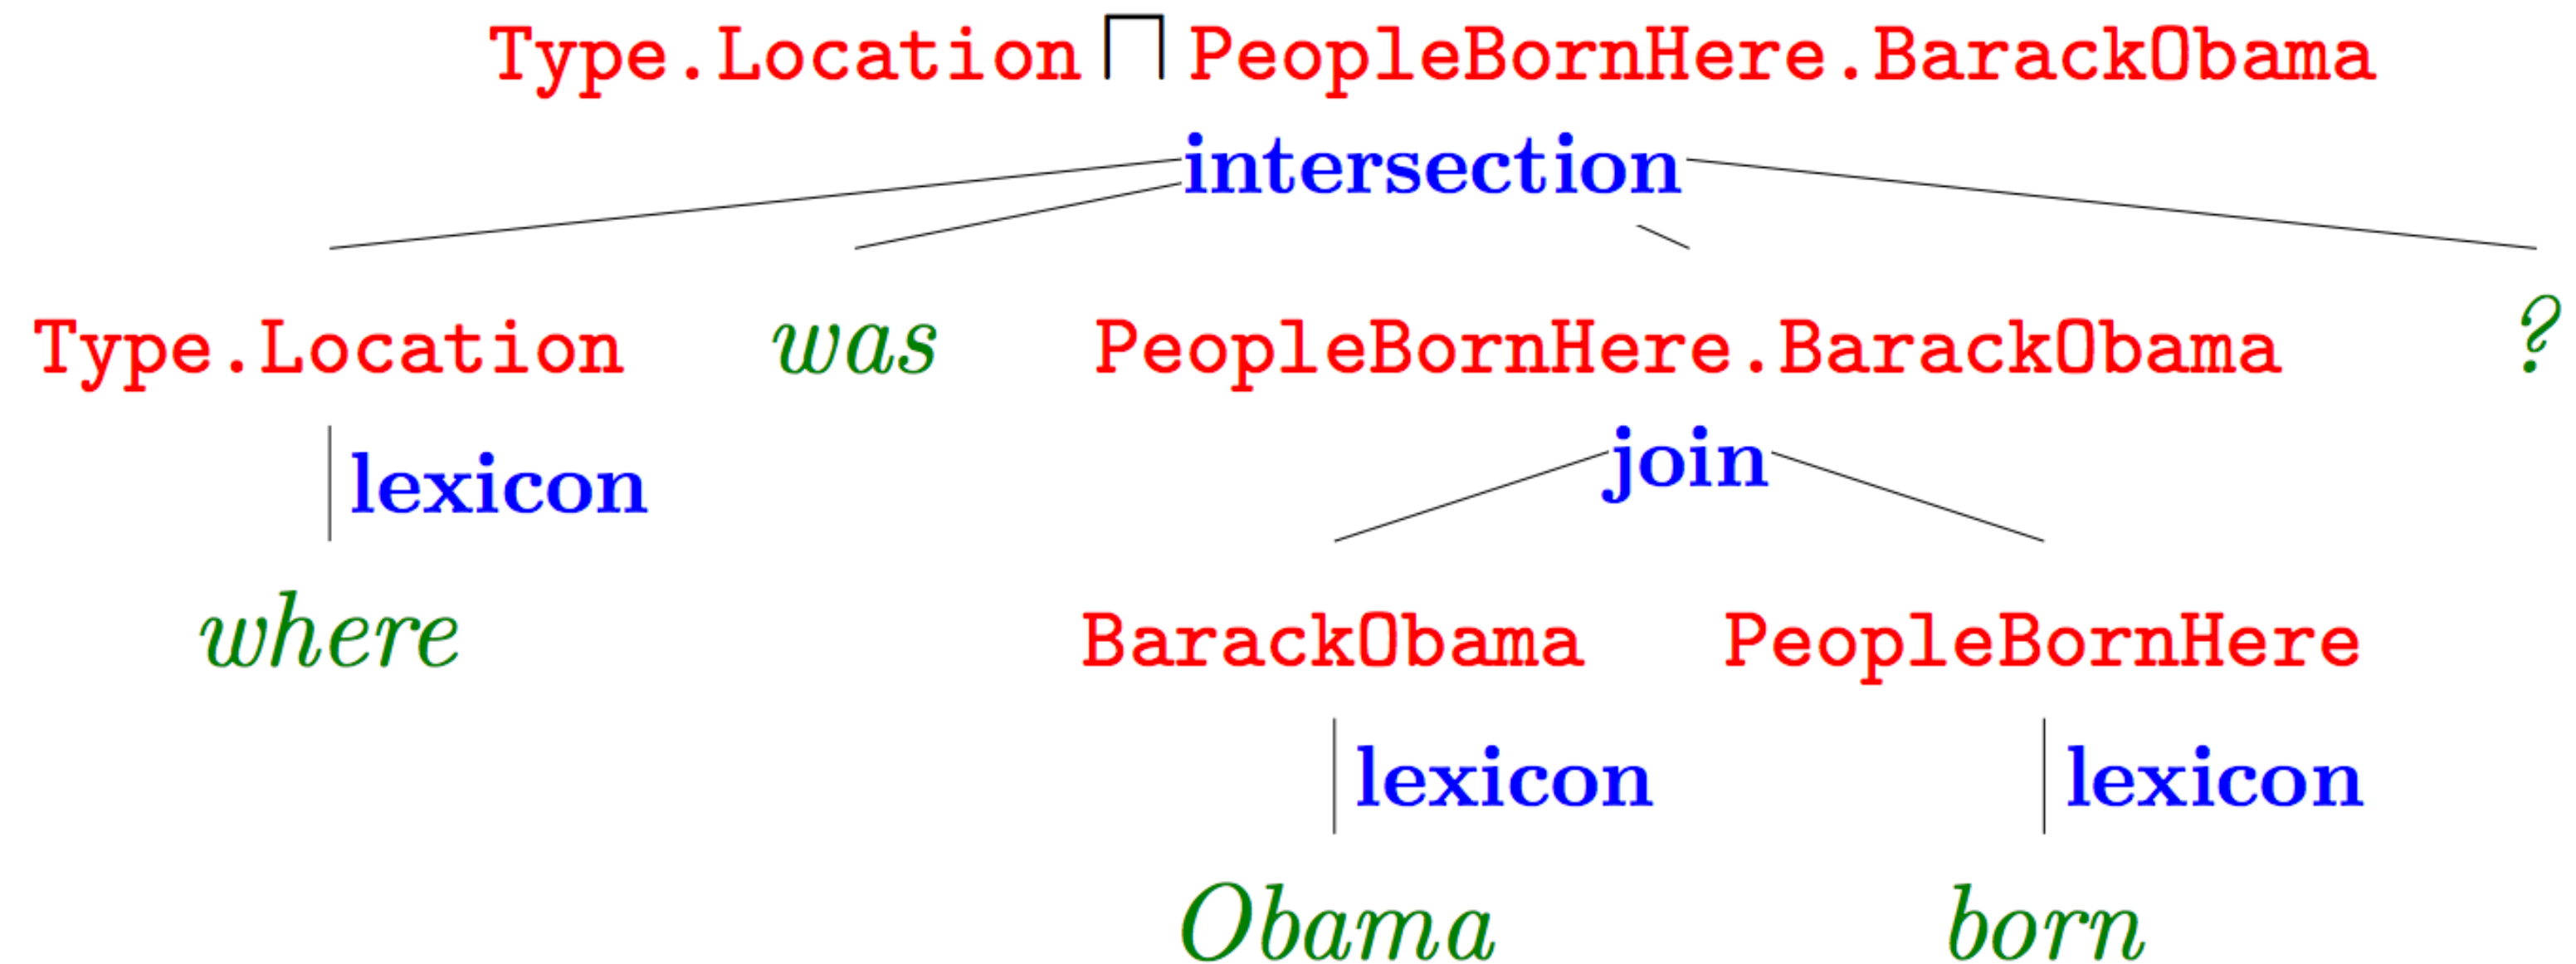
Slides were largely adapted from Prof. Greg Durrett @ UT Austin.
Some images/examples were from the two textbooks by (1) Jurafsky and Martin and (2) Eisenstein.

This Lecture

- ▶ Review QA & connect QA with dialogue
- ▶ Classic chatbots
- ▶ Neural chatbots
- ▶ Task-oriented dialogue

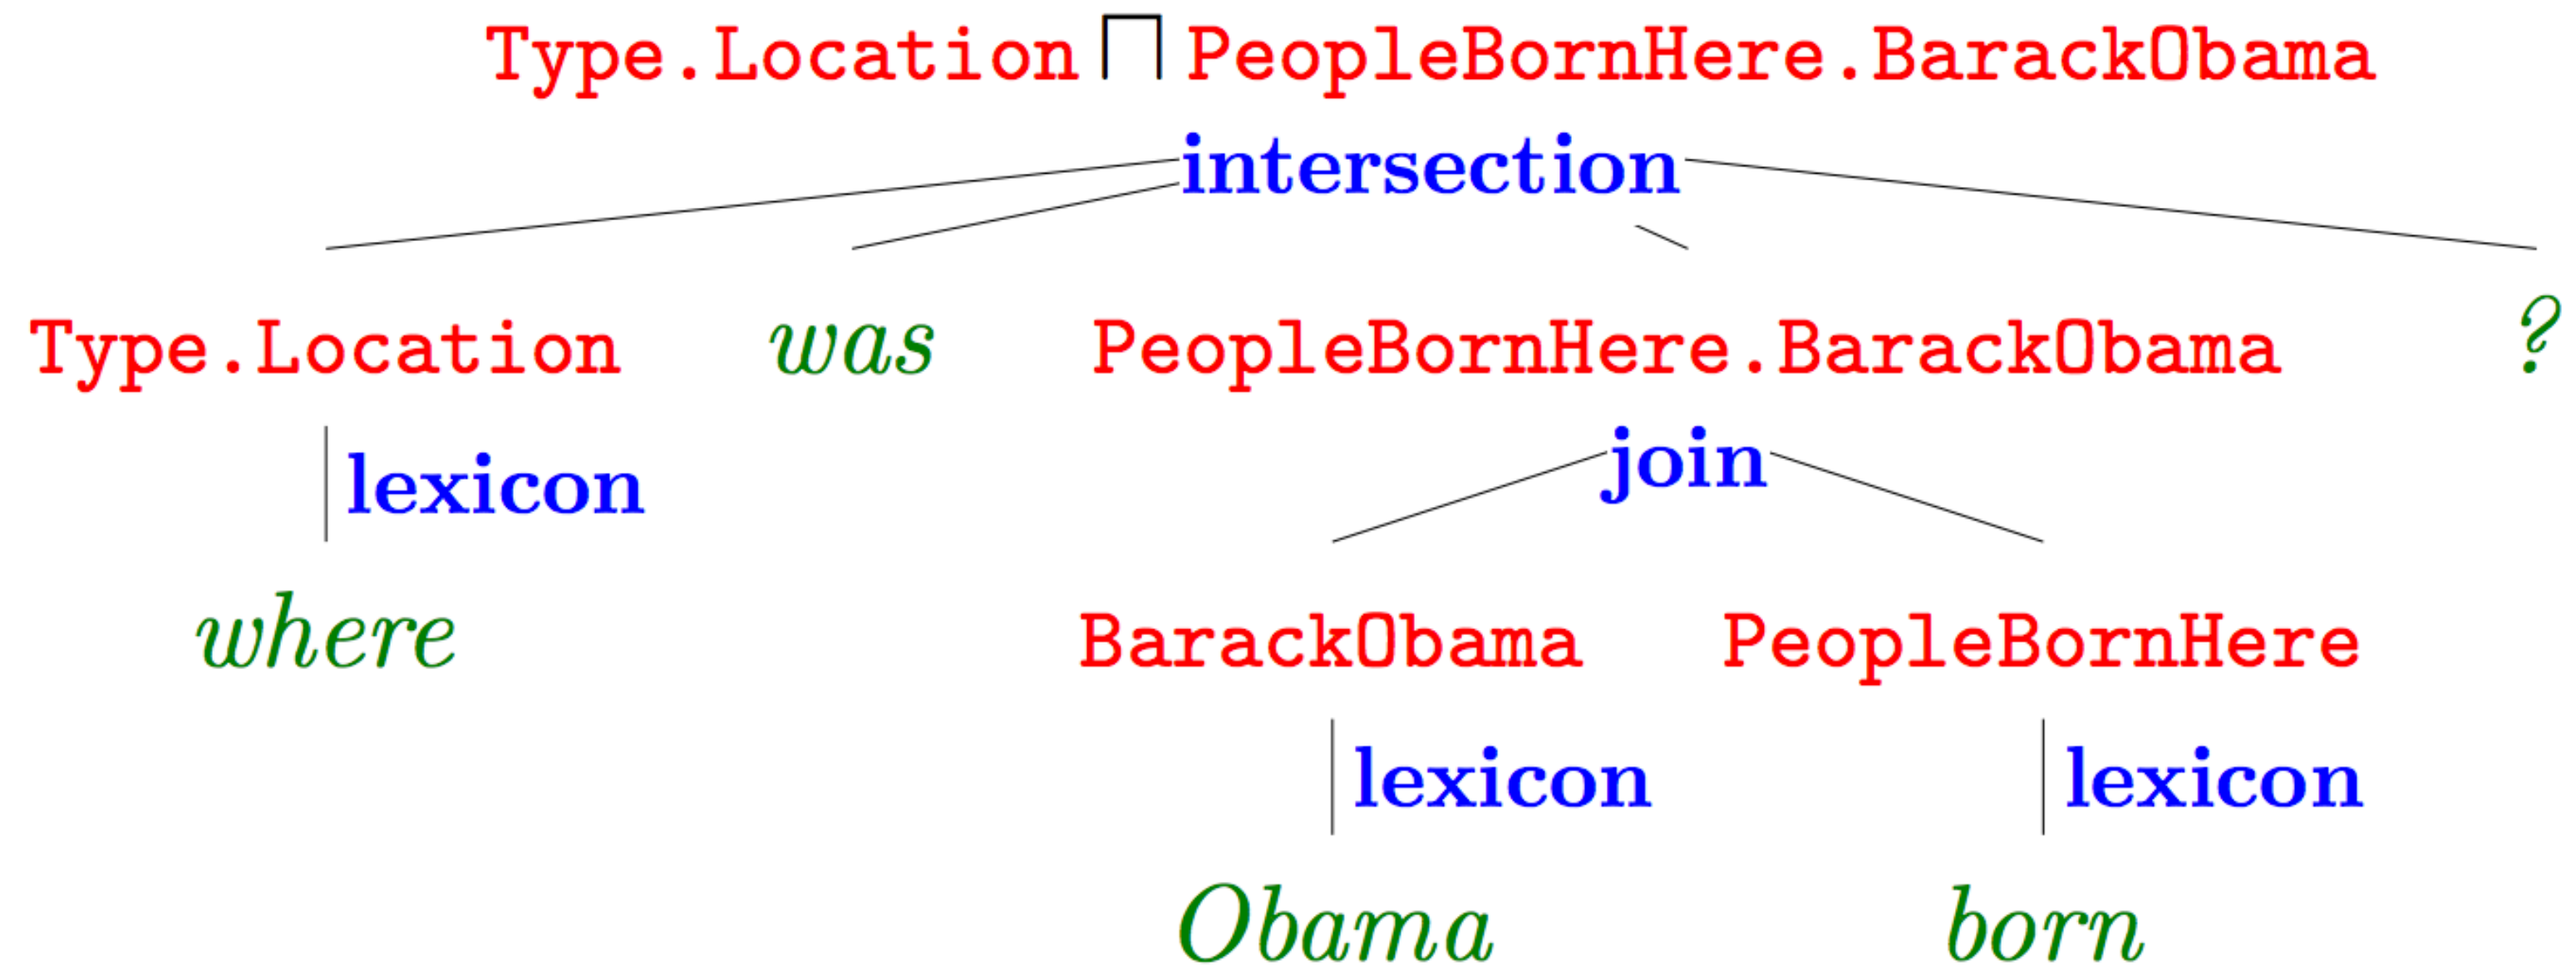
Recall: Semantic Parsing

- ▶ Parsing question *x* into derivation *d*:



Recall: Semantic Parsing

- ▶ Parsing question x into derivation d :



1. What approaches **cannot** be used for this task?
 - (1) Feature extraction + log-linear model
 - (2) Seq2seq model
 - (3) Classification of x with fixed categories

Recall: SQuAD dataset

One of the most famous people born in Warsaw was Maria Skłodowska-Curie, who achieved international recognition for her research on radioactivity and was the first female recipient of the Nobel Prize. Famous musicians include Władysław Szpilman and Frédéric Chopin. Though Chopin was born in the village of Żelazowa Wola, about 60 km (37 mi) from Warsaw, he moved to the city with his family when he was seven months old. Casimir Pulaski, a Polish general and hero of the American Revolutionary War, was born here in 1745.

What was Maria Curie the first female recipient of?

Ground Truth Answers: Nobel Prize

2. This dataset is sometimes regarded as *not* realistic. Why?

- (1) Questions are synthesized by computers
- (2) Answers are not accurate
- (3) It cannot be used to develop information retrieval models
- (4) Humans created questions when looking at the paragraph

Recall: Multi-hop dataset HotpotQA

Question: *What government position was held by the woman who portrayed **Corliss Archer** in the film *Kiss and Tell* ?*

Example picked from HotpotQA [Yang et al., 2018]

Doc 1 *Shirley Temple* Black was an American actress, businesswoman, and singer ...
As an adult, *she* served as *Chief of Protocol* of the United States

Same entity

...

Same entity

Doc 2 *Kiss and Tell* is a comedy film in which 17-year-old *Shirley Temple* acts as **Corliss Archer**.

