Dialog Systems

Lara J. Martin (she/they)

https://laramartin.net/neurosymbolic-text-gen

Slides modified from Chris Callison-Burch & João Sedoc

Learning Objectives

Interpret difficulties in human communication

Distinguish between chatbots and task-oriented dialog systems

Link reinforcement learning to chatbots

Determine how task-oriented systems can be made with neural/ML-based methods

Conversational Agents aka Dialog Systems

Digital Assistants

Answering questions on websites

Communicating with robots

Chatting for fun

Clinical uses

ChatGPT

Challenging properties of human conversation

- Turn taking
- Speech acts
- Grounding
- Dialogue structure
- Initiative
- Implicature

Turn taking

A conversation is a sequence of turns, where you take a turn and then I take a turn. A turn can be a sentence, or a single word.

A system must know when to start and stop talking.

Spoken dialogue systems must also detect whether a user is done speaking, so they can process the utterance and respond. This task of **endpoint detection** is tricky because people often pause mid-turn.

Speech acts

Constatives: committing the speaker to something's being the case (*answering*, *claiming*, *confirming*, *denying*, *disagreeing*, *stating*)
Directives: attempts by the speaker to get the addressee to do something (*advising*, *asking*, *forbidding*, *inviting*, *ordering*, *requesting*)
Commissives: committing the speaker to some future course of action (*promising, planning, vowing, betting, opposing*)
Acknowledgments: express the speaker's attitude regarding the hearer with respect to some social action (*apologizing*, *greeting*, *thanking*, *accepting an acknowledgment*)

Conversational Implicature

Agent: And, what day in May did you want to travel?

Client: I need to be there for a meeting that's from the 12th to the 15th.

Notice that the client does not in fact answer the agent's question. The speaker expects the hearer to draw certain inferences.

Implicature means certain kinds of inferences.

Gricean maxims. The *maxim of relevance* means that when the client mentions a meeting on the 12th, the agent reasons 'There must be some relevance for mentioning this meeting. What could it be?'

Two Classes of Dialog Systems

1. **Chatbots**
   - Systems designed for extended conversations
   - Chatting for fun and entertainment

2. **Task-Oriented Dialogue Agents**
   - Goal-Based Agents
   - Siri/Alexa, interface with robots, booking flights or hotels

Part I: Chatbots

Systems designed for extended conversations. Chatbots mimic unstructured conversations or ‘chats’ that are characteristic of informal human-human interaction

Architectures include:

Rule-Based

- Pattern-action rules

Corpus-Based

- Information Retrieval
- Neural networks

Part I: Chatbots

Systems designed for extended conversations. Chatbots mimic unstructured conversations or ‘chats’ that are characteristic of informal human-human interaction

Architectures include:

Rule-Based

- **Pattern-action rules**

Corpus-Based

- Information Retrieval
- Neural networks

ELIZA: Weizenbaum (1966)



Simulates a psychologist

Rogerian psychology is a rare type of conversation where "assume the pose of knowing almost nothing of the real world"

Draw the patient out by reflecting patient's statements back at them

ELIZA: Weizenbaum (1966)

Men are all alike.

IN WHAT WAY

They're always bugging us about something or other.

CAN YOU THINK OF A SPECIFIC EXAMPLE

Well my boyfriend made me come here.

YOUR BOYFRIEND MADE YOU COME HERE

He says I'm depressed much of the time.

I AM SORRY TO HEAR YOU ARE DEPRESSED

::

::

WHO ELSE IN YOUR FAMILY TAKES CARE OF YOU?

My father

YOUR FATHER

You are like my father in some ways.

ELIZA: Weizenbaum (1966)

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 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'M 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 Pattern Transform Rules

Pattern matching rules:

.* YOU **(.*)** ME **–>** WHAT MAKES YOU THINK I **$1** YOU?

This rule would transform

You **must really hate** me

into

WHAT MAKES YOU THINK I **MUST REALLY HATE** YOU?

Simplified sketch of ELIZA

function ELIZA GENERATOR(user *sentence*) **returns** *response*

Find the word *w* in *sentence* that has the highest keyword rank
  if *w* exists
    Choose the highest ranked rule *r* for *w* that matches *sentence*
    response Apply the transform in *r* to *sentence*
    if *w* = 'my'
      future Apply a transformation from the 'memory' rule list to *sentence*
      Push *future* onto memory stack
  else (no keyword applies) **either**
    response Apply the transform for the NONE keyword to *sentence*
  or
    response Pop the top response from the memory stack
return(*response*)

PARRY: Colby (1971)

ELIZA like regular expressions But a much richer –

- Control Structure
- Language Understanding Capabilities
- Mental Model : Parry has affective variables
  - Fear, Mistrust, Anger
  - If Anger level is high, respond with hostility

First system to pass the **Turing test (1971)**

Turing Test

“

I propose to consider the question, ‘Can machines think?’

”

- A human evaluator would judge text-based conversations between a human and a machine designed to generate human-like responses.
- If the evaluator cannot reliably tell the machine from the human, the machine is said to have passed the test.

A. M. TURING, I.—COMPUTING MACHINERY AND INTELLIGENCE, *Mind*, Volume LIX, Issue 236, October 1950, Pages 433–460, https://doi.org/10.1093/mind/LIX.236.433

Part I: Chatbots

Systems designed for extended conversations. Chatbots mimic unstructured conversations or ‘chats’ that are characteristic of informal human-human interaction

Architectures include:

Rule-Based

- Pattern-action rules (Eliza)

Corpus-Based

- **Information Retrieval**
- **Neural networks**

Conversational Data

Need: large collections of human conversations

Conversational threads on Twitter or Weibo (微博)

Retrieve dialog from movies, indexing subtitles

Recorded telephone conversations, collected for speech research

Crowdsourced conversations

What are some issues that might come from using each of these datasets?

Information Retrieval based Chatbots

Treat the human user's input as a query vector **q**

Search over a large corpus **C** of conversation to find the closest matching turn **t'** in those previous conversations.

Return the response **r** to that conversational turn.

$t' = \arg\max_{t \in C}$ cosine_similarity(q, t)

r = response(t')

q = Have you watched Doctor Who?

t' = Do you like Doctor Who?

r = Yes, I love SciFi shows!

IR-based Models

- Can use more features than just words in query q
  - User features - Information about the user or sentiment
  - Prior turns – Use conversation so far
  - Narrative (non-dialogue) text
    - COBOT chatbot (Isbell et al., 2000) :
      - Generate responses by selecting sentences from the Unabomber Manifesto by Theodore Kaczynski, articles on alien abduction, the scripts of “The Big Lebowski” and “Planet of the Apes”
      - Wikipedia Text