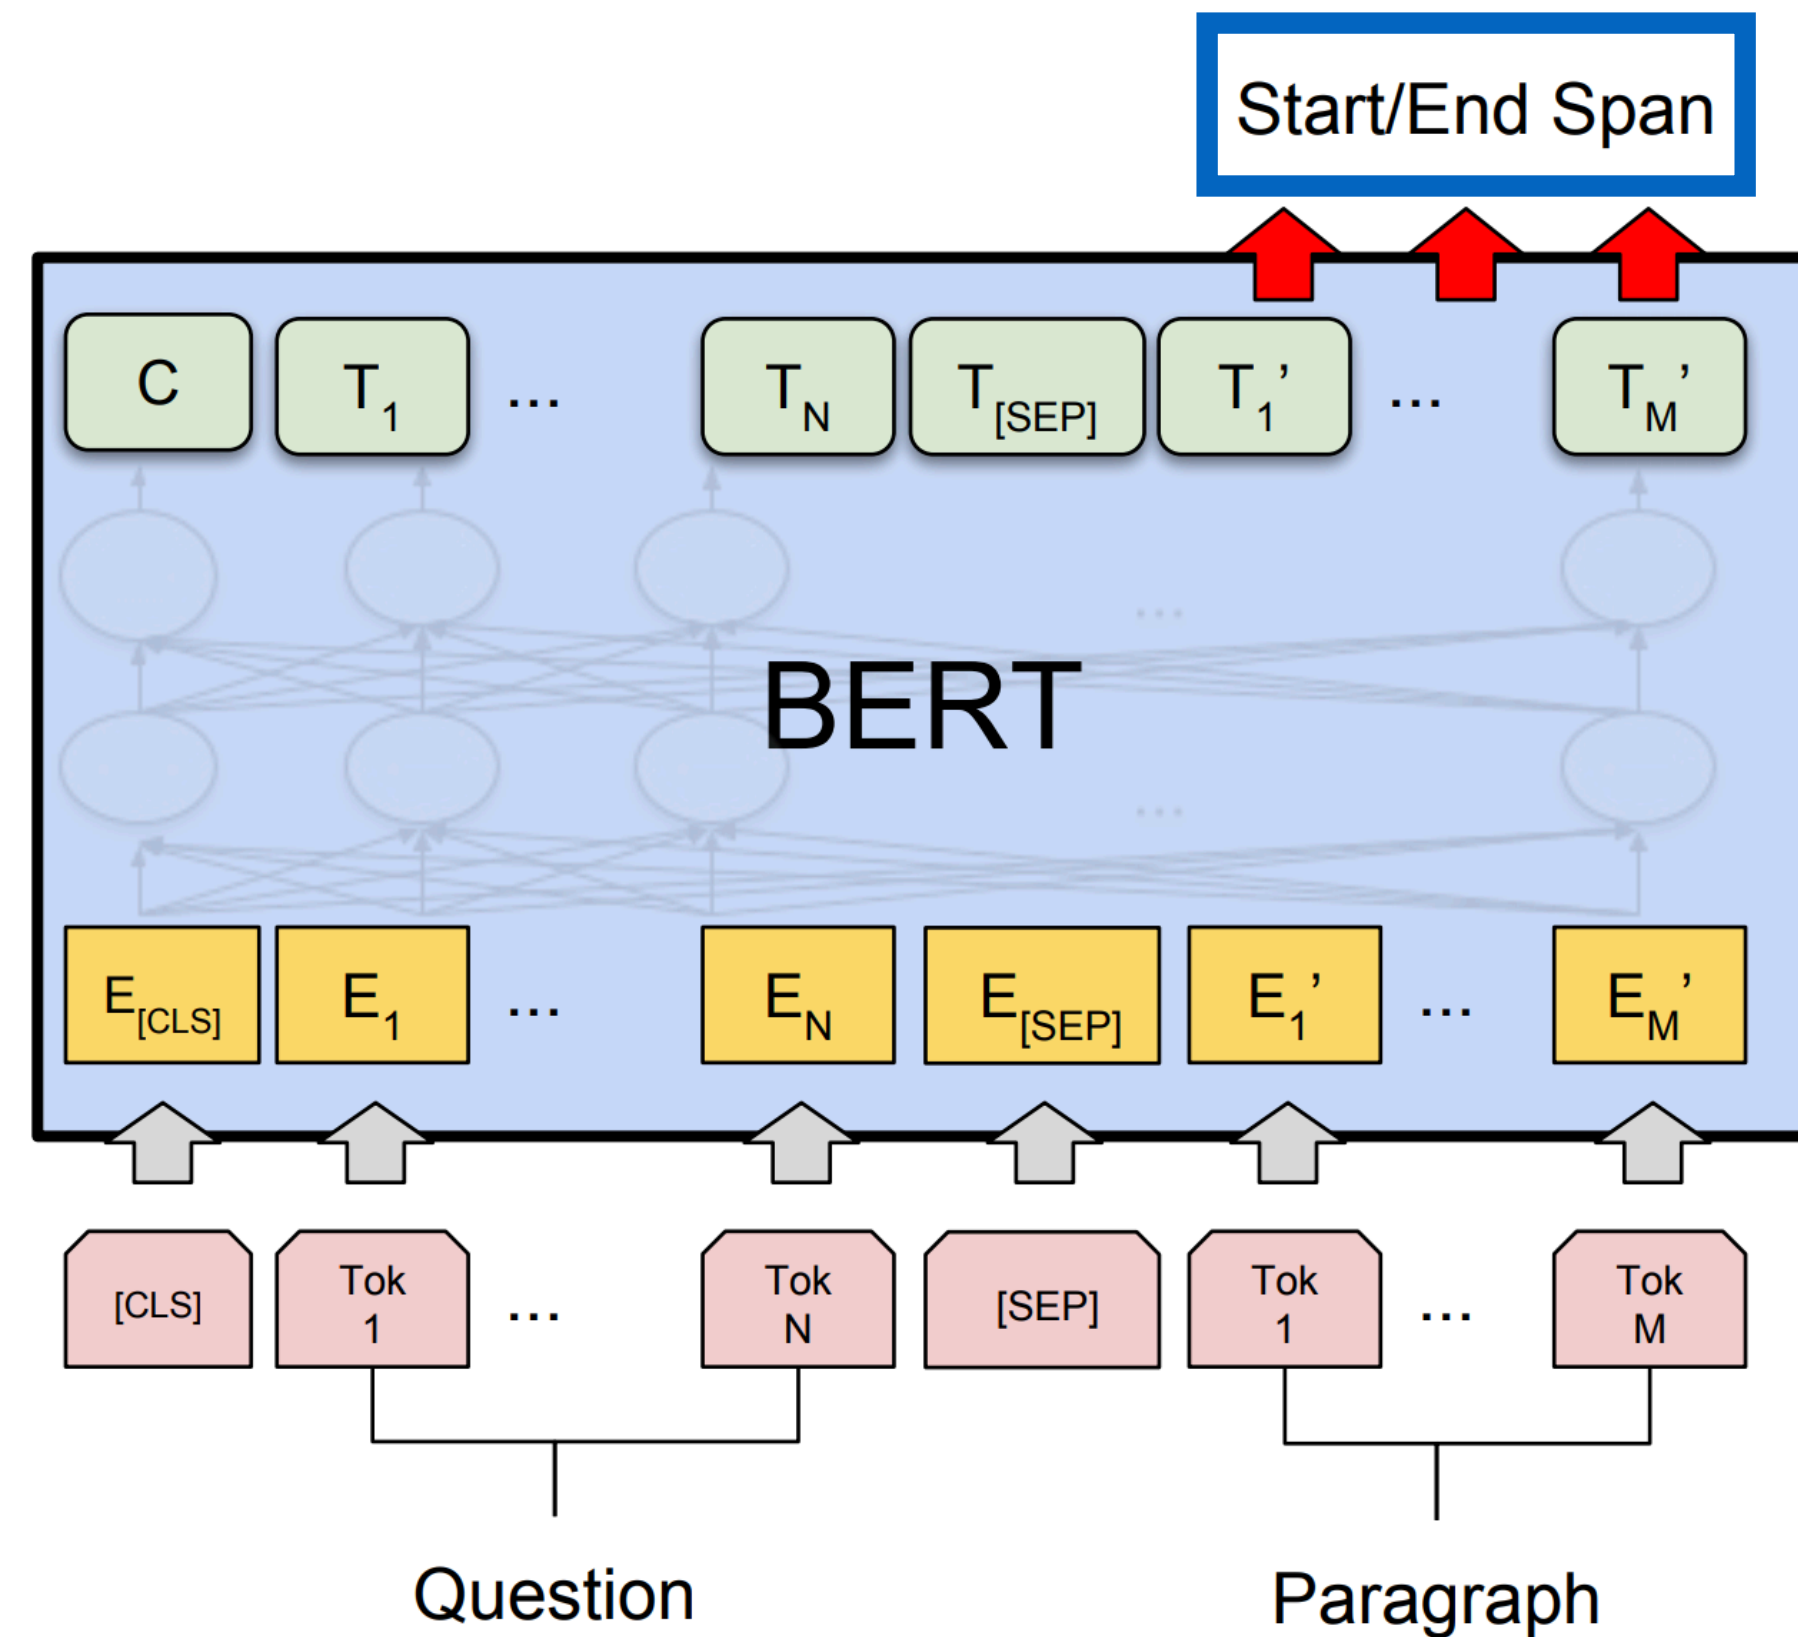
...

Doc 3 *Meet Corliss Archer* is an American television sitcom that aired on CBS ...

3. Which one below is **NOT** a drawback of this dataset?

- (1) High lexical word overlap between a question and a single Doc
- (2) Questions are too long and may involve multiple relations.
- (3) Answer spans may be found without context considered

Recall: BERT model for QA



4. Which QA dataset(s) is the above formulation more often used for?

- (1) SQuAD
- (2) HotpotQA
- (3) GeoQuery

Dialogue

Search/QA as Dialogue

- ▶ Dialogue is a very natural way to find information from a search engine or a QA system

Search/QA as Dialogue



movies chris pratt



All News Images Videos Shopping More Settings Tools

Chris Pratt > Movies



Guardians of
the Galaxy
2014



Jurassic World
2015



Passengers
2016



Guardians of
the Galaxy Vo...
2017

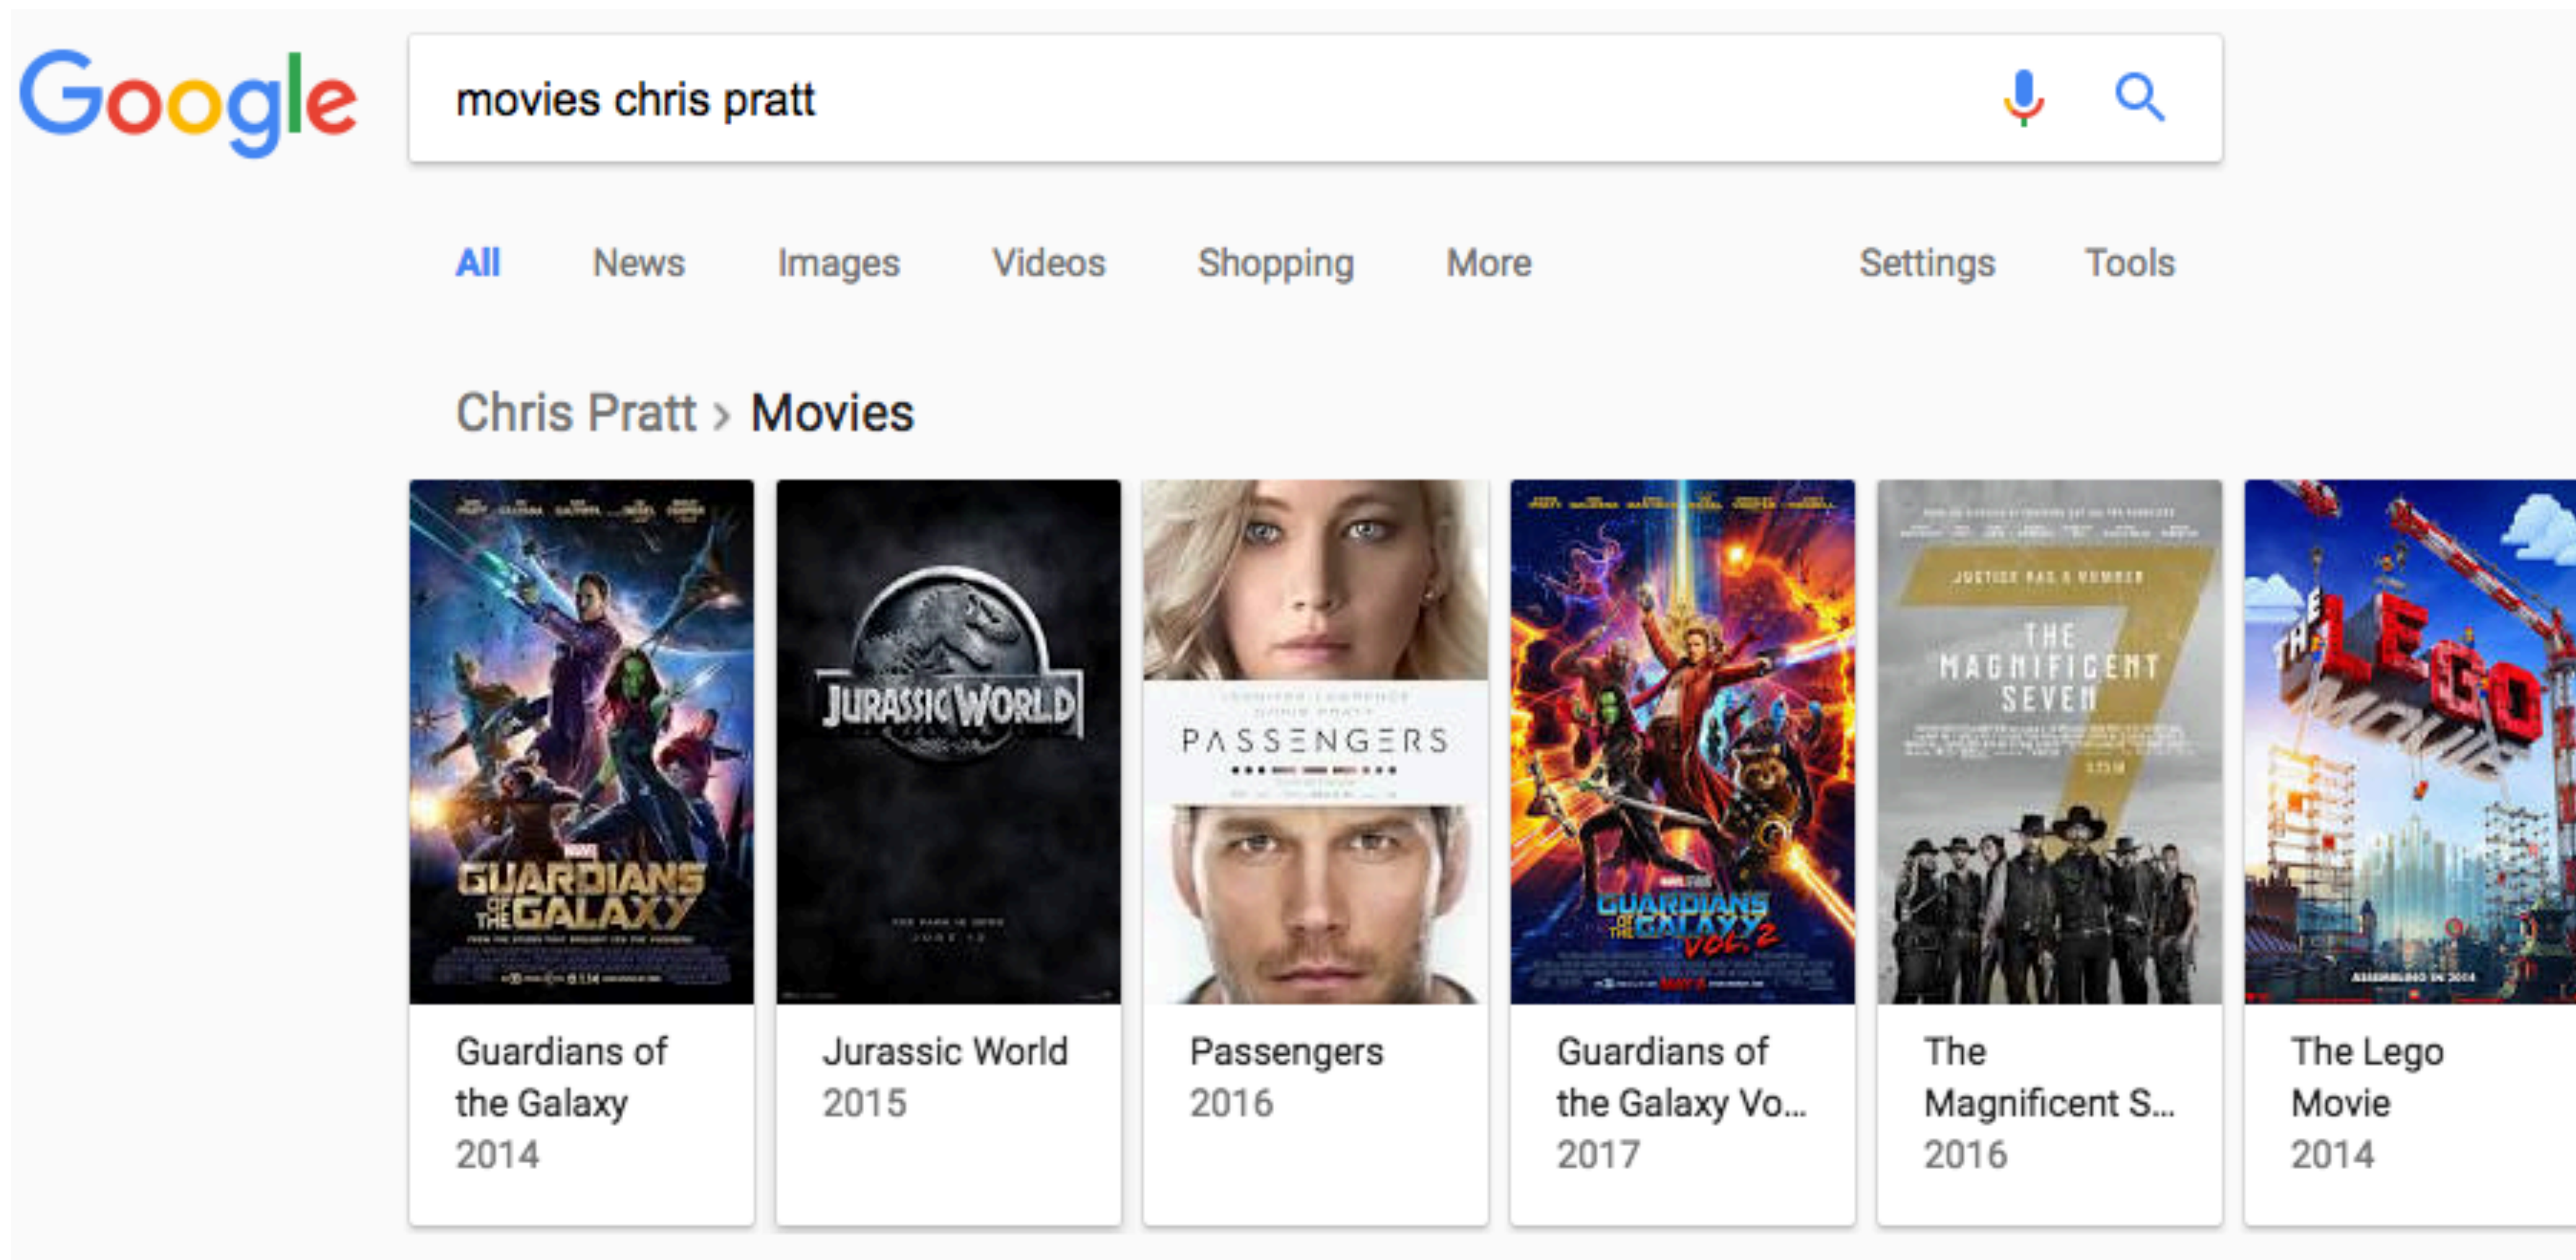


The
Magnificent S...
2016



The Lego
Movie
2014

Search/QA as Dialogue



The image shows a Google search interface. The search bar contains the text "movies chris pratt". Below the search bar, there are navigation tabs for "All", "News", "Images", "Videos", "Shopping", "More", "Settings", and "Tools". The search results are displayed under the heading "Chris Pratt > Movies". There are six movie cards shown, each with a poster and the movie title and year:

- Guardians of the Galaxy 2014
- Jurassic World 2015
- Passengers 2016
- Guardians of the Galaxy Vol. 2 2017
- The Magnificent Seven 2016
- The Lego Movie 2014

- ▶ “Has Chris Pratt won an Oscar?” / “Has *he* won an Oscar”

QA as Dialogue

- ▶ Dialogue is a very natural way to find information from a search engine or a QA system

Original intent:
What super hero from Earth appeared most recently?