Seq2Seq Encoder-Decoder Models

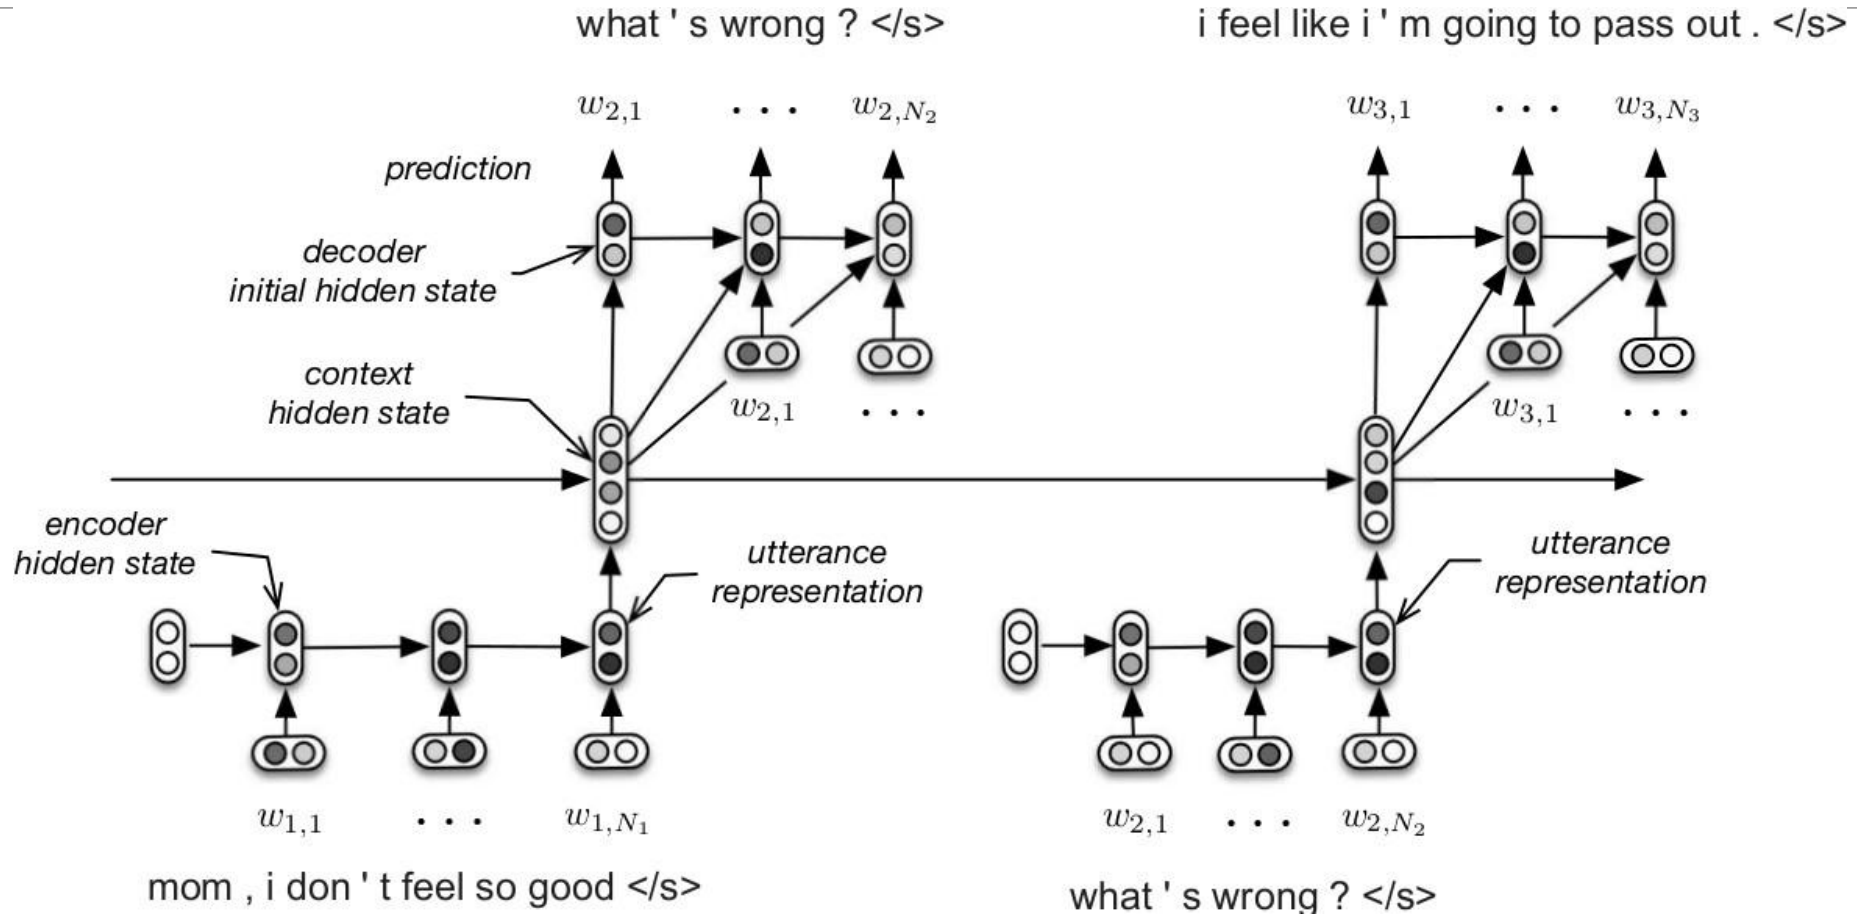

Sample Responses: Enc-Dec RNN Model

Input	tell me ... how long have you had this falling sickness ?
Vanilla-SEQ2SEQ	i'm not a doctor.
Adversarial	a few months, i guess .
Input	so i had the doctors test sammy 's response to conditioning .
Vanilla-SEQ2SEQ	sammy wrote the test sammy wrote the test .
Adversarial	so he took the pills .
Input	they didn 't have much success with this problem commander .
Vanilla-SEQ2SEQ	they 're not the only ones who have been in the system .
Adversarial	can we find someone else ?

Responses that are harder to distinguish from human responses

Remember RLHF?

Review: Human Preferences

A set of sampled completions for a prompt

Triples

Prompt

Sample A

Sample B

Sample C

$$D = \{x^i, y_w^i, y_l^i\}$$

Prompt

Preferred Response

Dispreferred Response

Review: REINFORCE

Sample a sequence from your model, score the sequence, and use the score to train the model.

$$L_{RL} = -\sum_{t=1}^{T} r(\hat{y}_t)\log P(\hat{y}_t|\{y^*\};\{\hat{y}\}_{<t})$$

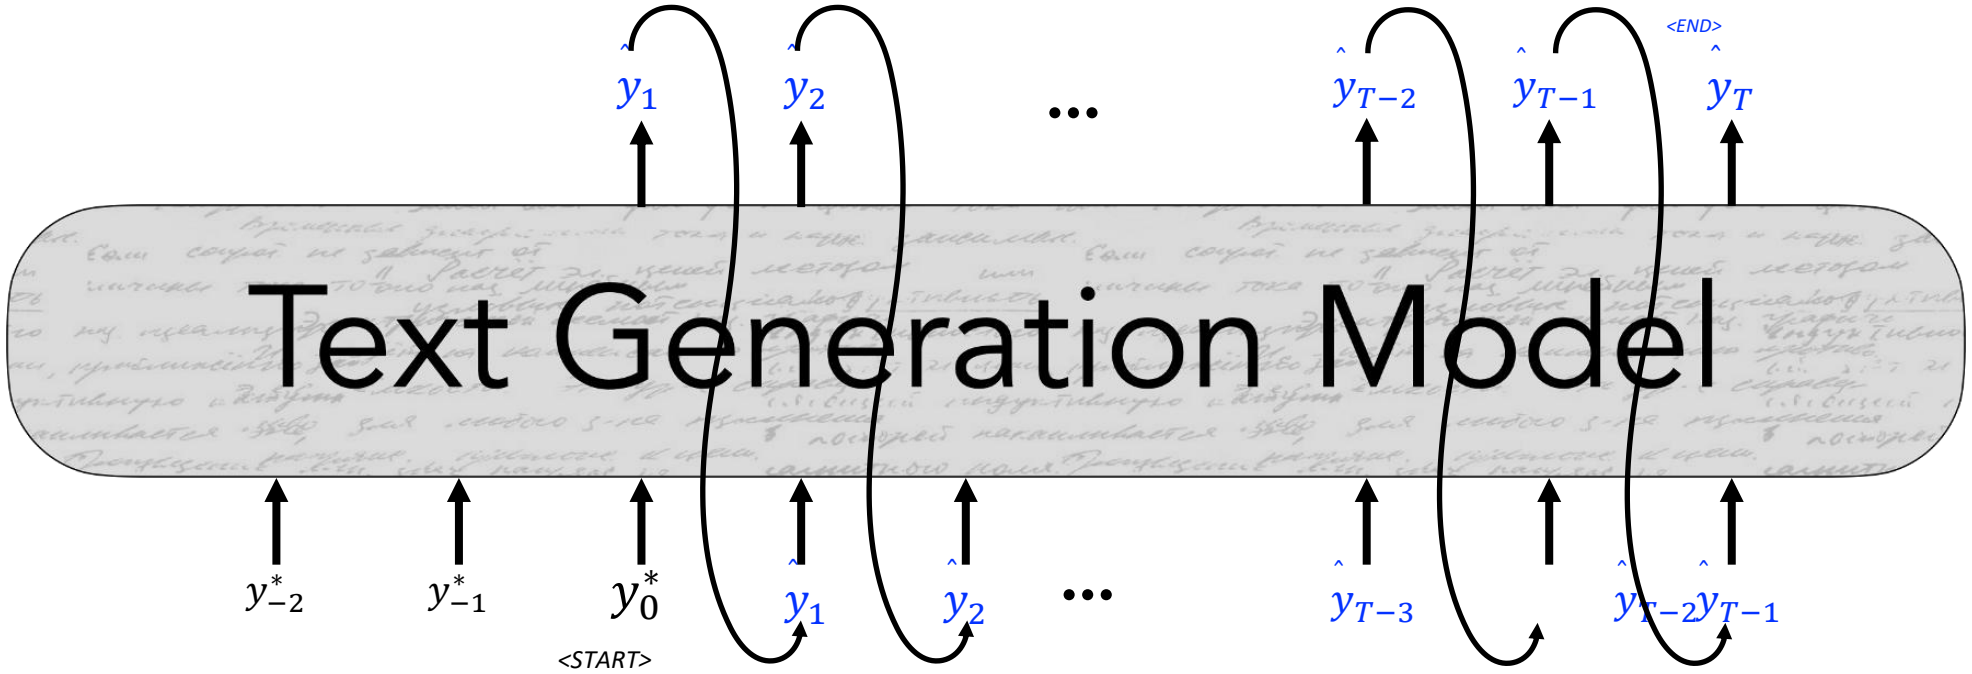

👱‍♀️ Can you help me write a resignation letter to my current employer, while leaving on good terms and expressing gratitude for the opportunities provided?

Review: Example Annotation

🤖 Here are two responses from the chatbot. (Please scroll down on the content to see the entire response if it is too long)





Annotator needs to choose whether they prefer A or B.

This is part of how ChatGPT was made

Two Classes of Dialog Systems

1. **Chatbots**
   - Systems designed for extended conversations
   - Chatting for fun and entertainment

2. **Task-Oriented Dialogue Agents**
   - Goal-Based Agents
   - Siri/Alexa, interface with robots, booking flights or hotels

Part II: Frame-based Dialogue Systems

- Task-based Dialogue Agents
- Based on “Domain Ontology”
  - A set of “Frames”
- Frame:
  - A knowledge structure representing user intentions
  - A collection of “slots”
  - Each “slot” having a set of “values”

Example: Travel Domain

Slot : Origin City

Type : City

Value : San Francisco

Slot	Type	Question
ORIGIN	city	What city are you leaving from?
DEST	city	Where are you going?
DEP DATE	date	What day would you like to leave?
DEP TIME	time	What time would you like to leave?
AIRLINE	line	What is your preferred airline?

Evaluation

1. Slot Error Rate for a Sentence

$$\frac{\text{\# of inserted/deleted/substituted slots}}{\text{\# of total reference slots for sentence}}$$

2. End-to-end evaluation (Task Success)

Evaluation

"Make an appointment with Lara at 10:30 in ITE 216"

Slot	Filler
PERSON	Lara
TIME	11:30 a.m.
ROOM	ITE 216

Slot error rate: 1/3

Task success: At end, was the correct meeting added to the calendar?

Frame-based Dialogue Agents

Artificial Intelligence Journal, 1977

GUS, A Frame-Driven Dialog System[1]

Daniel G. Bobrow, Ronald M. Kaplan, Martin Kay, Donald A. Norman, Henry Thompson and Terry Winograd

Xerox Palo Alto Research Center, 3333 Coyote Hill Road, Palo Alto, CA 94304, U.S.A.