1. Who are all of the super heroes?

2. Which of them come from Earth?

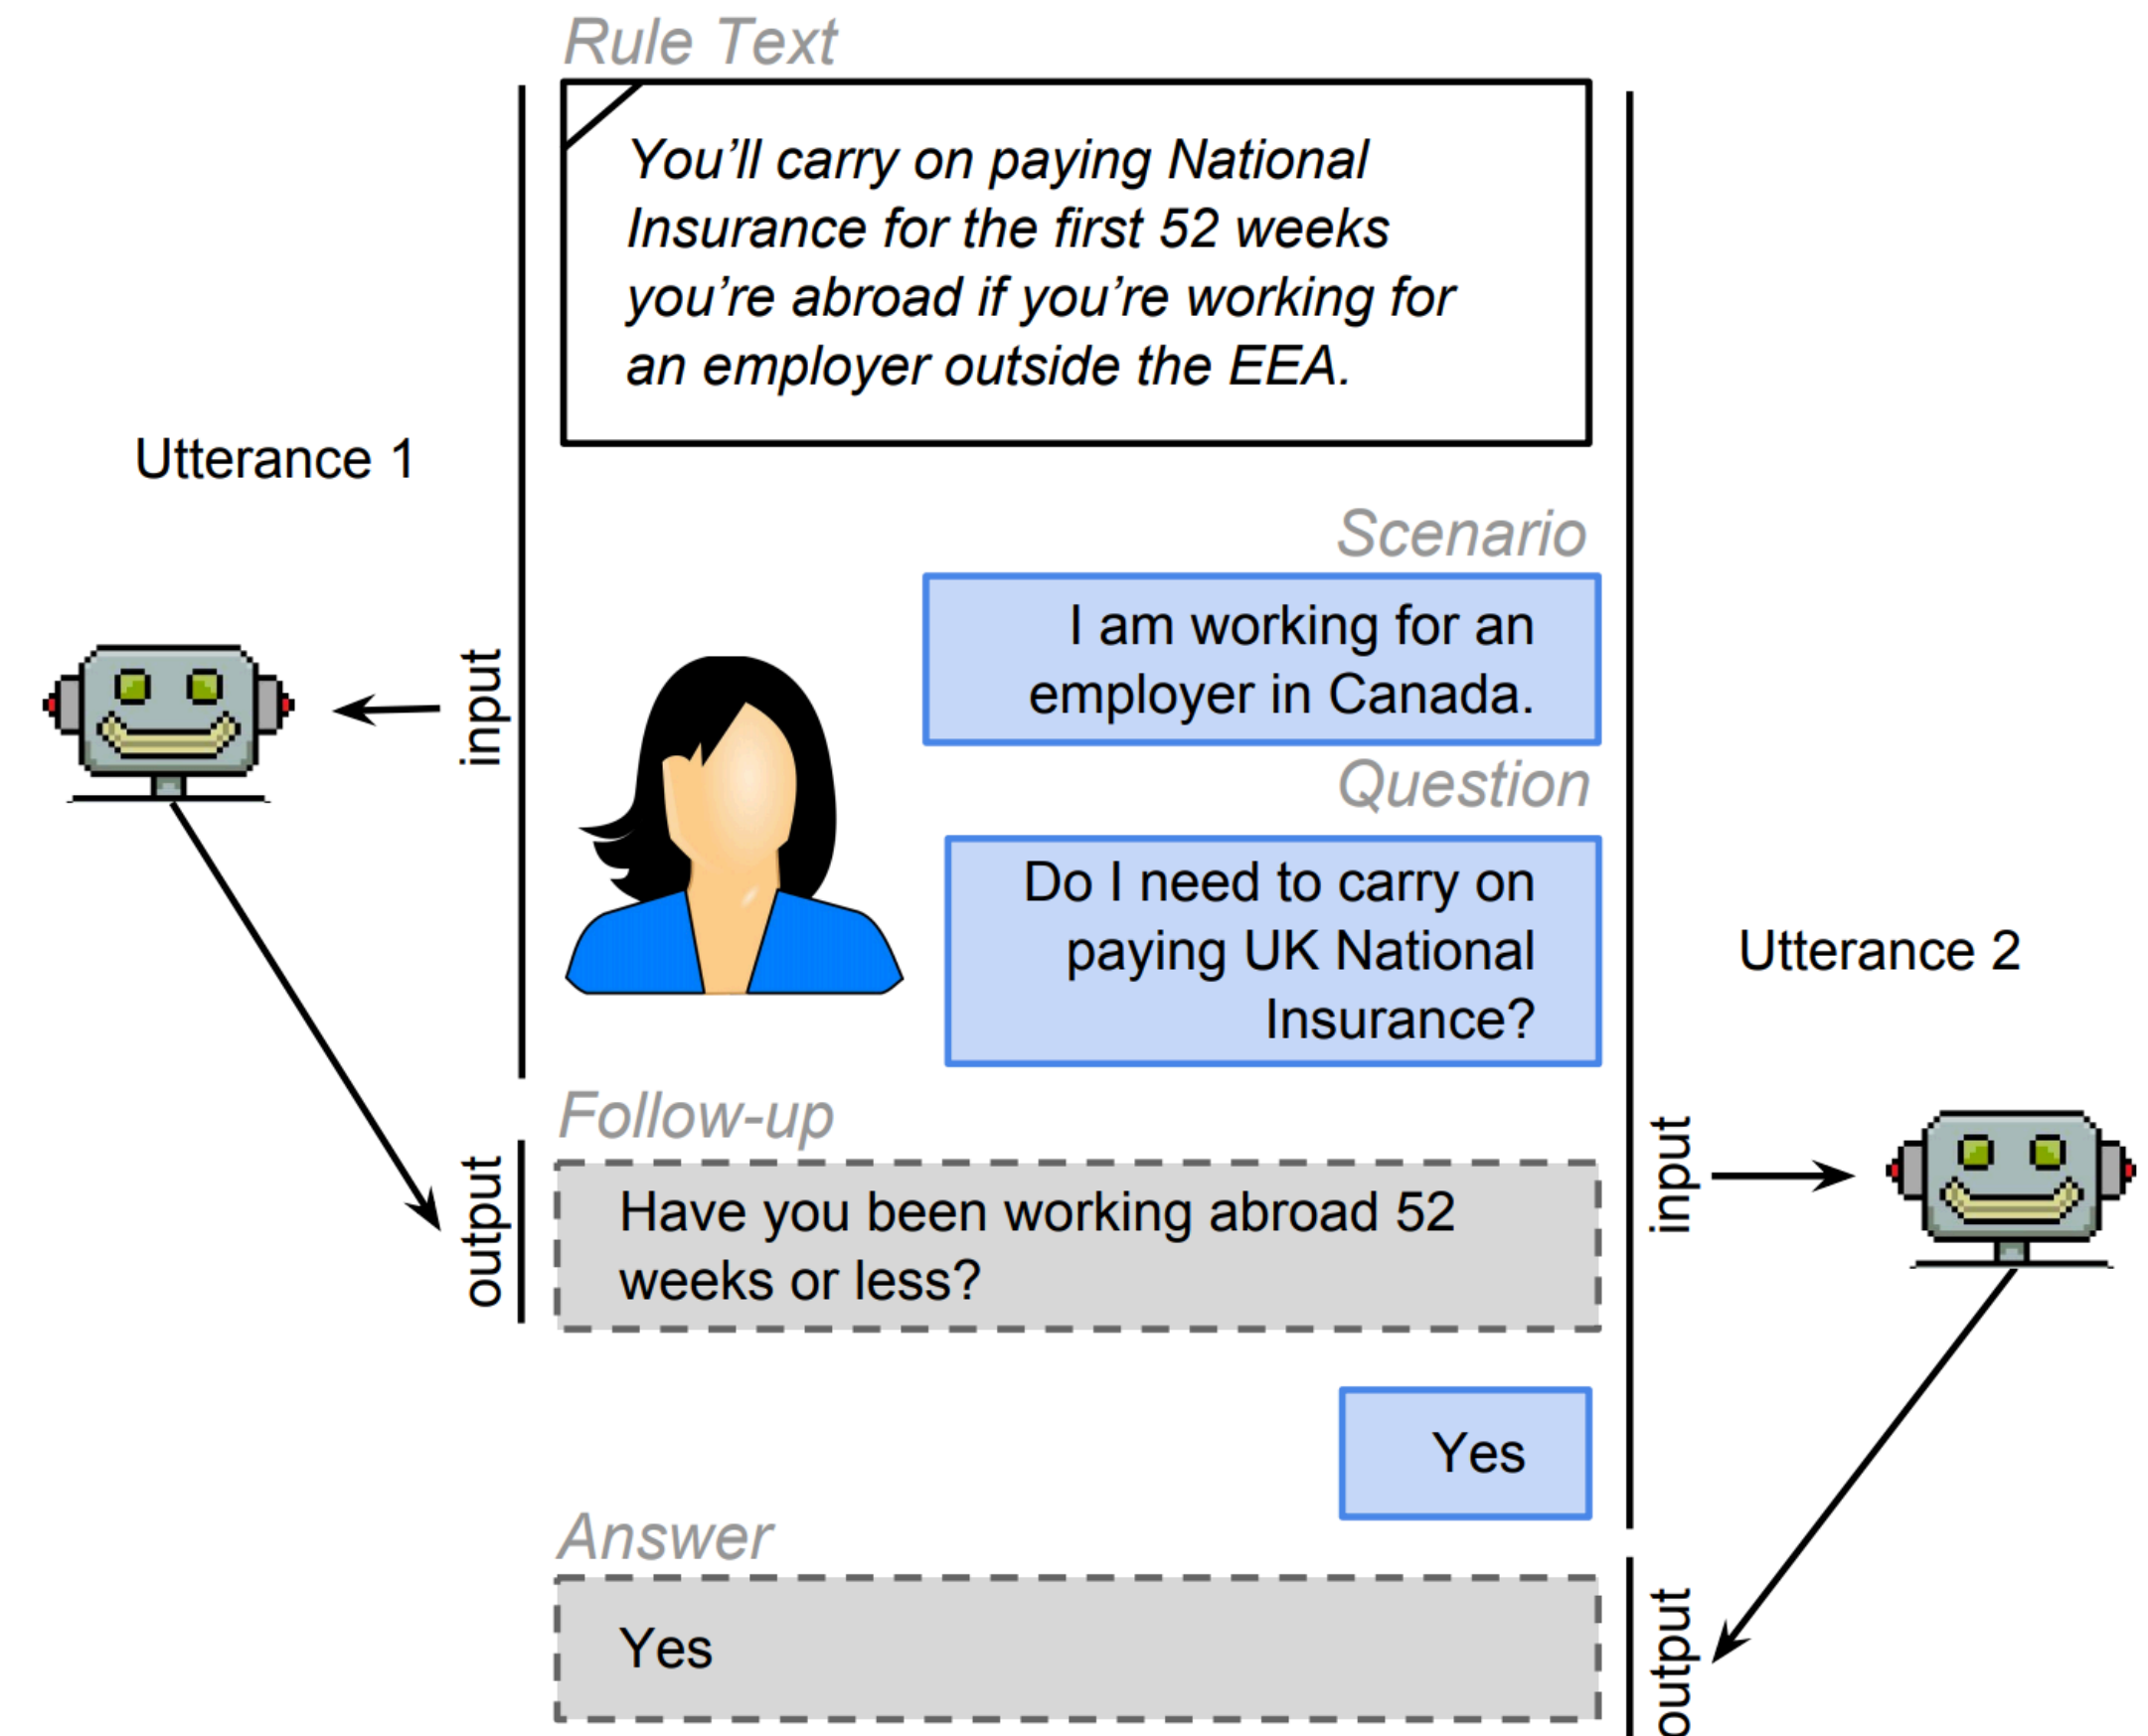
3. Of those, who appeared most recently?

Legion of Super Heroes Post-*Infinite Crisis*

<i>Character</i>	<i>First Appeared</i>	<i>Home World</i>	<i>Powers</i>
Night Girl	2007	Kathoon	Super strength
Dragonwing	2010	Earth	Fire breath
Gates	2009	Vyriga	Teleporting
XS	2009	Aarok	Super speed
Harmonia	2011	Earth	Elemental

QA as Dialogue

- ▶ **Conversational machine reading:** answer questions based on a passage
- ▶ Interesting and potentially useful idea



This Lecture

- ▶ QA as dialogue
- ▶ Classic chatbots
- ▶ Neural chatbots
- ▶ Task-oriented dialogue

Chatbots

ELIZA

- ▶ Created 1964-1966 at MIT, heavily scripted

You are like my father in some ways.
WHAT RESEMBLANCE DO YOU SEE
You are not very aggressive but I think you don't want me to notice that.
WHAT MAKES YOU THINK I AM NOT VERY AGGRESSIVE
You don't argue with me.
WHY DO YOU THINK I DON'T ARGUE WITH YOU
You are afraid of me.
DOES IT PLEASE YOU TO BELIEVE I AM AFRAID OF YOU
My father is afraid of everybody.
WHAT ELSE COMES TO MIND WHEN YOU THINK OF YOUR FATHER
Bullies.
DOES THAT HAVE ANYTHING TO DO WITH THE FACT THAT YOUR BOYFRIEND MADE YOU COME HERE

Eliza responds like a psychotherapist

Weizenbaum (1966)

ELIZA

- ▶ Created 1964-1966 at MIT, heavily scripted



You are like my father in some ways.
WHAT RESEMBLANCE DO YOU SEE
You are not very aggressive but I think you don't want me to notice that.
WHAT MAKES YOU THINK I AM NOT VERY AGGRESSIVE
You don't argue with me.
WHY DO YOU THINK I DON'T ARGUE WITH YOU
You are afraid of me.
DOES IT PLEASE YOU TO BELIEVE I AM AFRAID OF YOU
My father is afraid of everybody.
WHAT ELSE COMES TO MIND WHEN YOU THINK OF YOUR FATHER
Bullies.
DOES THAT HAVE ANYTHING TO DO WITH THE FACT THAT YOUR BOYFRIEND MADE YOU COME HERE

Featured in *Young Sheldon*: <https://www.youtube.com/watch?v=2V9skeVray4>

ELIZA

- ▶ Pattern matching and substitution, using **transformation rule**

(.*) you (.*) me



Why do you think I \$2 you?

You don't argue with me.

WHY DO YOU THINK I DON'T ARGUE WITH YOU

ELIZA

- ▶ Pattern matching and substitution, using **transformation rule**

My (.) (.*)



What else comes to mind when you think of your \$1?

My father is afraid of everybody.

WHAT ELSE COMES TO MIND WHEN YOU THINK OF
YOUR FATHER

ELIZA

- ▶ Identify keyword, identify context, apply **transformation rule**

(.*) you (.*) me



Why do you think I \$2 you?

My (.) (.*)



What else comes to mind when you think of your \$1?

You don't argue with me.

WHY DO YOU THINK I DON'T ARGUE WITH YOU

- ▶ Very little need to generate new content, but can only have one type of conversation; not coherent

Cleverbot

- ▶ Carpenter (1986), online system built in 2006

Cleverbot

- ▶ “Nearest neighbors”: A user says statement A -> Cleverbot finds a human response to statement A in human-human or human-computer chats -> repeat that

User: Hi there cleverbot.

Cleverbot: How are you?

User: Doing okay. what plans do you have today?

Cleverbot: I do revision.

User: What are you revising for?

Cleverbot: Maths history.

User: How old are you?

Cleverbot: That's none of your business.

User: What is maths history?

Cleverbot: I don't know, farming maybe

Check more:

<https://www.existor.com/products/cleverbot-data-for-machine-learning/>

Cleverbot

- ▶ Carpenter (1986), online system built in 2006
- ▶ “Nearest neighbors”: human says statement A, find a human response in human-human or human-computer chats to statement A, repeat that

User: Hi there cleverbot.

Cleverbot: How are you?

User: Doing okay. what plans do you have today?

Cleverbot: I do revision.

User: What are you revising for?

Cleverbot: Maths history.

User: How old are you?

Cleverbot: That's none of your business.

User: What is maths history?

Cleverbot: I don't know, farming maybe

- ▶ Can often give sensible answers, but the bot doesn't really impose high-level discourse structure

Data-Driven Approaches

Stimulus: I'm slowly making this soup
..... and it smells gorgeous!

Response: I'll bet it looks delicious too!
Haha

on Twitter

Data-Driven Approaches

- ▶ Can treat as a machine translation problem: “translate” from current utterance to next one

Stimulus: I’m slowly making this soup
..... and it smells gorgeous!

Response: I’ll bet it looks delicious too!
Haha

Observation: structural relationship between Stimulus and Response

Data-Driven Approaches

- ▶ Can treat as a machine translation problem: “translate” from current utterance to next one

Stimulus: I’m slowly making this soup
..... and it smells gorgeous!

Response: I’ll bet it looks delicious too!
Haha

Observation: structural relationship between Stimulus and Response

What words/phrases are aligned in this example?

Ritter et al. (2011)

Data-Driven Approaches

- ▶ Can treat as a machine translation problem: “translate” from current utterance to next one

what	.	.	.	■	■	-	name	■	■	■	.
time	.	.	.	■	■	3	kevin	■	■	■	.
u	■	costner	movies	■	■	■	.
get	.	■	.	.	.	that	dont	■	■	■	.
out	.	.	■	.	.	suck	.	■	■	■	.
?	■
	i	get	off	at	5		easier	.	question	.	.
							please

- ▶ How to decide larger phrase (not just word-word) alignment?

Data-Driven Approaches

- ▶ Can treat as a machine translation problem: “translate” from current utterance to next one

what	.	.	.	■	■	-	name	■	■	■	.
time	.	.	.	■	■	3	kevin	■	■	■	.
u	■	costner	movies	■	■	■	.
get	.	■	.	.	.	that	dont	■	■	■	.
out	.	.	■	.	.	suck	.	■	■	■	.
?	■
	i	get	off	at	5		easier	question	please	.	.

- ▶ How to decide larger phrase (not just word-word) alignment?
 - ▶ use statistical measures to filter out phrase pairs with low correlation

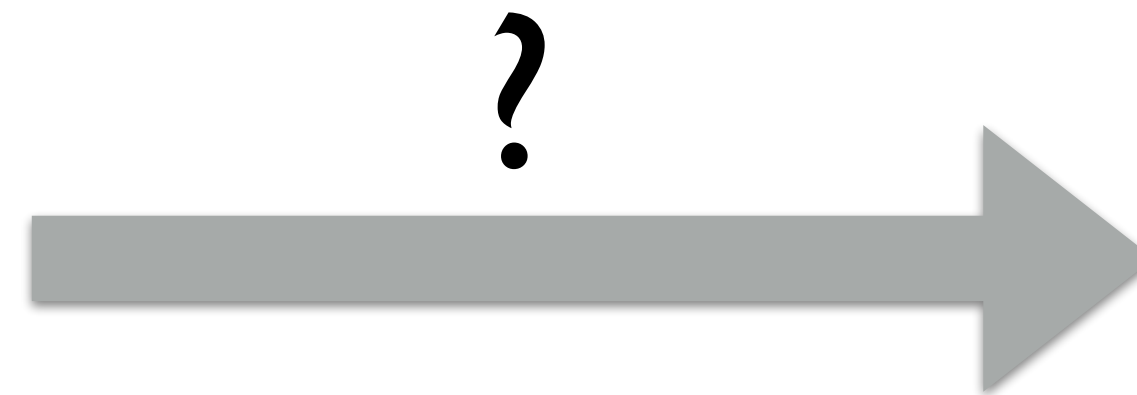
Data-Driven Approaches

Status	MT-CHAT	MT-BASELINE	HUMAN
riding to the arc on my bike. Hopefully I won't melt into the pavement.	good luck on your bike ride arc	i ca n't wait for it ! i 'm arc is a year .	make sure to ride in the middle of the street
I really hate watching movies in the living room.	me too , what movies are you doing in room ?	me too ! i love those ?) in my mouth .	Why? no comfy spots or just too open to public
Hello Did I miss David Getting his teen choice award? Or did they not show it :(hello miss ? * he was getting his will probably win . <u>proposed method</u>	it 's just like what you just did the exact date hello ? <u>phrase alignment is weaker</u>	nothing yet...he presented though he is so darn cute

Ritter et al. (2011)

What other approaches can you think of?

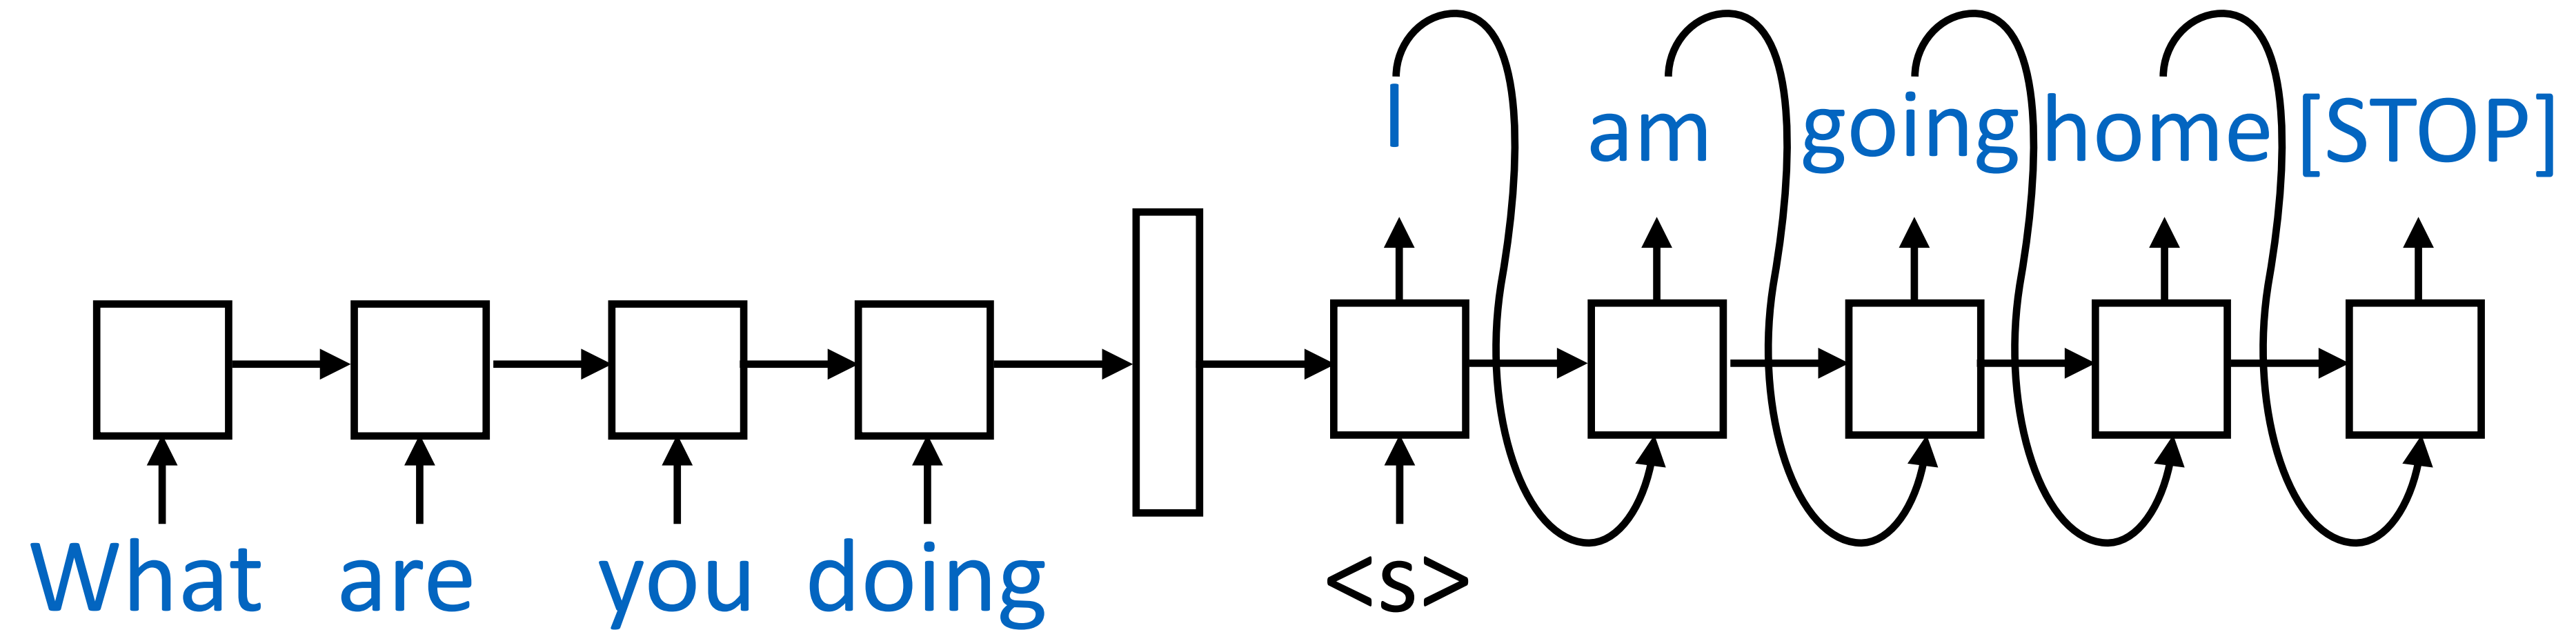
Status
riding to the arc on my bike. Hopefully I won't melt into the pavement.
I really hate watching movies in the living room.
Hello Did I miss David Getting his teen choice award? Or did they not show it :(



HUMAN
make sure to ride in the middle of the street
Why? no comfy spots or just too open to public
nothing yet...he presented though he is so darn cute
Ritter et al. (2011)

Neural Chatbots

Seq2seq models



- ▶ Just like the machine translation task mentioned earlier
- ▶ can train seq2seq models for this task

Subtitles Data for Model Development

do you want to meet your sponsor for the last 10 years ?
of course ! but he doesn't want to see me !

and where had you been just before ?
i' d been to the palace of the legion of honor , the art gallery .

yeah , we were just going to hit up taco bell .
well , it' s my pleasure .

Subtitles Data

do you want to meet your sponsor for the last 10 years ?
of course ! but he doesn't want to see me !

and where had you been just before ?
i' d been to the palace of the legion of honor , the art gallery .

yeah , we were just going to hit up taco bell .
well , it' s my pleasure .

- ▶ Why might this model perform poorly? What might it be bad at?

Lack of Diversity

► Why?

Input: What are you doing?	
-0.86 I don't know.	-1.09 Get out of here.
-1.03 I don't know!	-1.09 I'm going home.
-1.06 Nothing.	-1.09 Oh my god!
-1.09 Get out of the way.	-1.10 I'm talking to you.
Input: what is your name?	
-0.91 I don't know.	...
-0.92 I don't know!	-1.55 My name is Robert.
-0.92 I don't know, sir.	-1.58 My name is John.
-0.97 Oh, my god!	-1.59 My name's John.
Input: How old are you?	
-0.79 I don't know.	...
-1.06 I'm fine.	-1.64 Twenty-five.
-1.17 I'm all right.	-1.66 Five.
-1.17 I'm not sure.	-1.71 Eight.

Lack of Diversity

- ▶ Training to maximize likelihood gives a system that prefers common responses:

$\max P(\text{response}|\text{input})$

Input: What are you doing?

-0.86	I don't know.	-1.09	Get out of here.
-1.03	I don't know!	-1.09	I'm going home.
-1.06	Nothing.	-1.09	Oh my god!
-1.09	Get out of the way.	-1.10	I'm talking to you.

Input: what is your name?

-0.91	I don't know.	...	
-0.92	I don't know!	-1.55	My name is Robert.
-0.92	I don't know, sir.	-1.58	My name is John.
-0.97	Oh, my god!	-1.59	My name's John.

Input: How old are you?

-0.79	I don't know.	...	
-1.06	I'm fine.	-1.64	Twenty-five.
-1.17	I'm all right.	-1.66	Five.
-1.17	I'm not sure.	-1.71	Eight.

Lack of Diversity

- ▶ Solution: mutual information criterion; response R should be predictive of user utterance U as well
- ▶ Standard conditional likelihood: $\log P(R|U)$
- ▶ Mutual information: $\log \frac{P(R, U)}{P(R)P(U)} = \log P(R|U) - \log P(R)$
- ▶ $\log P(R)$ = probabilities under a language model

We use a generalization of the MMI objective which introduces a hyperparameter λ that controls how much to penalize generic responses:

$$\hat{T} = \arg \max_T \{ \log p(T|S) - \lambda \log p(T) \} \quad (9)$$

An alternate formulation of the MMI objective uses Bayes' theorem:

$$\log p(T) = \log p(T|S) + \log p(S) - \log p(S|T)$$

which lets us rewrite Equation 9 as follows:

$$\begin{aligned} \hat{T} &= \arg \max_T \{ (1 - \lambda) \log p(T|S) \\ &\quad + \lambda \log p(S|T) - \lambda \log p(S) \} \\ &= \arg \max_T \{ (1 - \lambda) \log p(T|S) + \lambda \log p(S|T) \} \end{aligned} \quad (10)$$

Lack of Diversity

(maximum mutual info)

message	SEQ2SEQ	MMI $P(T S) - \lambda P(T)$
I did not get the report from an mi-6 agent.	I don't think that's a good idea.	You did the right thing, did you?
You haven't been given an assignment in this case.	I don't know what you are talking about.	I've been looking all over for you.
I'm losing my grip.	I don't know what you are talking about.	I'm the only one in the world.
I am ready to help.	Come on, come on	I have something we need to talk about.
You programmed me to gather intelligence. That's all I've ever done.	You know that, don't you?	You do have fun, don't you?
I mean, we'd have to talk to him.	I mean, I don't know.	I mean, he's a good guy.
This has nothing to do with him.	You don't know anything about him, do you?	I know him as much.
How come you never say it?	I don't know	Because I don't want to hurt you

► OpenSubtitles data

Another idea: Specificity

- ▶ Train a specificity classifier on labeled data

I don't know => SPECIFICITY=1

Going to the store => SPECIFICITY=3

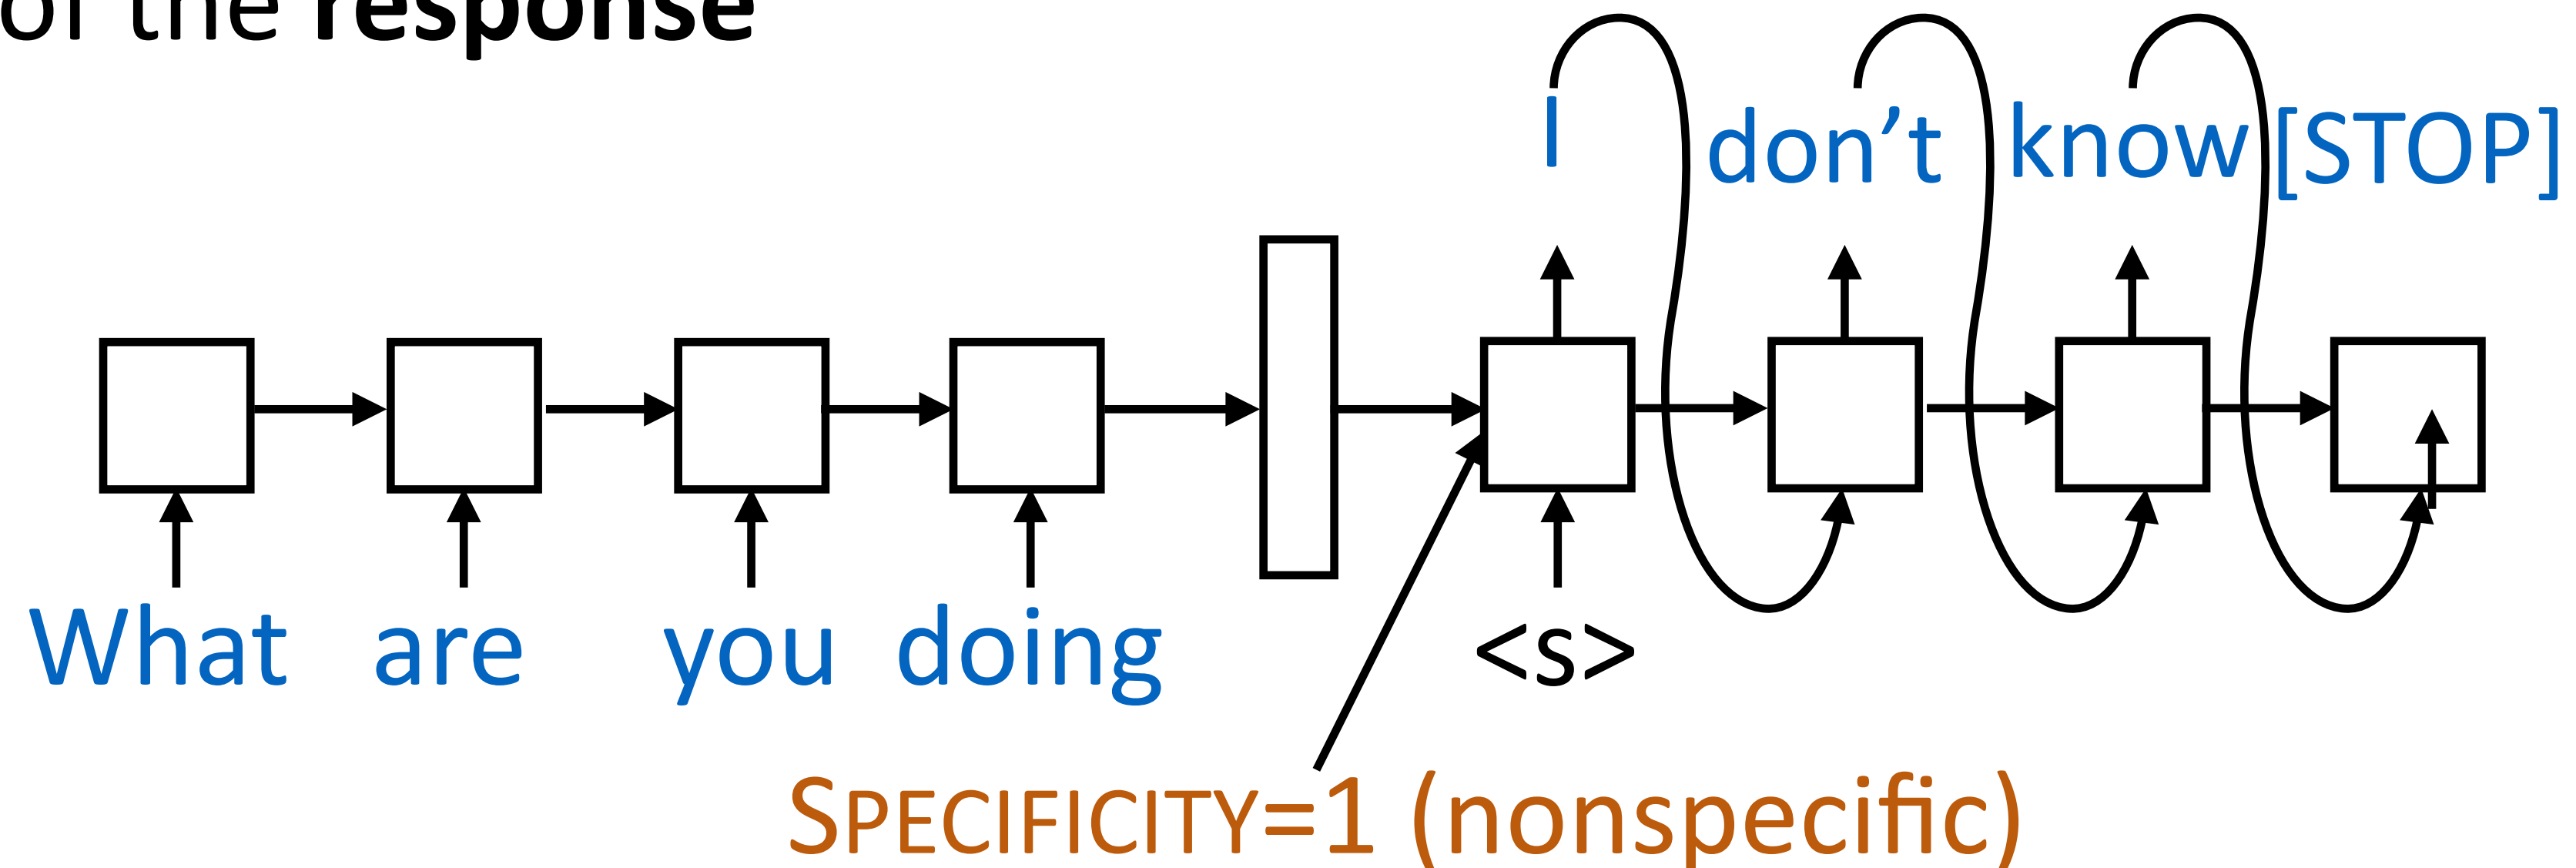
Specificity

- ▶ Train a specificity classifier on labeled data

I don't know => SPECIFICITY=1

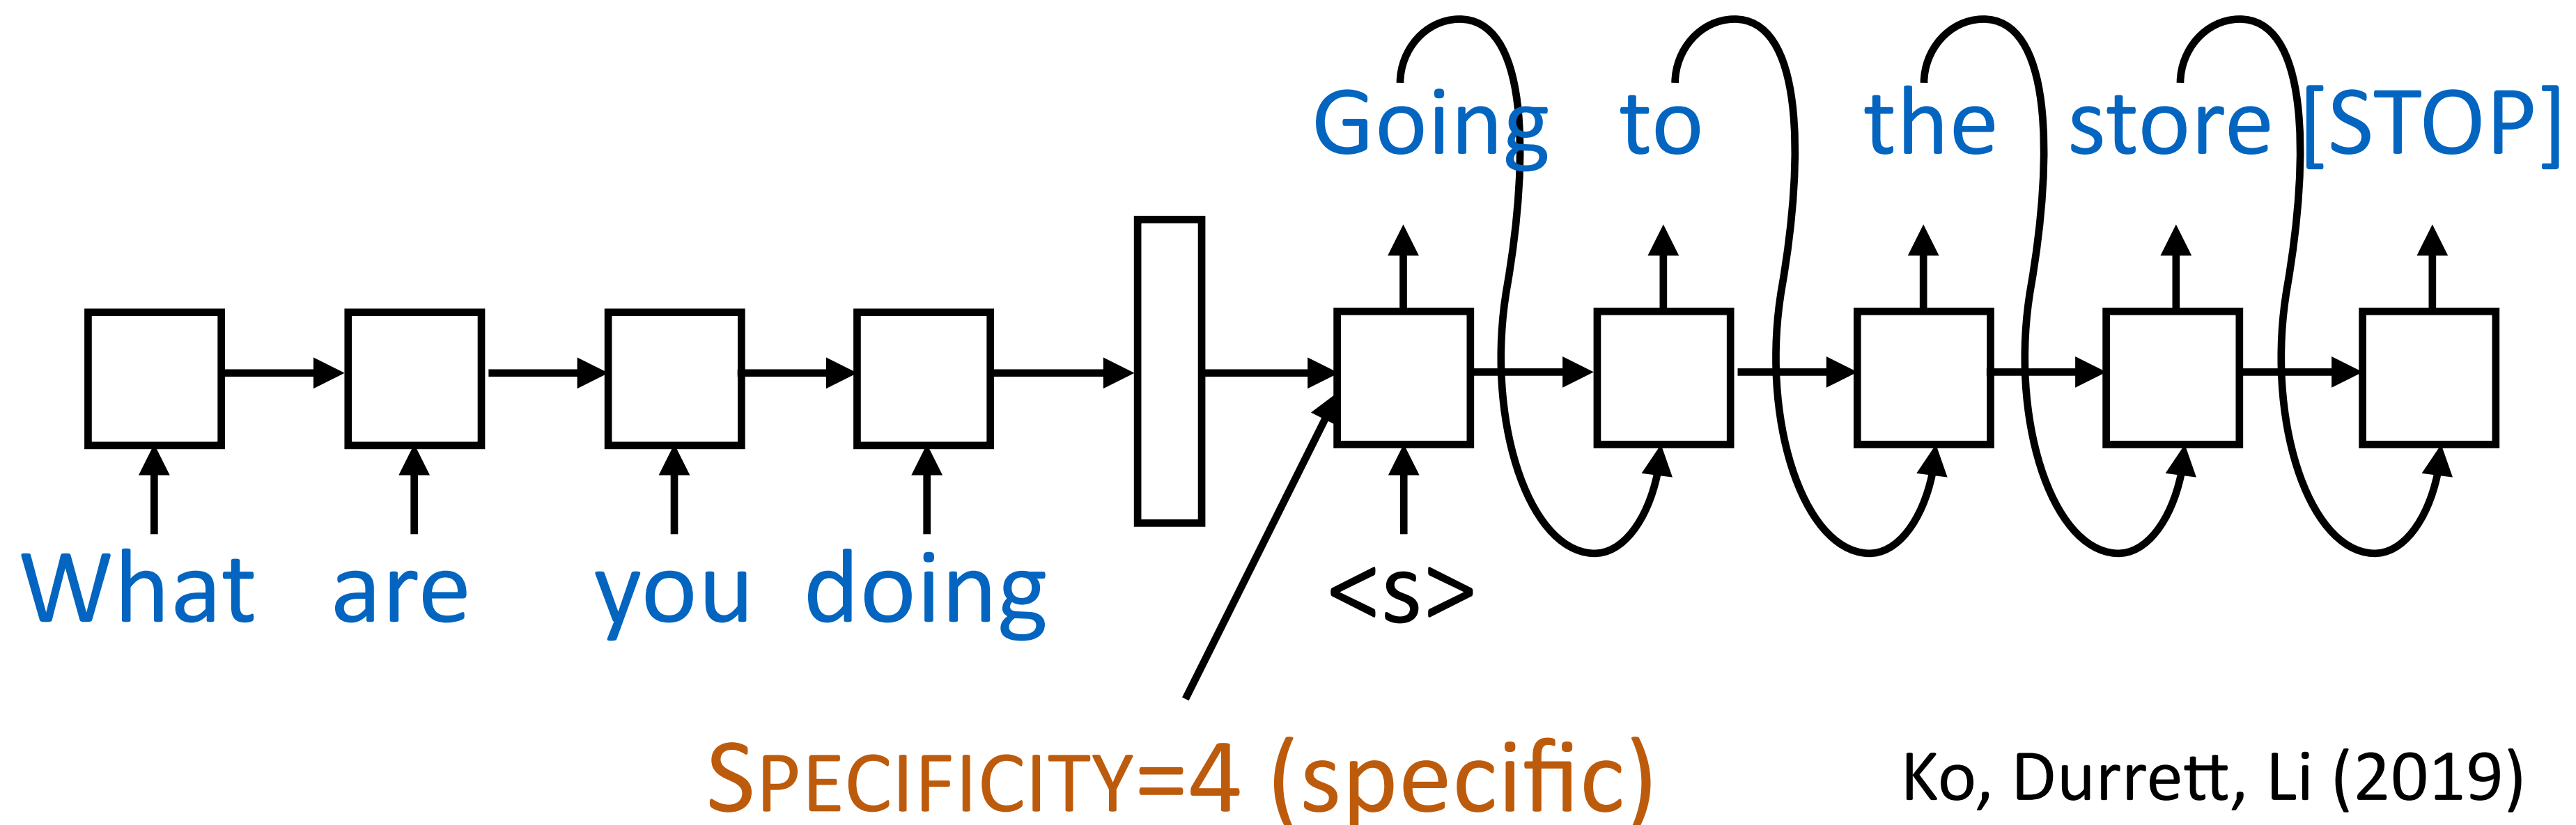
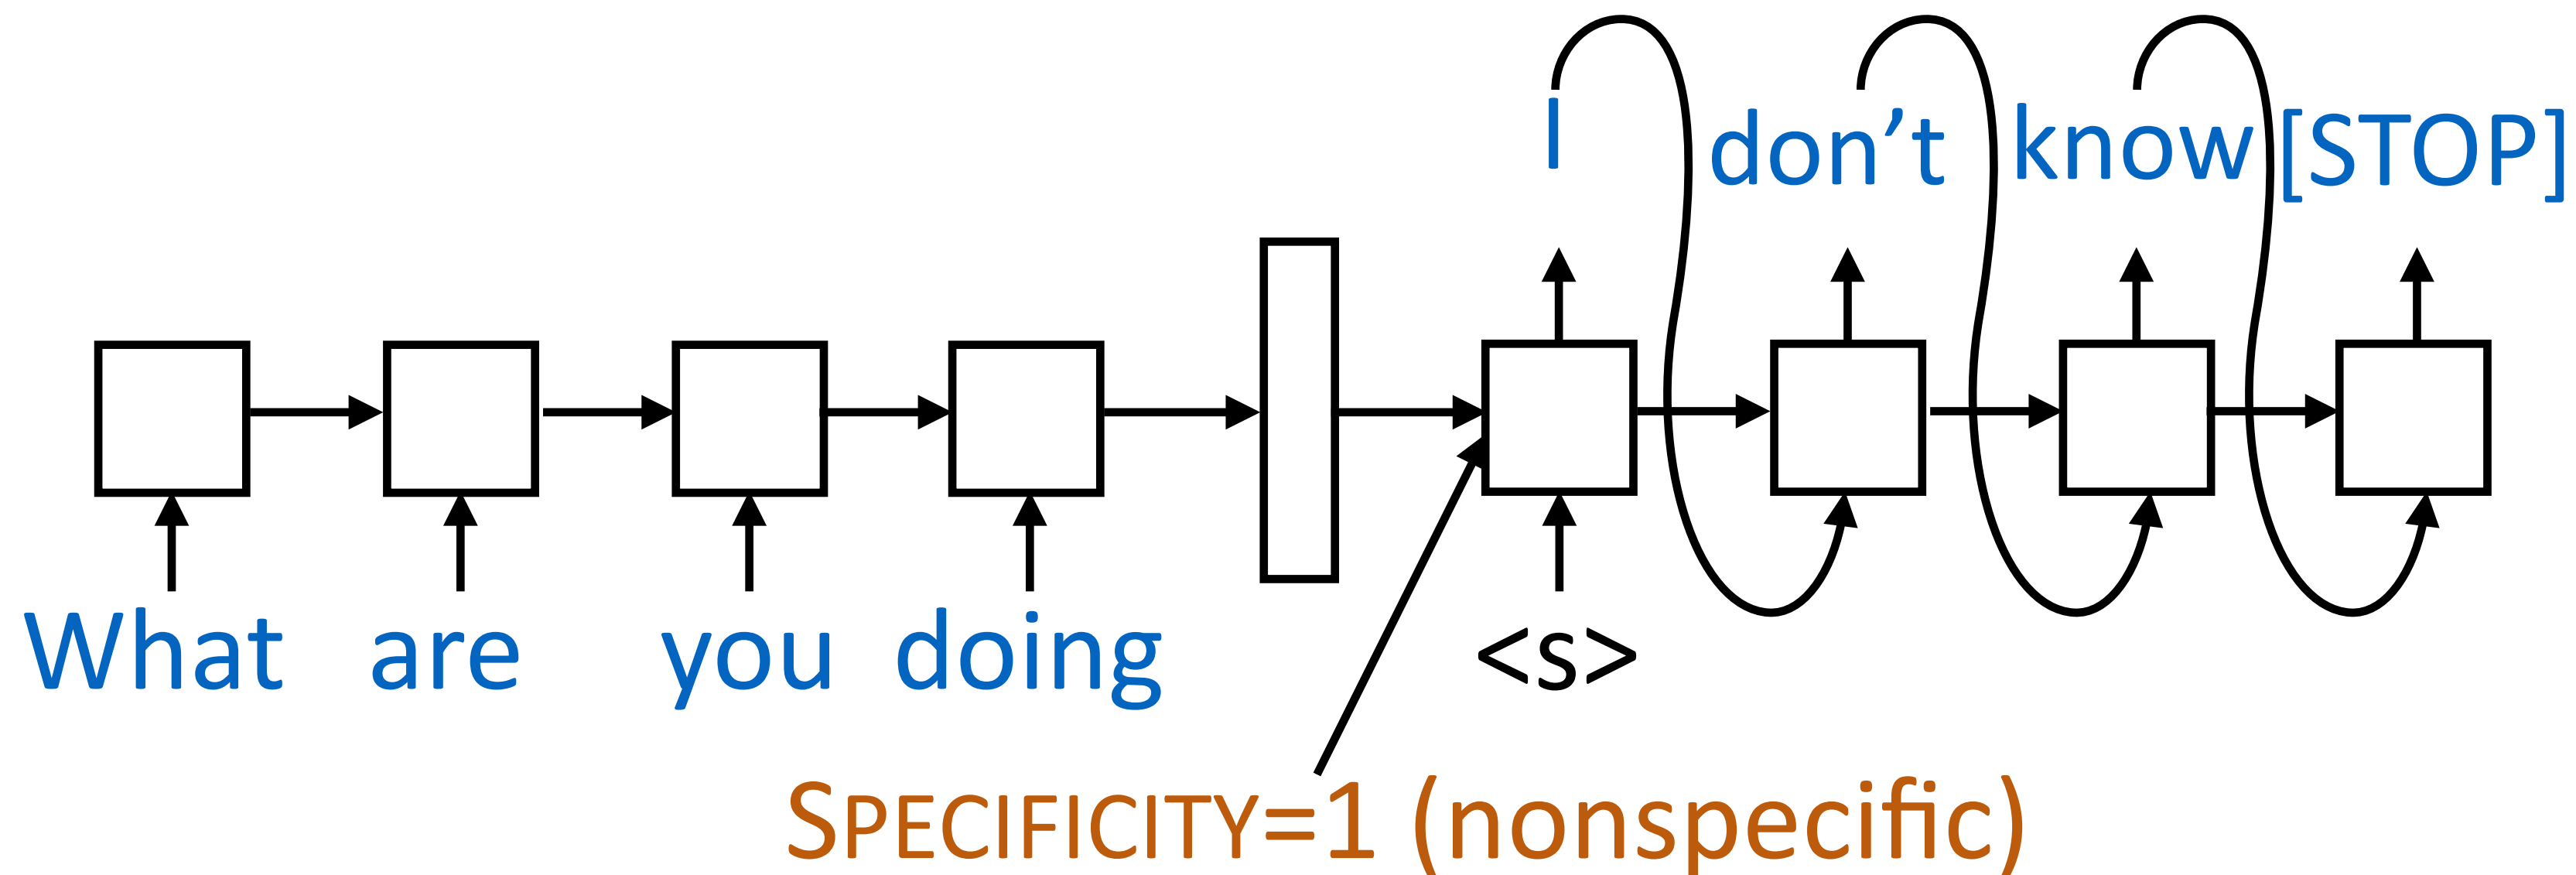
Going to the store => SPECIFICITY=3

- ▶ When training the decoder, condition on the **automatically predicted specificity** of the response



Specificity

- ▶ At test time, set the specificity level higher to get less generic responses



Specificity

Conflicting	i understand. i am not sure if i can afford a babysitter, i am a millionaire
Wrong connective	i am an animal phobic, but i do not like animals
Wrong pronoun	my mom was a social worker, he was an osteopath.
Wrong noun	cool. i work at a non profit organization that sells the holocaust.
Repeating	my favorite food is italian, but i also love italian food, especially italian food.

issues in the generated responses

- ▶ Can use other models to try to fix these issues. But **the facts are still all made up, even if they make sense**

PersonaChat

Persona 1	Persona 2
I like to ski My wife does not like me anymore I have went to Mexico 4 times this year I hate Mexican food I like to eat cheetos	I am an artist I have four children I recently got a cat I enjoy walking for exercise I love watching Game of Thrones

[PERSON 1:] Hi

[PERSON 2:] Hello ! How are you today ?

[PERSON 1:] I am good thank you , how are you.

[PERSON 2:] Great, thanks ! My children and I were just about to watch Game of Thrones

[PERSON 1:] Nice ! How old are your children?

[PERSON 2:] I have four that range in age from 10 to 21. You?

[PERSON 1:] I do not have children at the moment.

[PERSON 2:] That just means you get to keep all the popcorn for yourself.

[PERSON 1:] And Cheetos at the moment!

[PERSON 2:] Good choice. Do you watch Game of Thrones?

[PERSON 1:] No, I do not have much time for TV.

[PERSON 2:] I usually spend my time painting: but, I love the show.

Zhang et al. (2018)

State of Chatbots

- ▶ Can force chatbots to give consistent answers with a persona, but still probably not very interesting
- ▶ “Wizard of Wikipedia:” chatbot that can discuss topics by retrieving from Wikipedia [Dinan et al., 2019]
- ▶ Xiaoice: Microsoft chatbot in Chinese; according to [1], 660 million online users, over 30 billion conversations...people do seem to like talking to them...?



[1] <https://news.microsoft.com/apac/features/much-more-than-a-chatbot-chinas-xiaoice-mixes-ai-with-emotions-and-wins-over-millions-of-fans/>

Task-Oriented Dialogue

Task-Oriented Dialogue

- ▶ Personal assistants / API front-ends:



Siri, find me a good sushi restaurant in Chelsea

Sushi Seki Chelsea is a sushi restaurant in Chelsea with 4.4 stars on Google



Task-Oriented Dialogue

- ▶ Personal assistants / API front-ends:

Siri, find me a good sushi restaurant in Chelsea

Sushi Seki Chelsea is a sushi restaurant in Chelsea with 4.4 stars on Google

How expensive is it?


Entrees are around \$30 each

Find me something cheaper



Task-Oriented Dialogue

▶ Customer service:



Hey Alexa, why isn't my Amazon order here?

Let me retrieve your order.
Your order was scheduled to arrive at 4pm today.

It never came

Okay, I can put you through to customer service.



Air Travel Information Service (ATIS)

- ▶ Given an utterance, predict a domain-specific semantic interpretation

Utterance	<i>How much is the cheapest flight from Boston to New York tomorrow morning?</i>
Goal:	<i>Airfare</i>
Cost_Relative	<i>cheapest</i>
Depart_City	<i>Boston</i>
Arrival_City	<i>New York</i>
Depart_Date.Relative	<i>tomorrow</i>
Depart_Time.Period	<i>morning</i>

- ▶ Can formulate as semantic parsing, but **simple “slot-filling” solutions** (classifiers) work well too

DARPA (early 1990s), Figure from Tur et al. (2010)

Intents

▶ 29 different intents

which flights go from cleveland to indianapolis on april fifth

Intent: flight

does tacoma airport offer transportation from the airport to the downtown area

Intent: ground_service

what days of the week do flights from san jose to nashville fly on

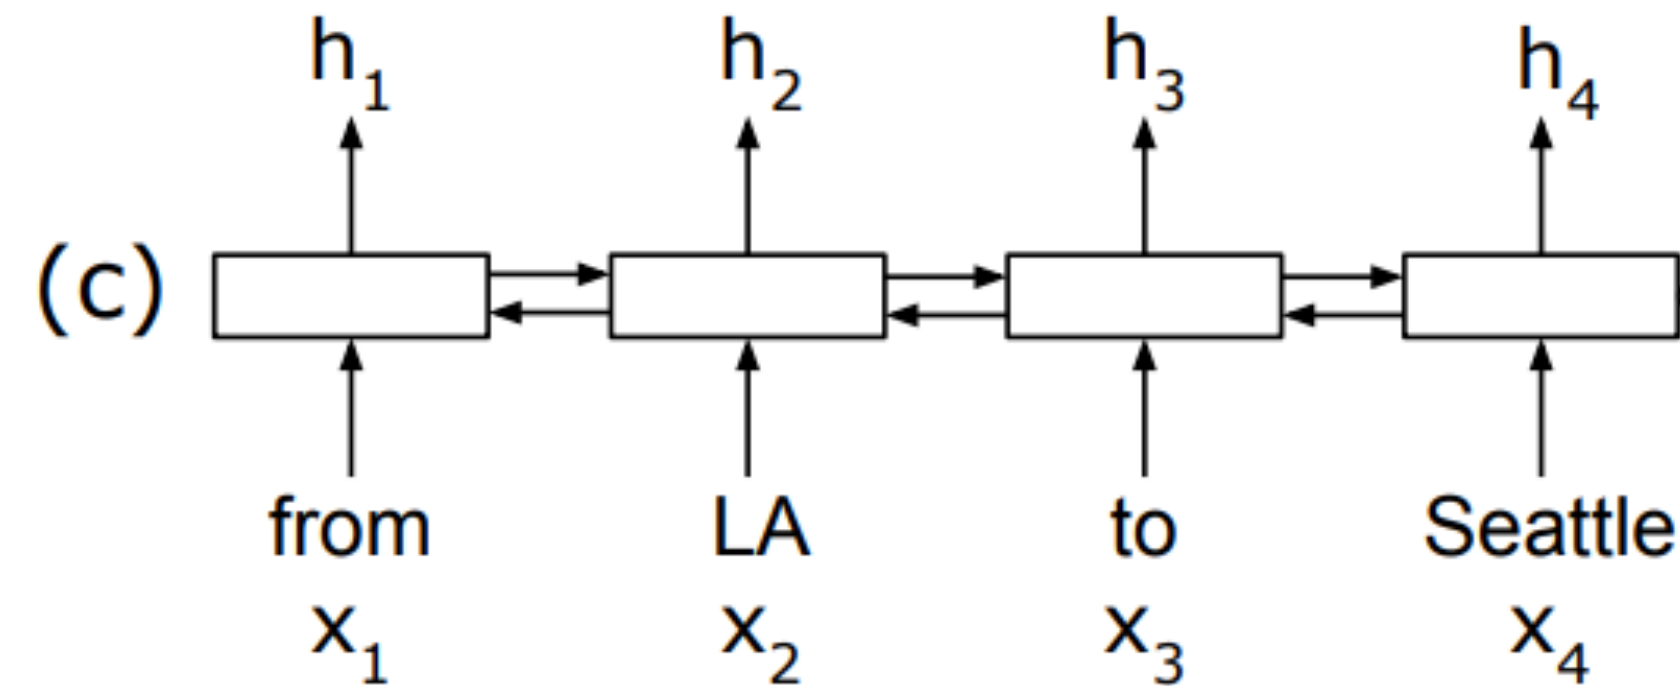
Intent: day_name

what meals are served on american flight 811 from tampa to milwaukee

Intent: meal

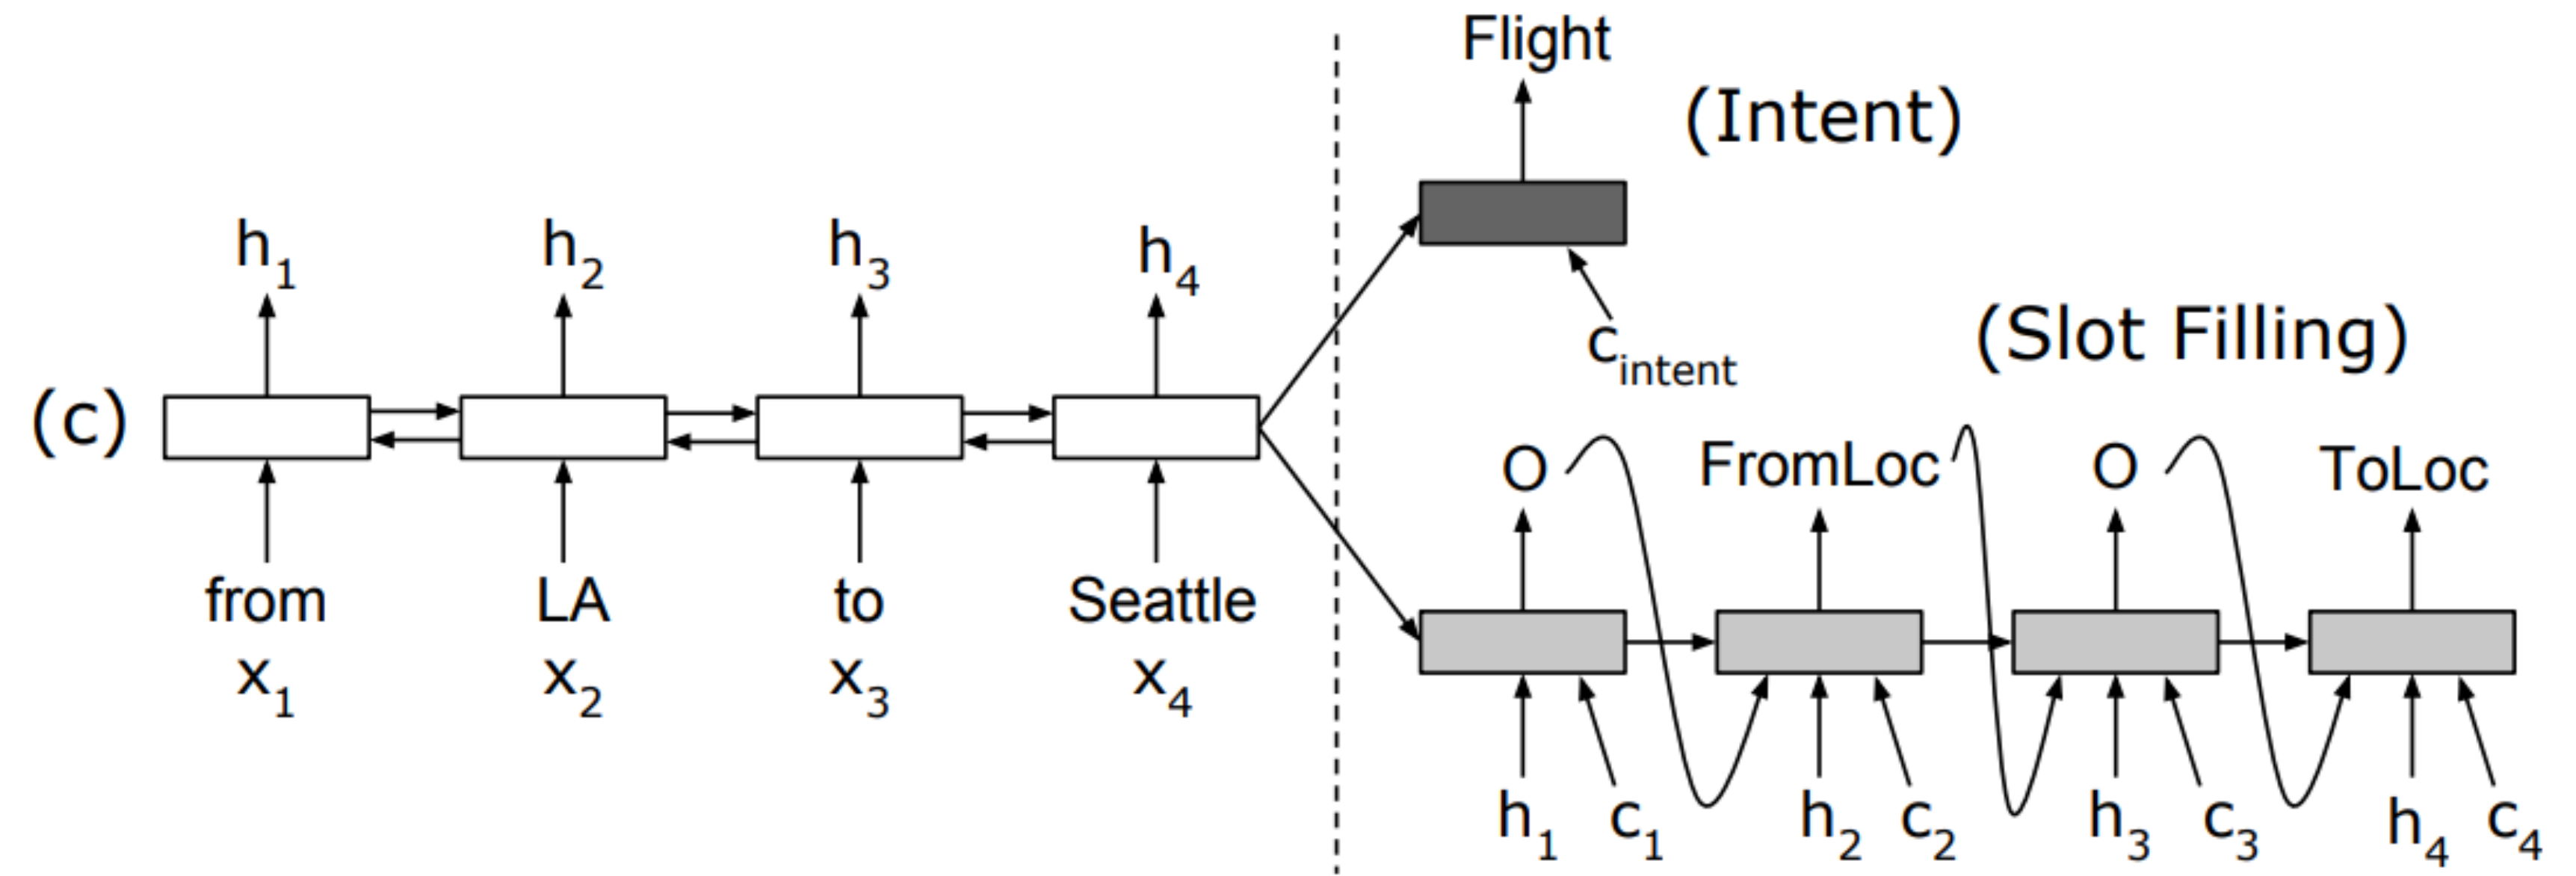
Joint Intent Classification and Tagging

- ▶ RNN jointly predicts intent and slot tags



Joint Intent Classification and Tagging

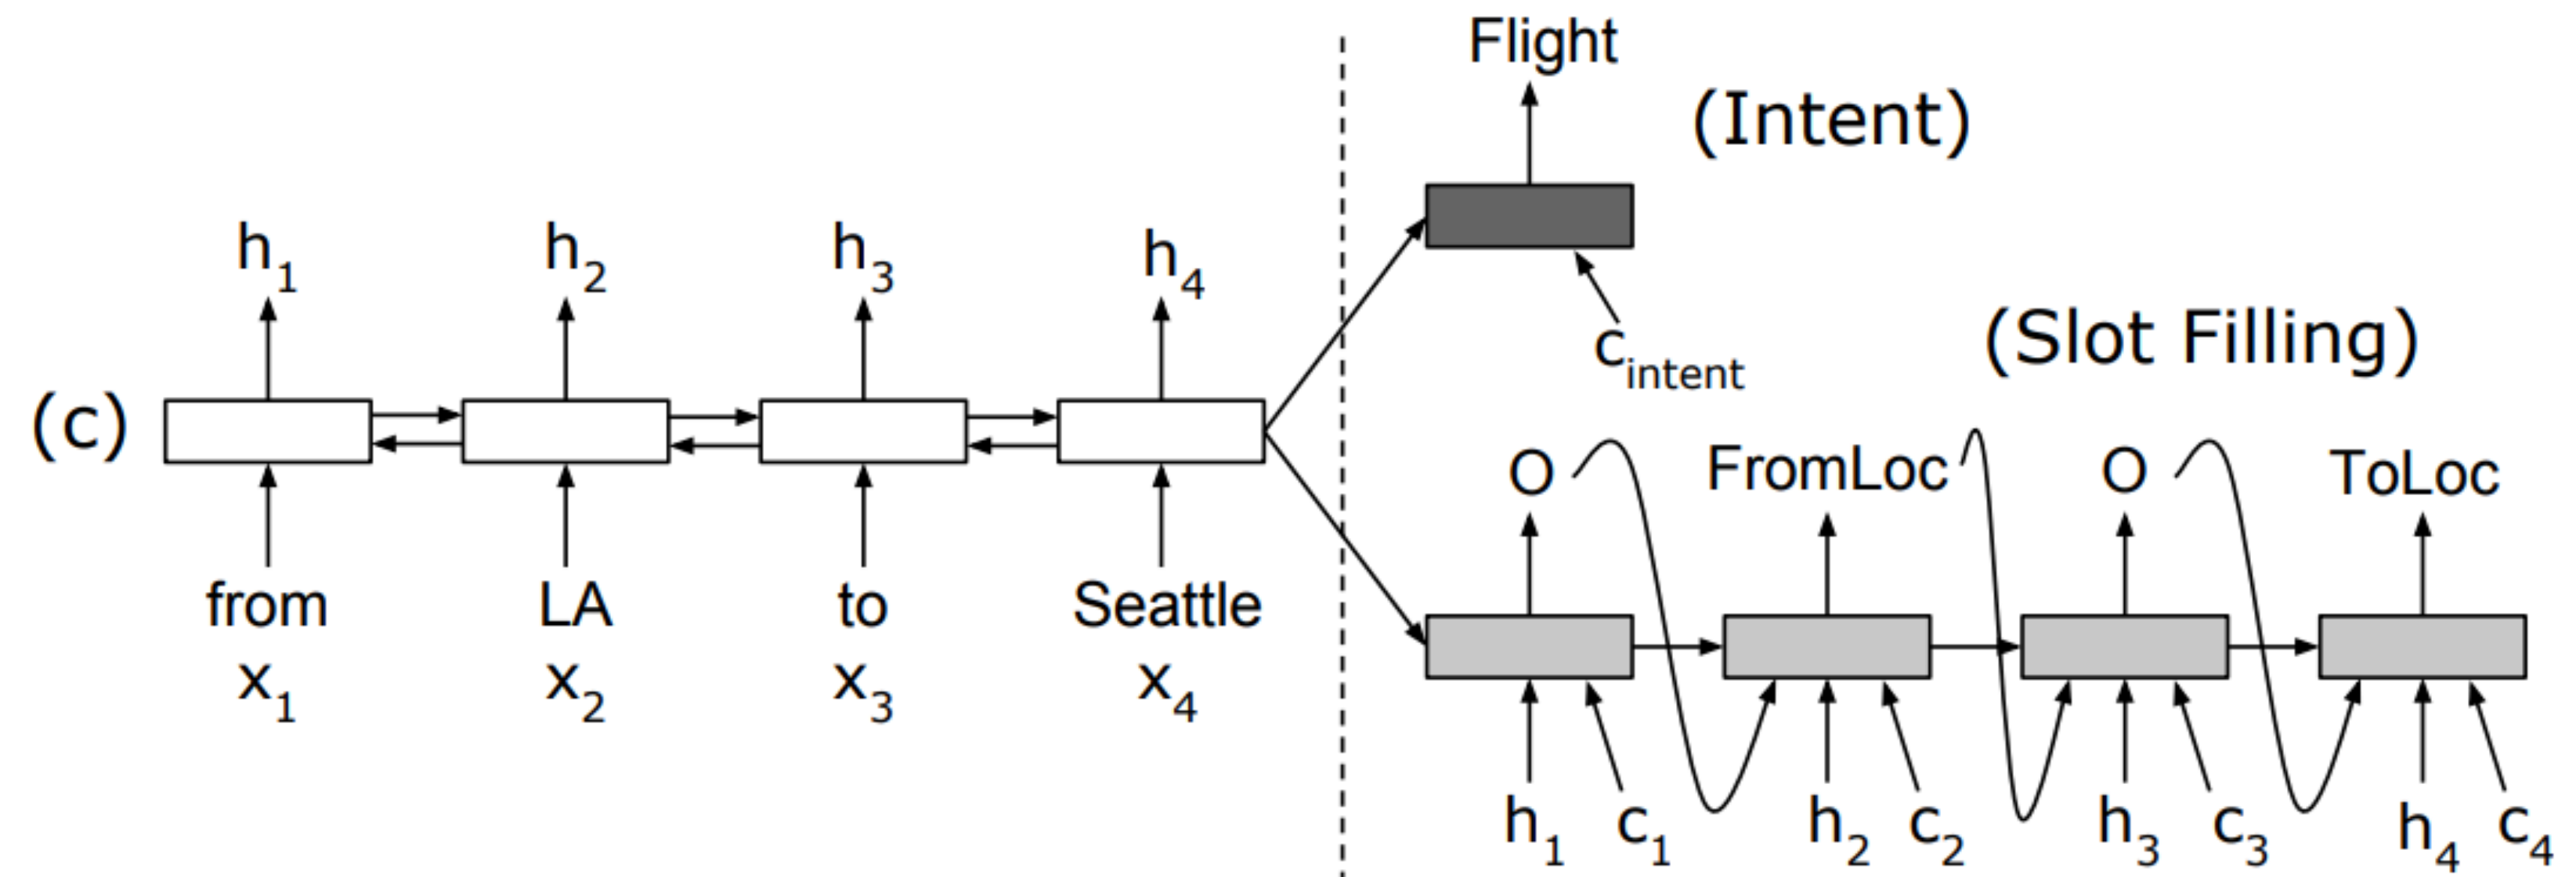
- ▶ RNN jointly predicts intent and slot tags



what is c here?

Joint Intent Classification and Tagging

- ▶ RNN jointly predicts intent and slot tags



Model	F1 Score	Intent Error (%)
RecNN [8]	93.22	4.60
RecNN+Viterbi [8]	93.96	4.60
Attention Encoder-Decoder NN (with aligned inputs)	95.87	1.57
Attention BiRNN	95.98	1.79

Air Travel Information Service (ATIS) as conversational semantic parsing

\bar{x}_1 : *show me flights from seattle to boston next monday*

\bar{y}_1 : (SELECT DISTINCT flight.flight_id FROM flight WHERE (flight.from_airport IN (SELECT airport_service.airport_code FROM airport_service WHERE airport_service.city_code IN (SELECT city.city_code FROM city WHERE city.city_name = 'SEATTLE')))) AND (flight.to_airport IN (SELECT airport_service.airport_code FROM airport_service WHERE airport_service.city_code IN (SELECT city.city_code FROM city WHERE city.city_name = 'BOSTON')))) AND (flight.flight_days IN (SELECT days.days_code FROM days WHERE days.day_name IN (SELECT date_day.day_name FROM date_day WHERE date_day.year = 1993 AND date_day.month_number = 2 AND date_day.day_number = 8))));

\bar{x}_2 : *on american airlines*

\bar{y}_2 : (SELECT DISTINCT flight.flight_id FROM flight WHERE (flight.airline_code = 'AA') AND (flight.from_airport IN (SELECT airport_service.airport_code FROM airport_service WHERE airport_service.city_code IN (SELECT city.city_code FROM city WHERE city.city_name = 'SEATTLE')))) AND (flight.to_airport IN (SELECT airport_service.airport_code FROM airport_service WHERE airport_service.city_code IN (SELECT city.city_code FROM city WHERE city.city_name = 'BOSTON')))) AND (flight.flight_days IN (SELECT days.days_code FROM days WHERE days.day_name IN (SELECT date_day.day_name FROM date_day WHERE date_day.year = 1993 AND date_day.month_number = 2 AND date_day.day_number = 8))));

Air Travel Information Service (ATIS) as conversational semantic parsing

\bar{x}_1 : *show me flights from seattle to boston next monday*

\bar{y}_1 : (SELECT DISTINCT flight.flight_id FROM flight WHERE (flight.from_airport IN (SELECT airport_service.airport_code FROM airport_service WHERE airport_service.city_code IN (SELECT city.city_code FROM city WHERE city.city_name = 'SEATTLE')))) AND (flight.to_airport IN (SELECT airport_service.airport_code FROM airport_service WHERE airport_service.city_code IN (SELECT city.city_code FROM city WHERE city.city_name = 'BOSTON')))) AND (flight.flight_days IN (SELECT days.days_code FROM days WHERE days.day_name IN (SELECT date_day.day_name FROM date_day WHERE date_day.year = 1993 AND date_day.month_number = 2 AND date_day.day_number = 8))));

\bar{x}_2 : *on american airlines*

\bar{y}_2 : (SELECT DISTINCT flight.flight_id FROM flight WHERE (flight.airline_code = 'AA') AND (flight.from_airport IN (SELECT airport_service.airport_code FROM airport_service WHERE airport_service.city_code IN (SELECT city.city_code FROM city WHERE city.city_name = 'SEATTLE')))) AND (flight.to_airport IN (SELECT airport_service.airport_code FROM airport_service WHERE airport_service.city_code IN (SELECT city.city_code FROM city WHERE city.city_name = 'BOSTON')))) AND (flight.flight_days IN (SELECT days.days_code FROM days WHERE days.day_name IN (SELECT date_day.day_name FROM date_day WHERE date_day.year = 1993 AND date_day.month_number = 2 AND date_day.day_number = 8))));

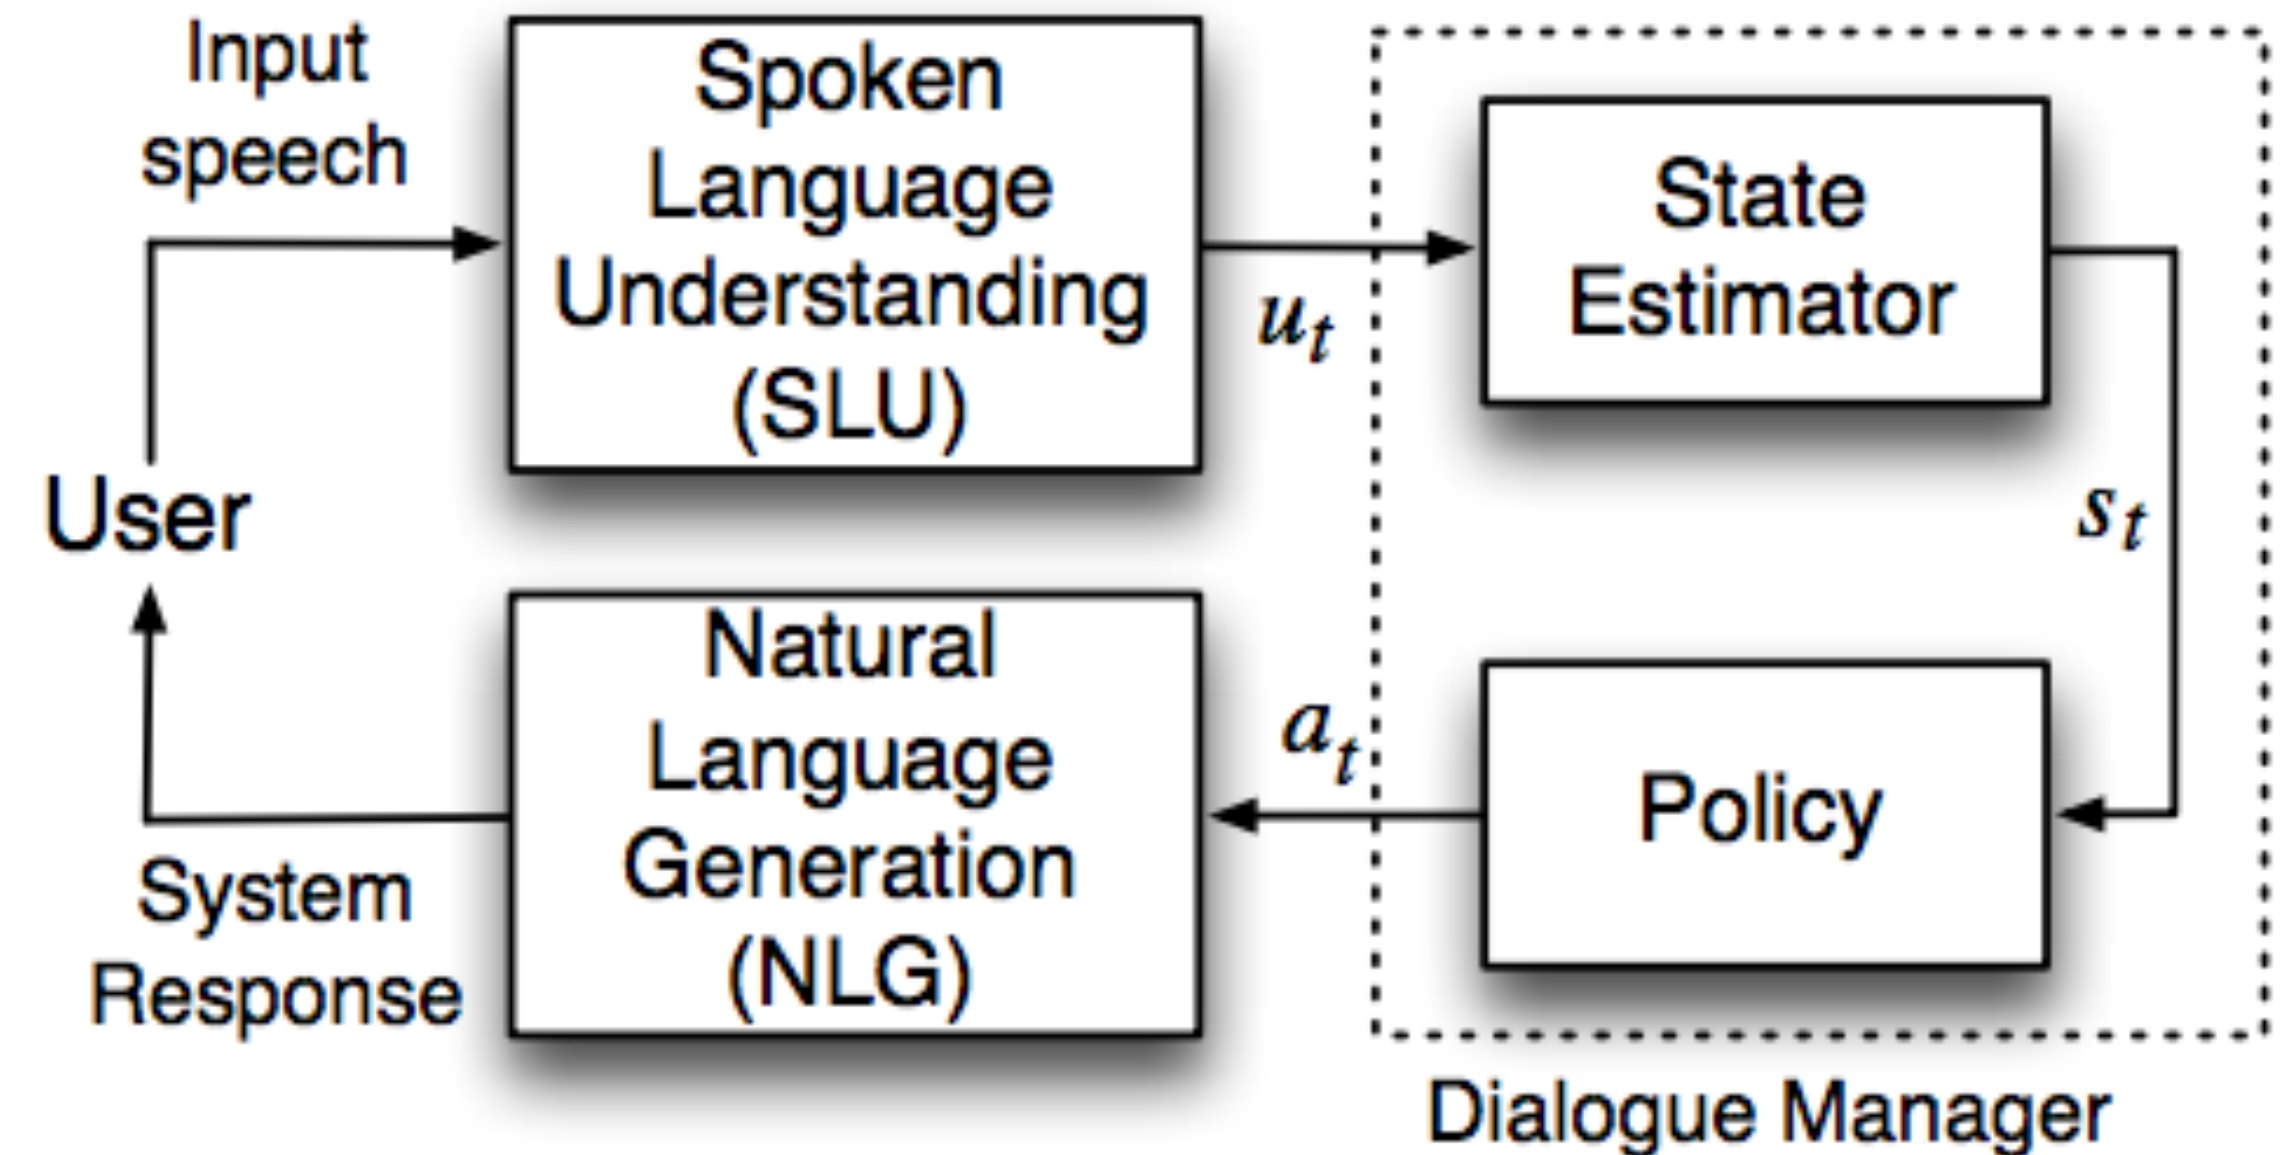
- ▶ Need to **use dialogue context** to do the right thing. Here we're appending American Airlines as a constraint to the previous query
- ▶ seq2seq model mapping to query with copy mechanism

Full Dialogue Task

- ▶ Parsing / language understanding is just one piece of a system

Full Dialogue Task

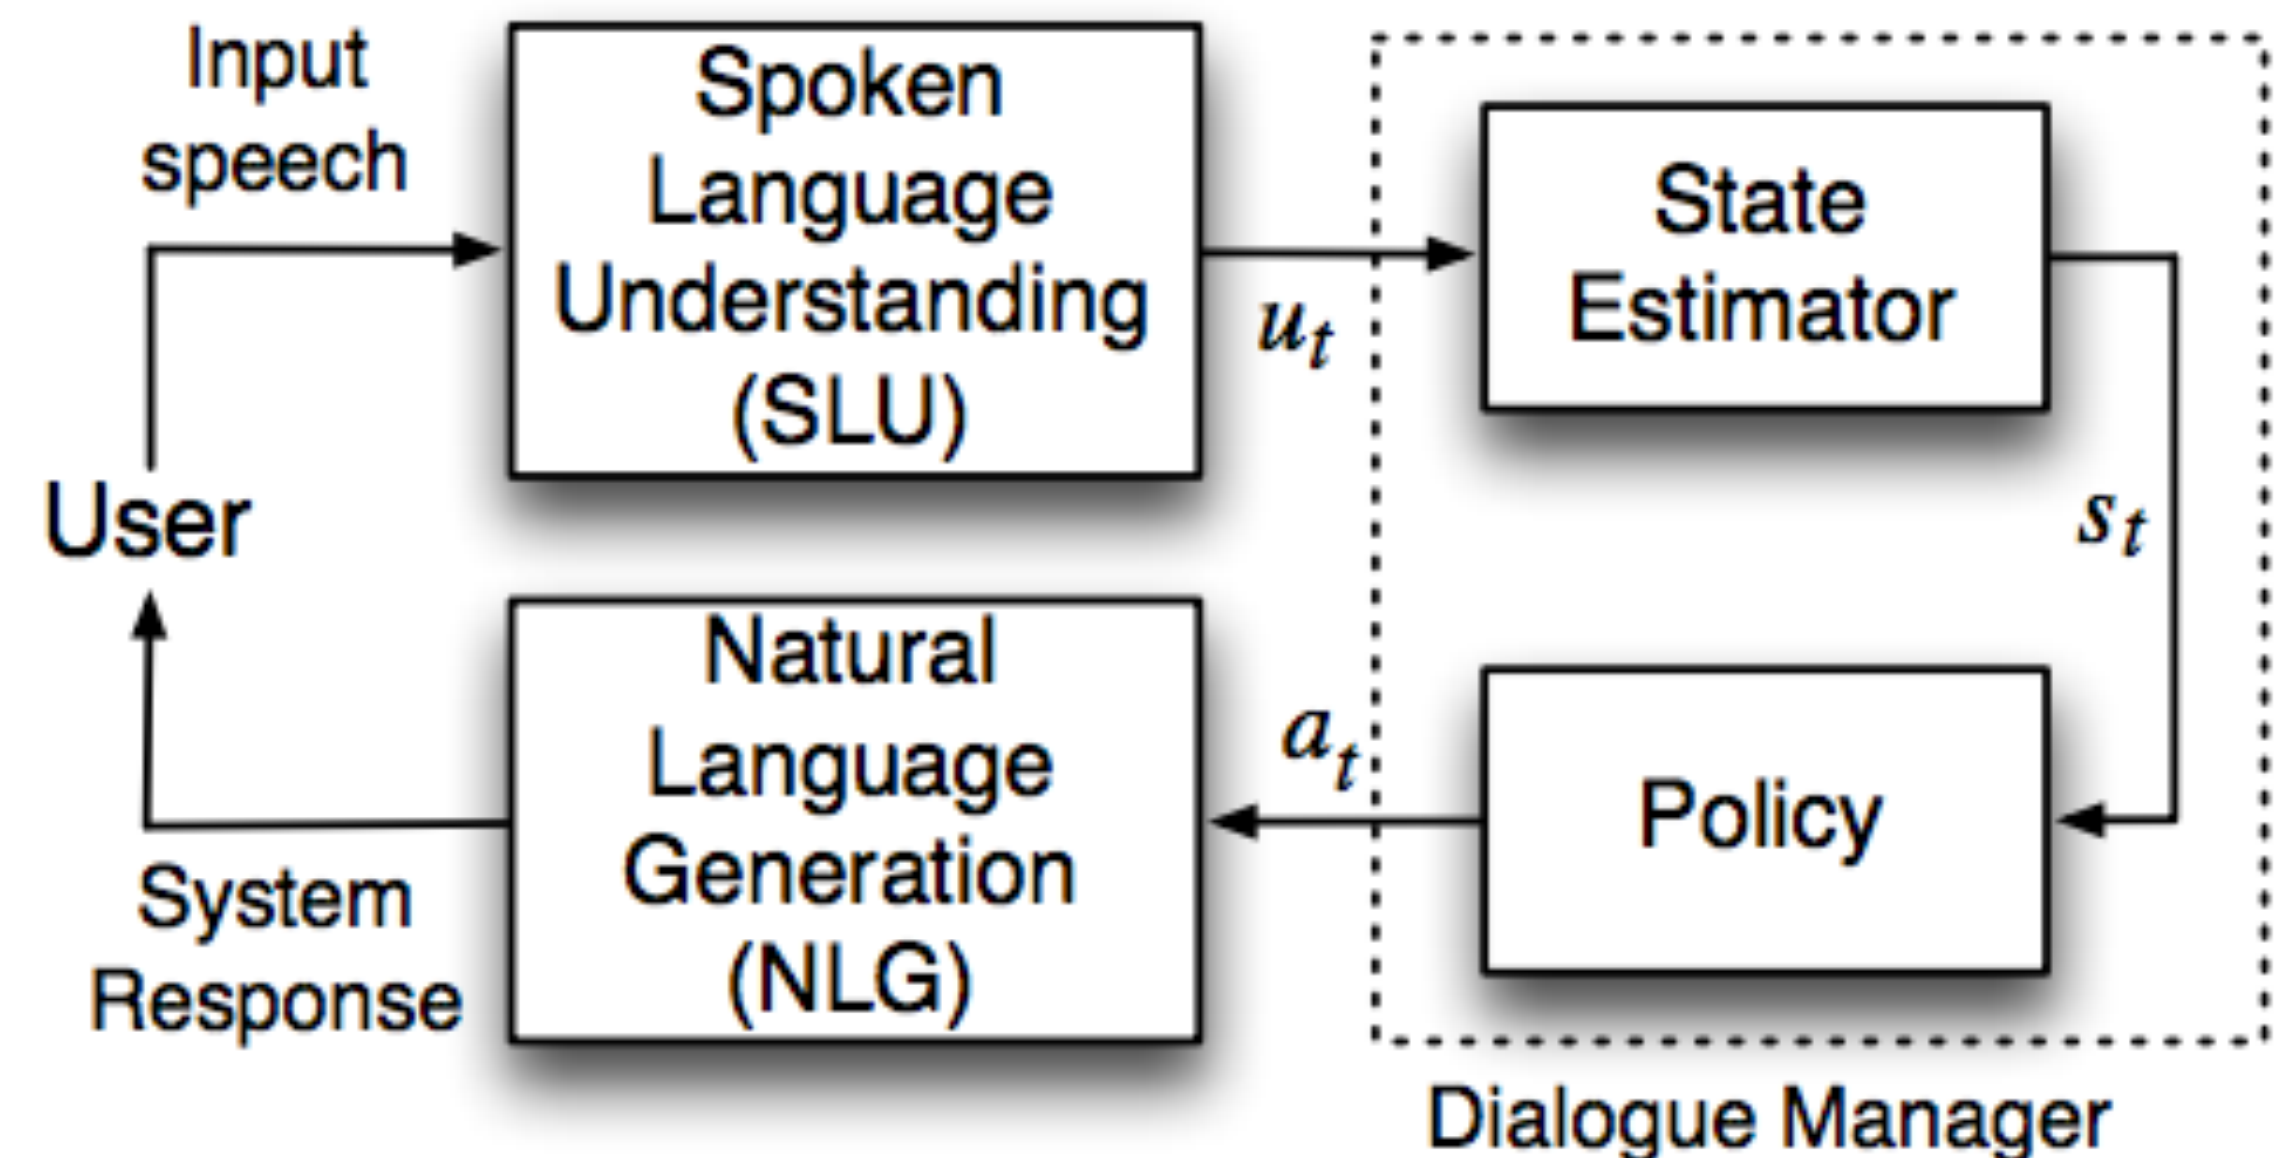
- ▶ Parsing / language understanding is just one piece of a system



- ▶ User utterance -> update dialogue state -> take an action decided by policy (e.g., query the restaurant database) -> say something

Full Dialogue Task

- ▶ Parsing / language understanding is just one piece of a system
- ▶ Dialogue **state**: reflects any information about the conversation (e.g., search history)



- ▶ User utterance -> **update dialogue state** -> **take an action decided by policy** (e.g., query the restaurant database) -> **say something**
- ▶ Much more complex than chatbots!

- ▶ **Full dialogue:** User utterance -> update dialogue state -> take an action decided by policy (e.g., query the restaurant database) -> say something
-

“Find me a good sushi restaurant in Chelsea”

```
restaurant_type <- sushi
```

```
location <- Chelsea
```

```
curr_result <- execute_search()
```

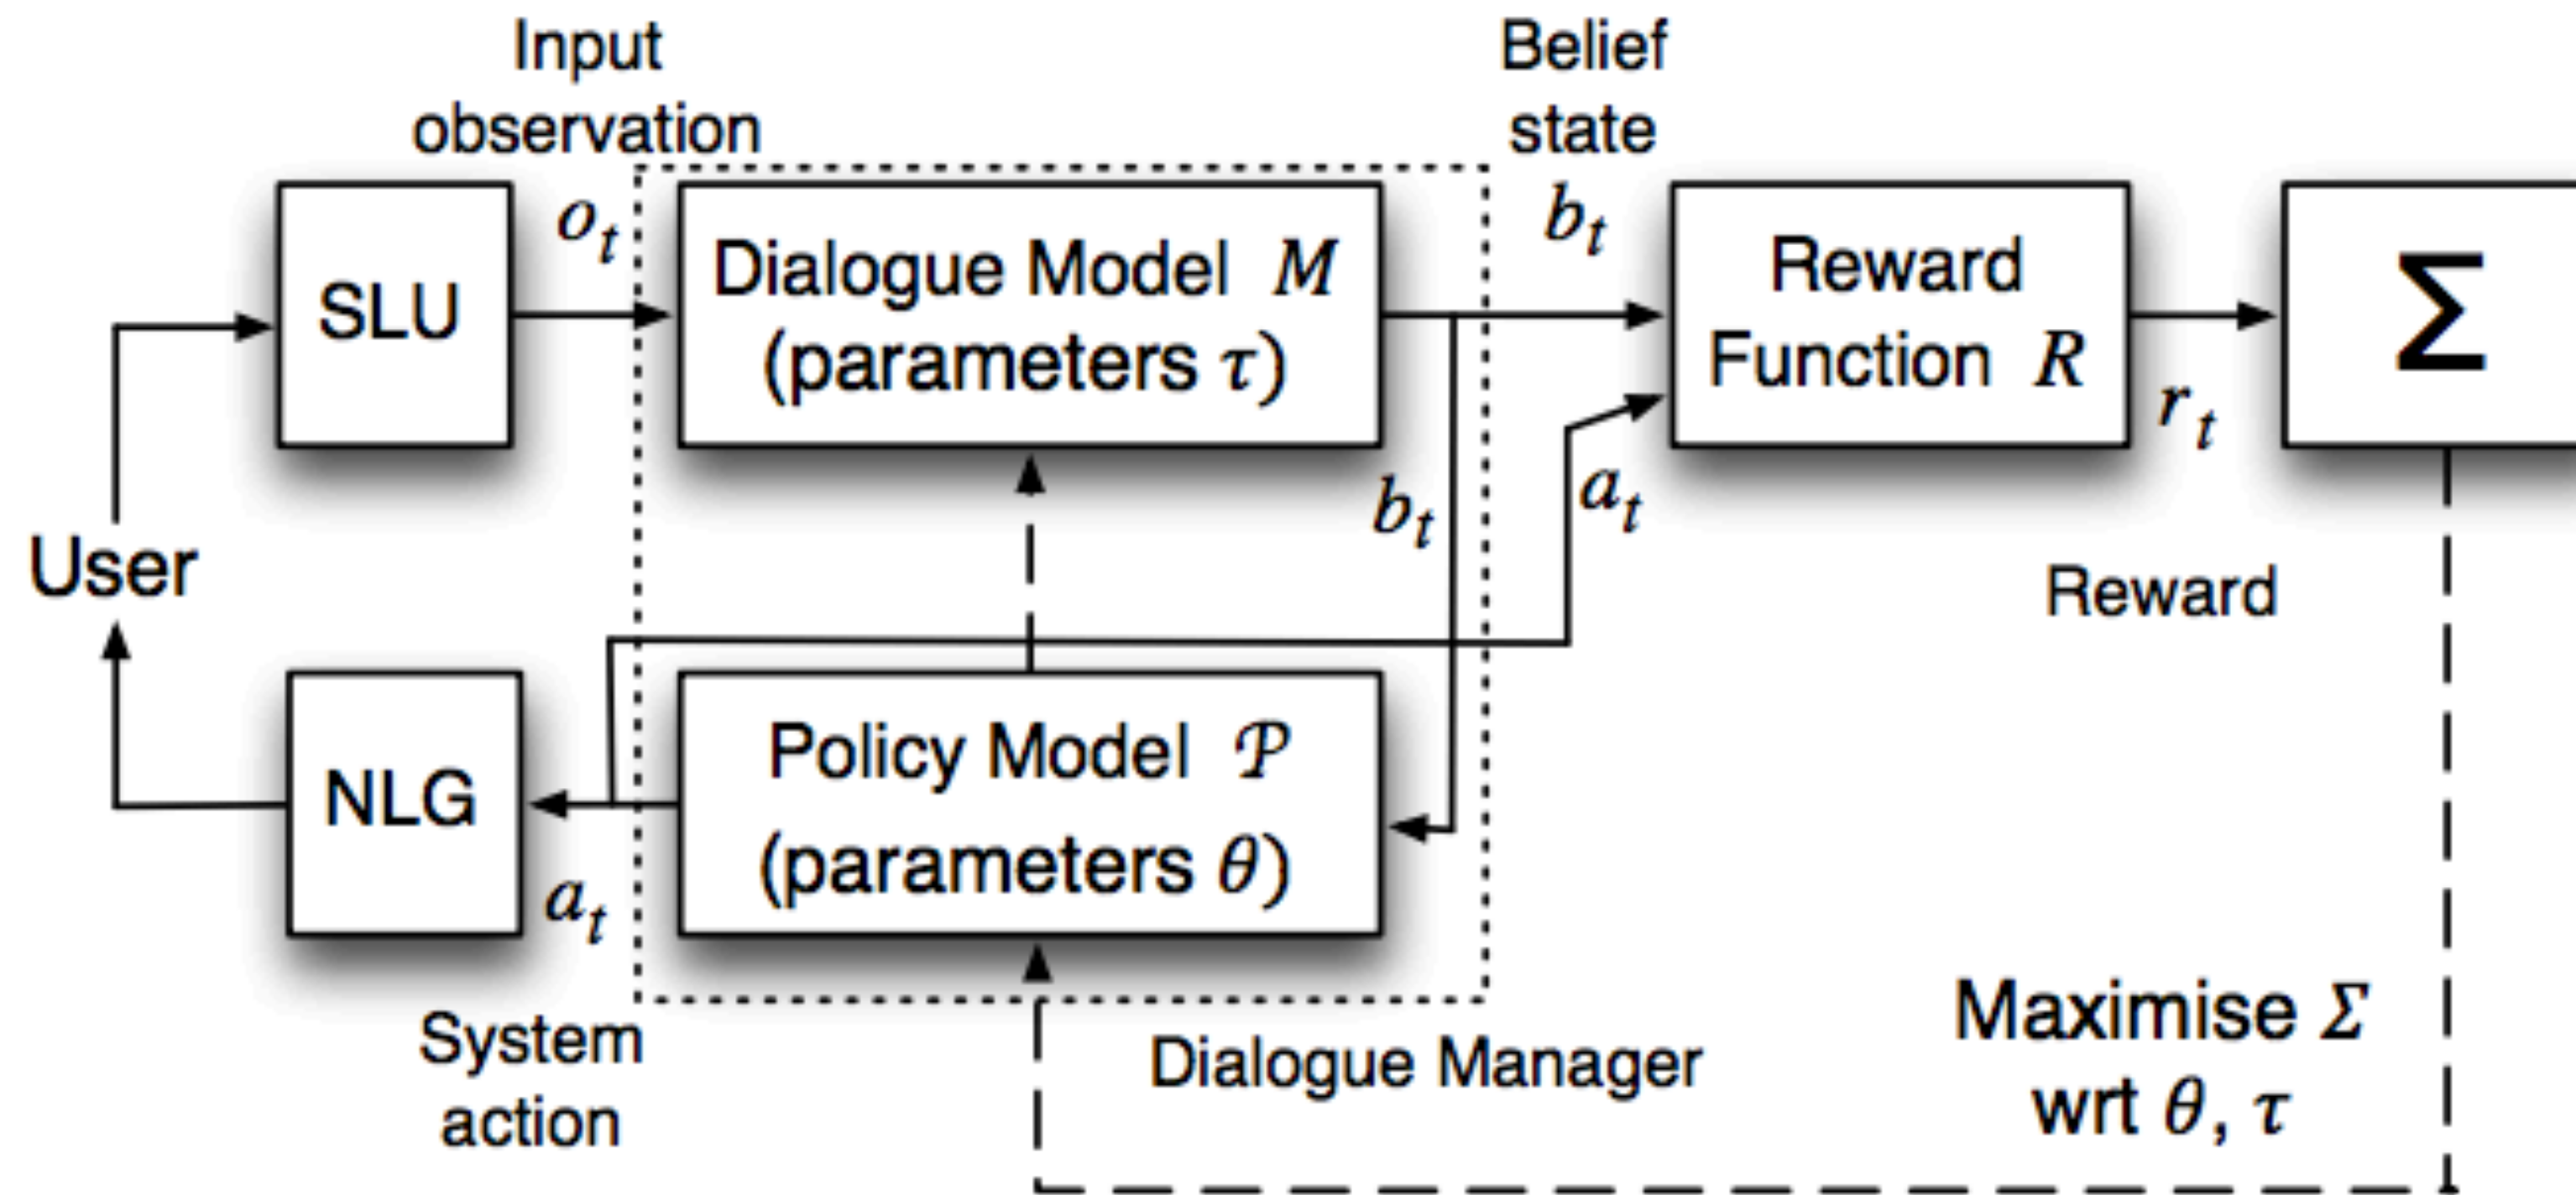
“Sushi Seki Chelsea is a sushi restaurant in Chelsea with 4.4 stars on Google”

“How expensive is it?”

```
get_value(cost, curr_result)
```

“Entrees are around \$30 each”

POMDP-based Dialogue Systems

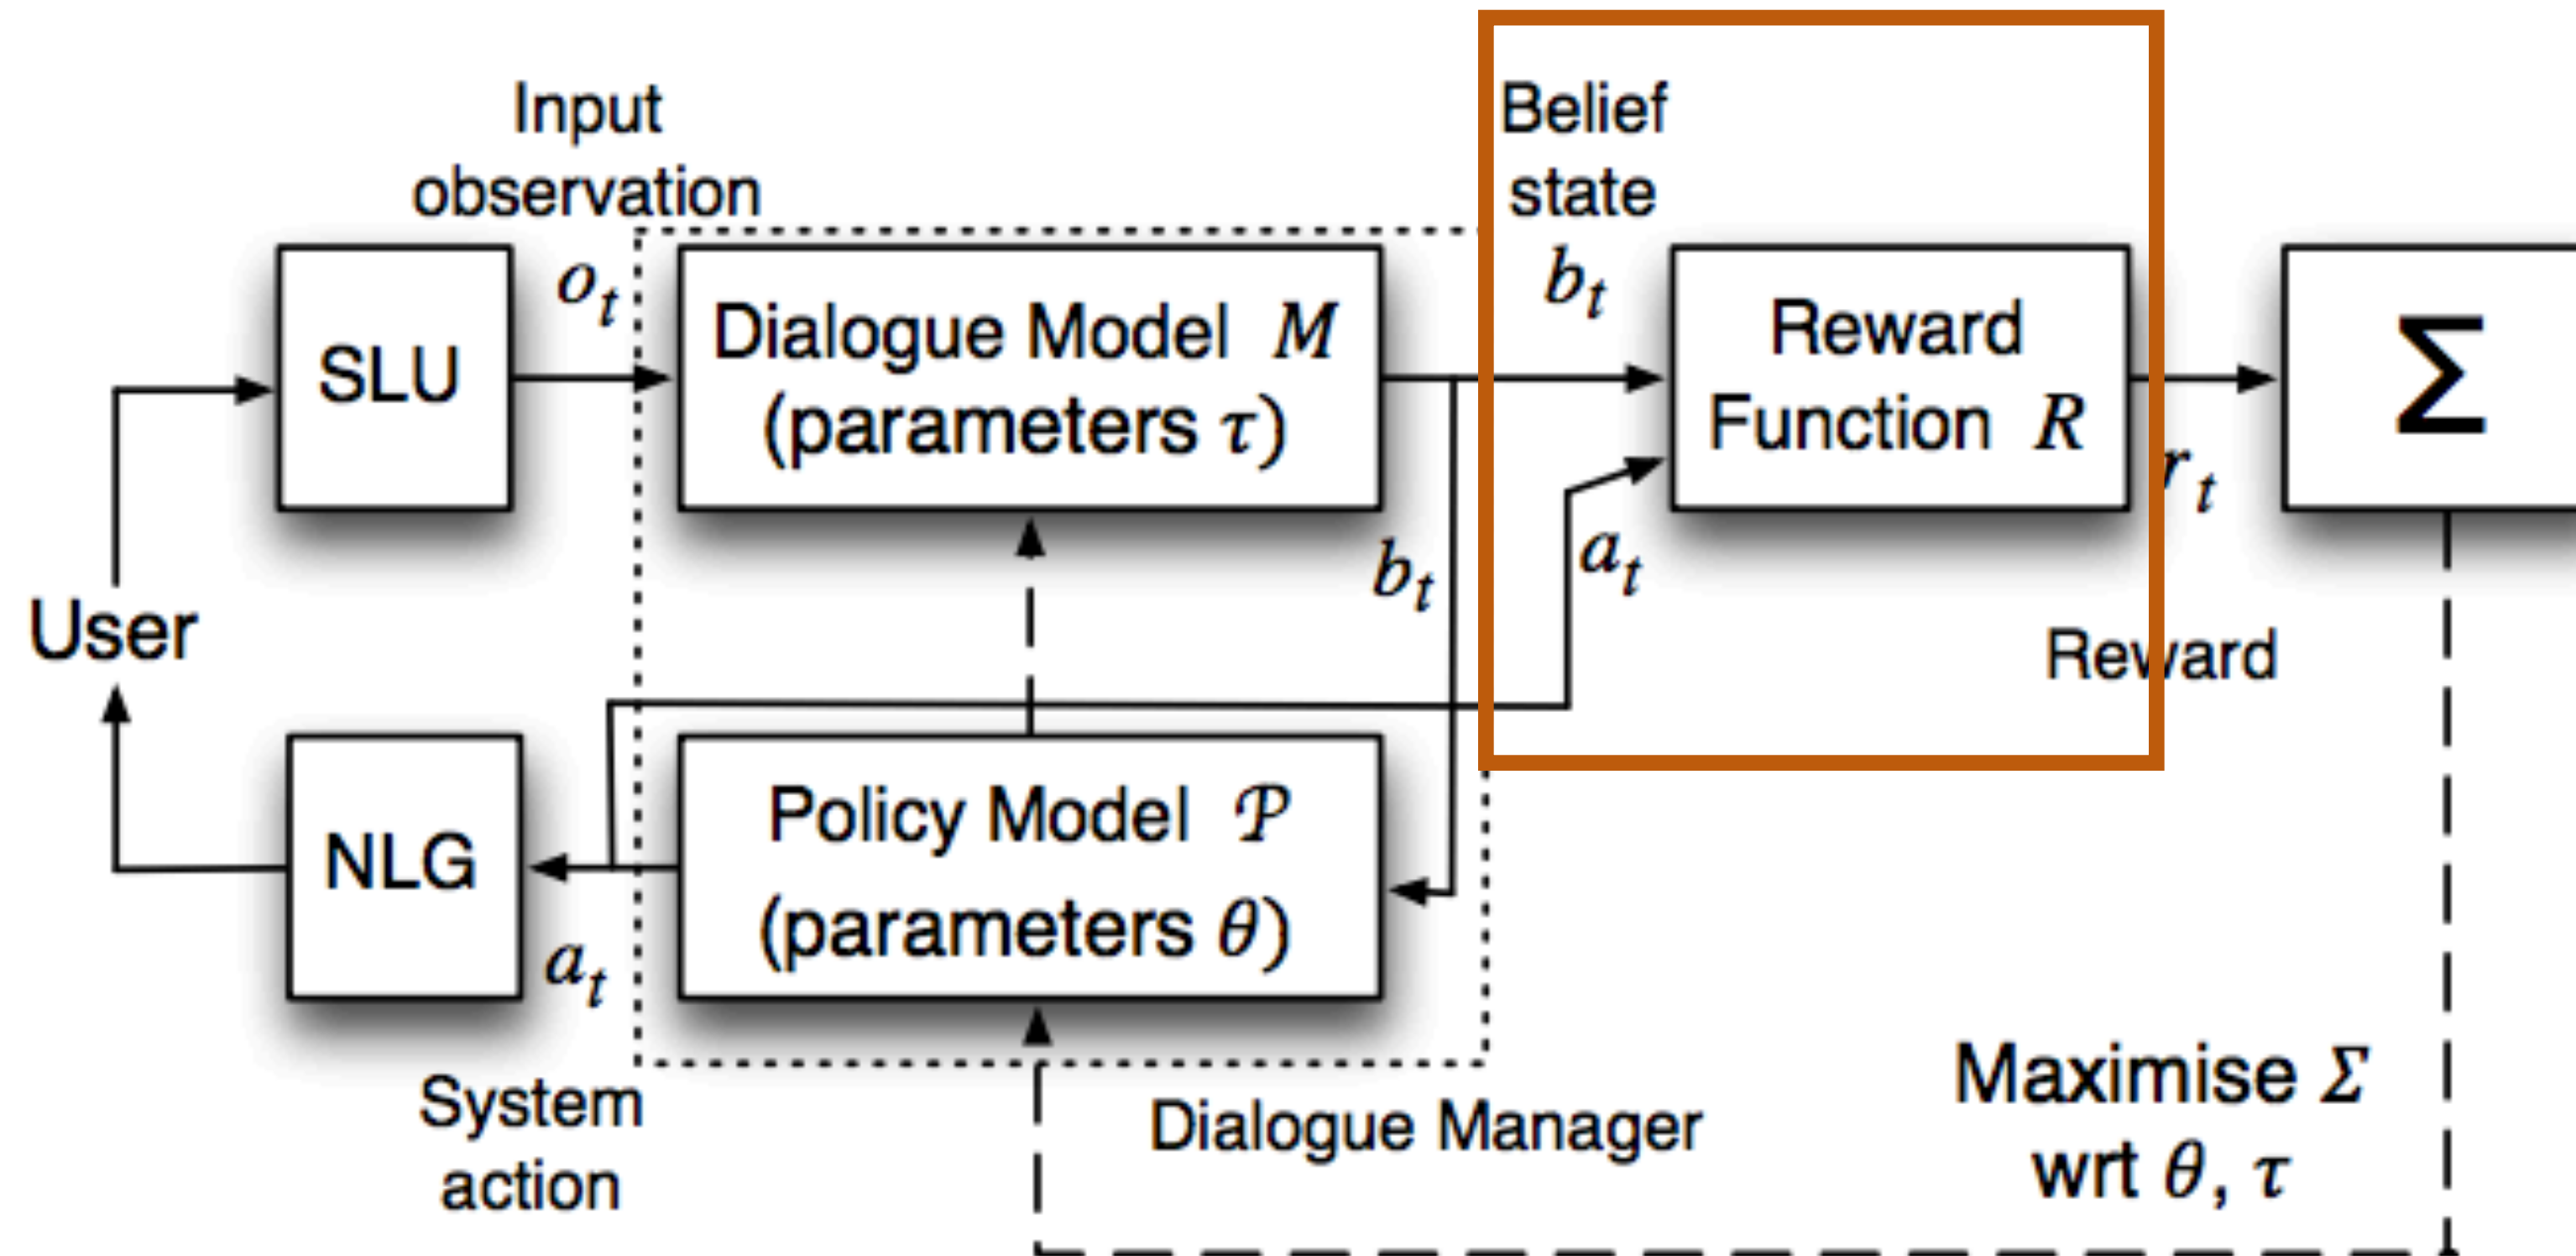


- ▶ POMDP: user is the “environment,” an utterance is a noisy signal of state
- ▶ Dialogue model: can look like a parser or any kind of encoder model
- ▶ Generator: use templates or seq2seq model

Reinforcement learning type of approach

Young et al. (2013)

POMDP-based Dialogue Systems



- ▶ POMDP: user is the “environment,” an utterance is a noisy signal of state
- ▶ Dialogue model: can look like a parser or any kind of encoder model
- ▶ Generator: use templates or seq2seq model
- ▶ **Where do rewards come from?**

Reward for completing task?

Find me a good sushi restaurant in Chelsea

```
restaurant_type <- sushi
```

```
location <- Chelsea
```

```
curr_result <- execute_search()
```

Sushi Seki Chelsea is a sushi restaurant in Chelsea with
4.4 stars on Google

How expensive is it?

...

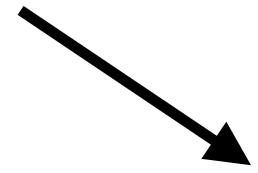
Okay make me a reservation!

```
+1 make_reservation(curr_result)
```

Reward for completing task?

Find me a good sushi restaurant in Chelsea

Very indirect signal
of what should
happen up here



```
restaurant_type <- sushi
```

```
location <- Chelsea
```

```
curr_result <- execute_search()
```

Sushi Seki Chelsea is a sushi restaurant in Chelsea with
4.4 stars on Google

How expensive is it?

...

Okay make me a reservation!

```
+1 make_reservation(curr_result)
```


User gives reward?

Find me a good sushi restaurant in Chelsea

How does the user
know the right
search happened?

↘ **+1**

```
restaurant_type <- sushi
```

```
location <- Chelsea
```

```
curr_result <- execute_search()
```

Sushi Seki Chelsea is a sushi restaurant in Chelsea with
4.4 stars on Google

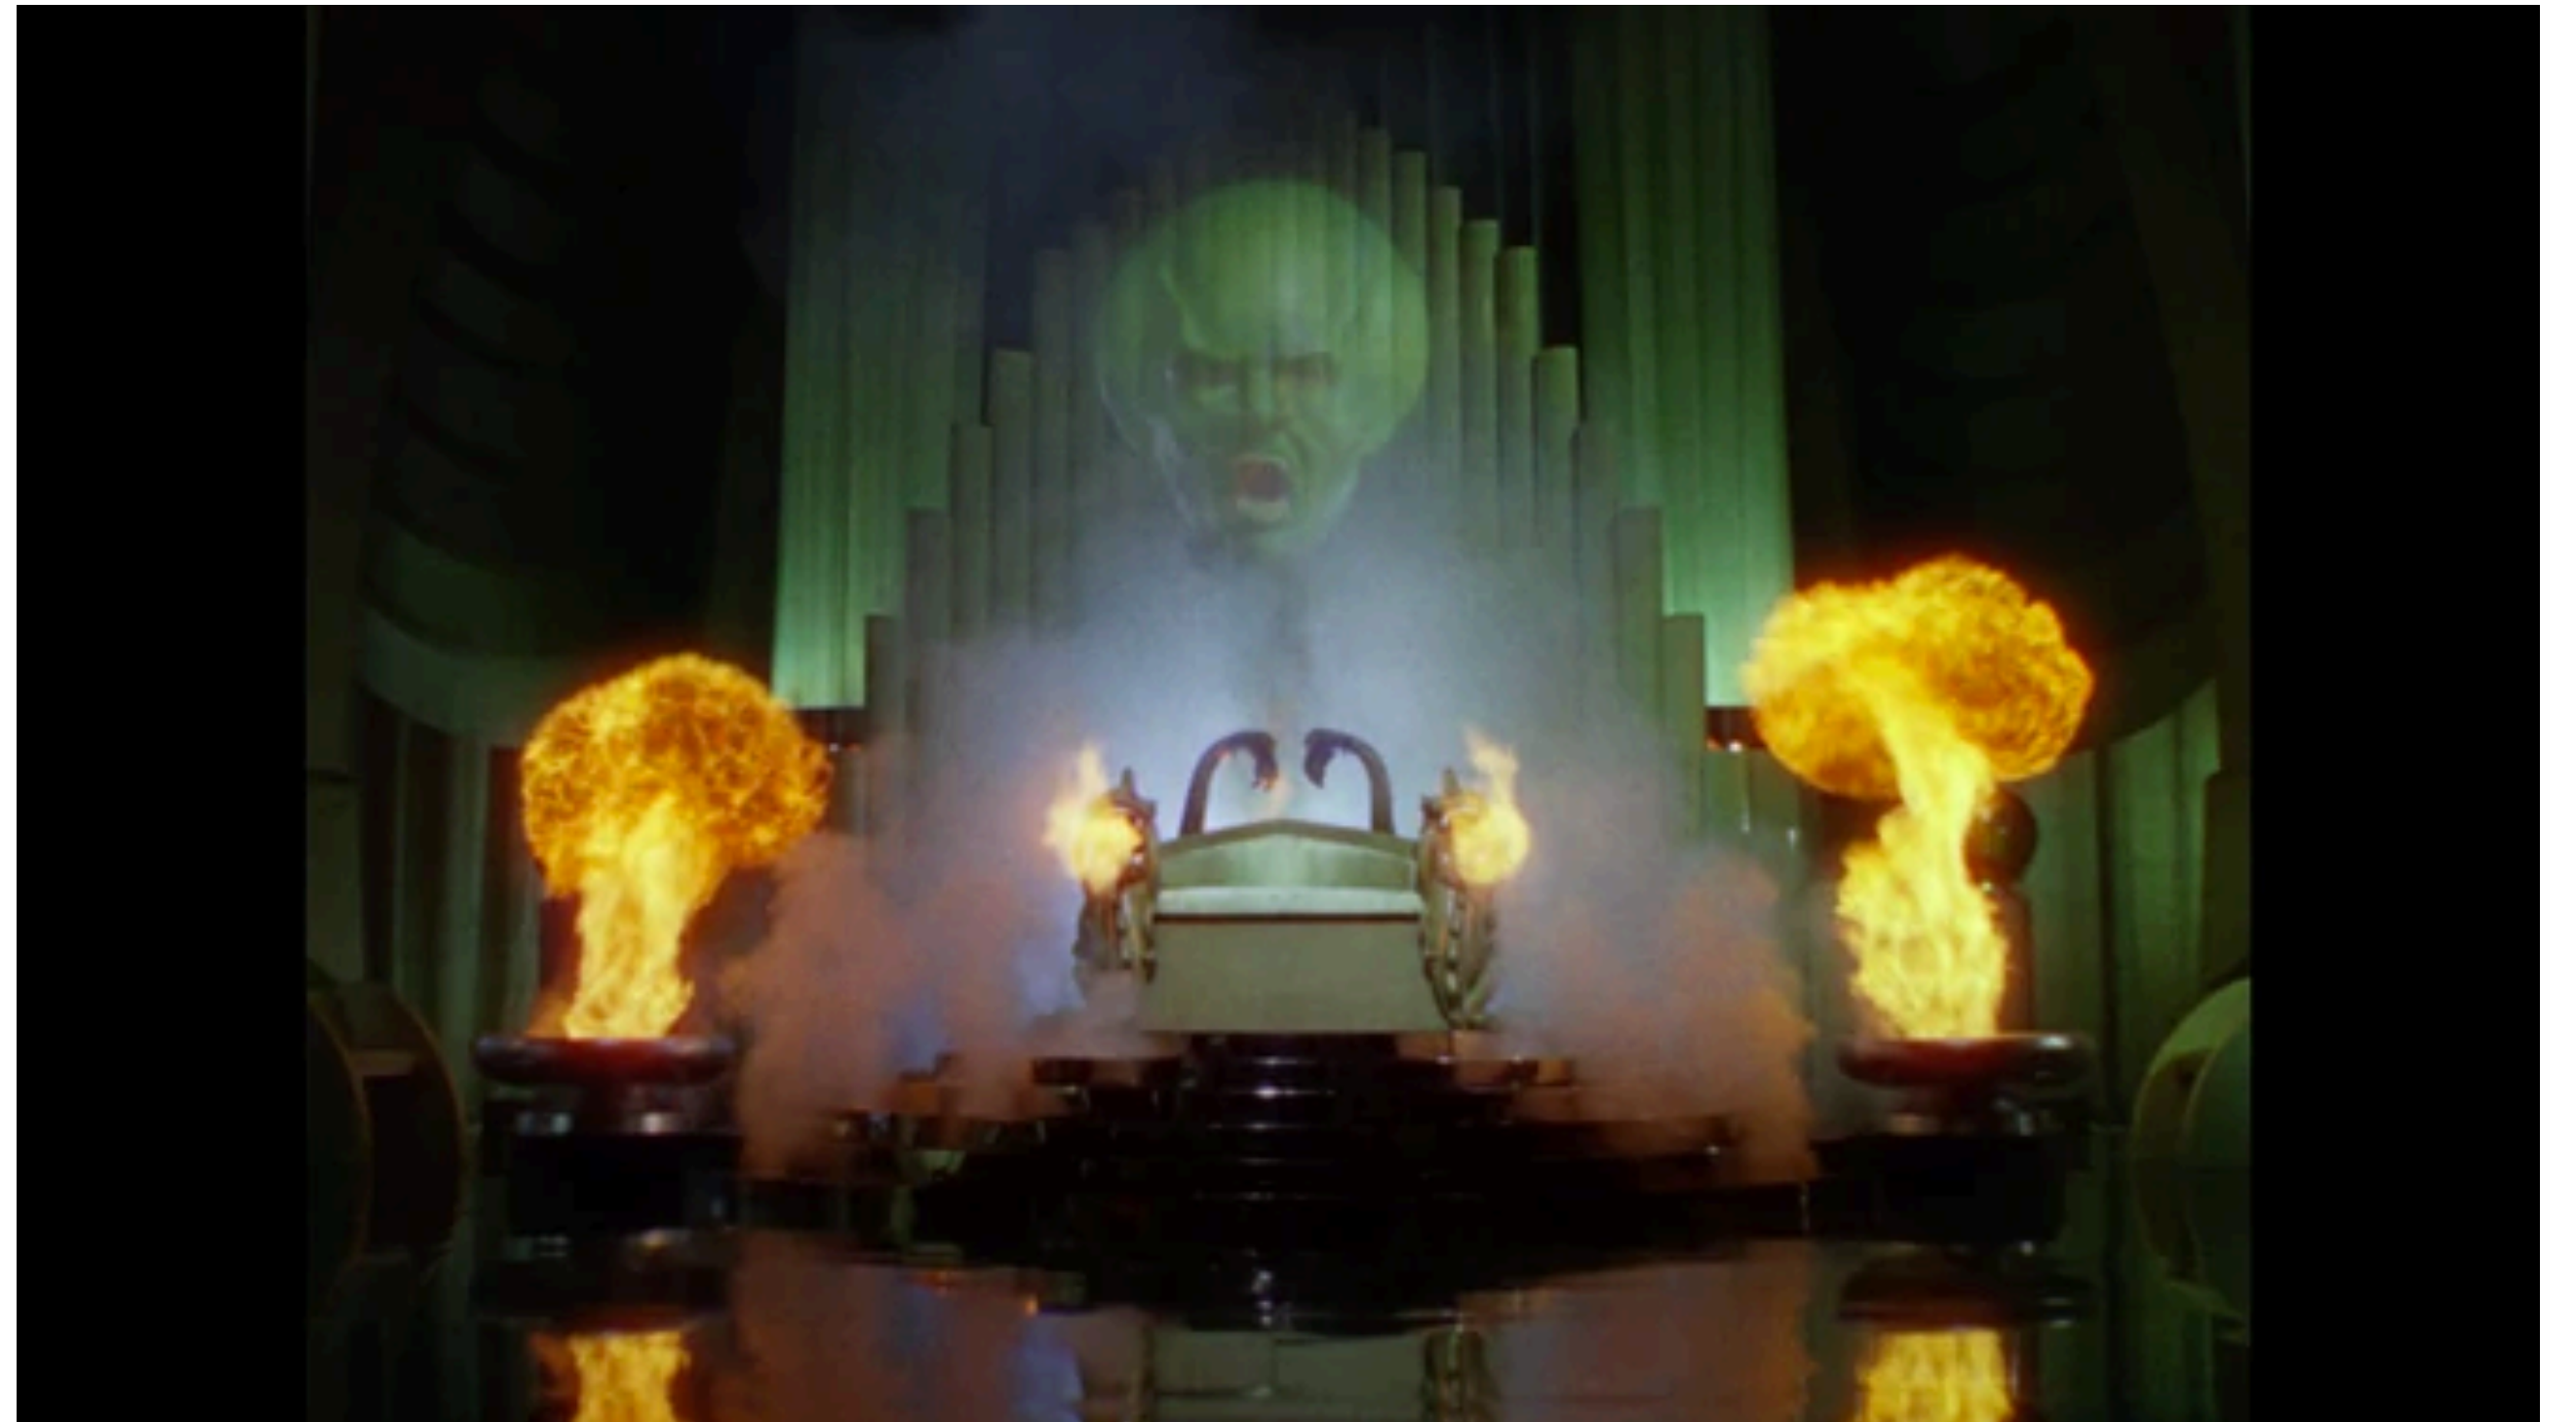
How expensive is it?

```
get_value(cost, curr_result)
```

+1 Entrees are around \$30 each

Wizard-of-Oz

- ▶ Learning from demonstrations: Assume “wizard” can demonstrate what the dialogue system should do (e.g., how to update its state and take actions)



Kelley (early 1980s), Ford and Smith (1982)

Full Dialogue Task

Find me a good sushi restaurant in Chelsea

wizard enters these

```
restaurant_type <- sushi  
location <- Chelsea  
curr_result <- execute_search()
```

wizard types this out or invokes templates

```
Sushi Seki Chelsea is a sushi restaurant in Chelsea with  
4.4 stars on Google
```

- ▶ Wizard can be a trained expert and know exactly what the dialogue systems is supposed to do

Learning from Static Traces

Hi!

Hello, what can I help you with today?

I'd like to book a table for six people in an expensive price range with British food.

I'm on it!

Where should it be?

In London.

Ok, let me look into some options for you.

```
api_call(British, London, Six, Expensive)
```

Actually i would prefer for four.

Sure. Is there anything else to update?

No.

Ok let me look into some options for you.

```
api call(British, London, Four, Expensive)
```

Task 1

Issuing API calls

Task 2

Updating API calls

- ▶ Using either wizard-of-Oz or other annotations, can collect static traces and train from these

Other issues: Language Ambiguity & Model Adaptivity

Find me a good sushi restaurant in Chelsea

```
restaurant_type <- sushi  
location <- Chelsea  
stars <- 4+  
curr_result <- execute_search()
```

- ▶ User asked for a “good” restaurant — does that mean we should change our model to filter by star rating? What does “good” mean?
- ▶ Hard to change system behavior if training from static traces, especially if system capabilities or desired behavior change

Goal-oriented Dialogue

- ▶ Tons of industry interest!
- ▶ Dozens of startups + medium-sized companies in this space
- ▶ Big Companies: Apple Siri, Google Assistant, Amazon Alexa, Microsoft Cortana, Facebook, Samsung Bixby, Tencent WeChat, ASAPP
- ▶ Lots of cool work that's not public yet

Takeaways

- ▶ Some decent chatbots, but unclear how to make these more sophisticated than they are right now
- ▶ Task-oriented dialogue systems are growing in scope and complexity — really exciting systems on the way
- ▶ More and more problems are being formulated as dialogue — interesting applications but challenging to get working well