- Still the industrial state of the art
- Siri/Alexa based on GUS architecture

GUS system : An actual dialogue

GUS:	Hello. My name is GUS. I can help you plan a simple trip by air. Where do you want to go ?
Client:	I want to go to San Diego on May 28
GUS:	What time do you want to leave?
Client:	I must be in San Diego before 10 am
GUS:	Would you like P.S.A. flight 102 that arrives at 9.15 am ?
Client:	What is the next flight ?
GUS:	Air California flight 310 that leaves at 8.30am and arrives at 10.00am
Client:	I'll take the first one
GUS:	In what name should I make a reservation, please?
Client:	Dan Bobrow
GUS:	I have confirmed the following flight: P.S.A. flight 102 on Wednesday May 28 from San Jose to San Diego that leaves at 7.30 am and arrives at 9.15 am.
GUS:	What date do you want to return on ?
Client:	On Friday in the evening.
GUS:	Would you like the flight that leaves at 7.45 pm ?
Client:	That's fine.
GUS:	I have confirmed the following flight: P.S.A. flight 307 on Friday May 30 from San Diego to San Jose that leaves at 7.45 pm and arrives at 9.30 pm Thank you for calling. Goodbye

Slot types can be complex

The type *DATE*

```
DATE
   MONTH:NAME    YEAR:INTEGER     DAY:(BOUNDED-INTEGER 1 31)
   WEEKDAY:(MEMBER (Sunday Monday Tuesday Wednesday
                    Thursday Friday Saturday))
```

How do you know what slot types or frames to use for your dialog system?

Control structure for frame-based dialog

Consider a trivial airline travel system:

- Ask the user for a departure city
- Ask for a destination city
- Ask for a time
- Ask whether the trip is round-trip or not

Rule-based Slot-filling

- Semantic Grammar Rules or Regular Expressions

```
Wake me (up) | set (the|an) alarm | get me up
```

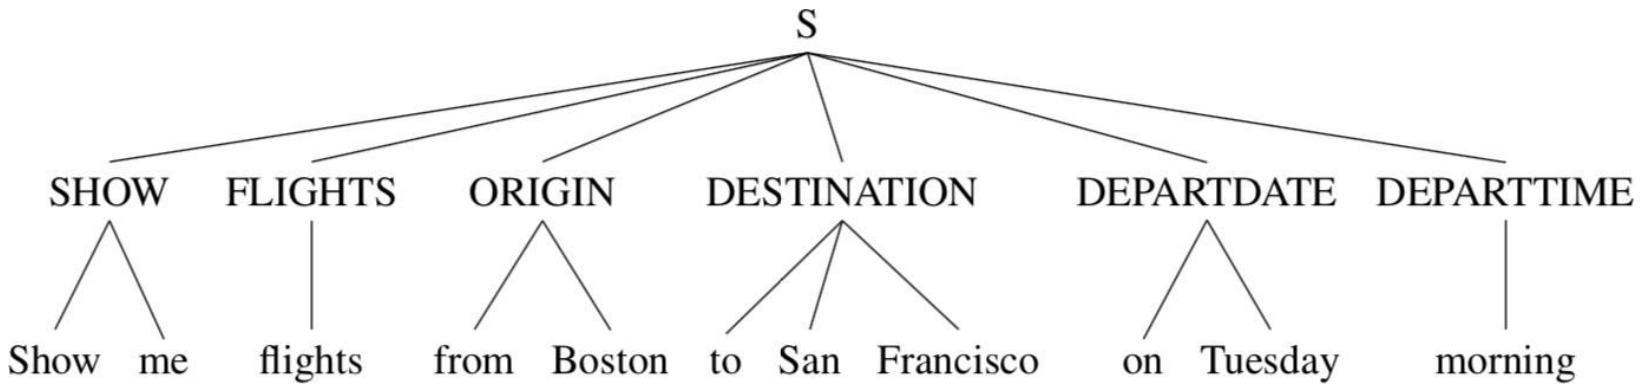

A semantic grammar parse for a user sentence, using slot names as the internal parse tree nodes

Rule Sets

- Collections of **rules** consisting of:
  - condition
  - action
- When user input is processed, facts added to store and
  - rule conditions are evaluated
  - relevant actions executed

Dialogue-State Architecture

More sophisticated version of frame-based architecture

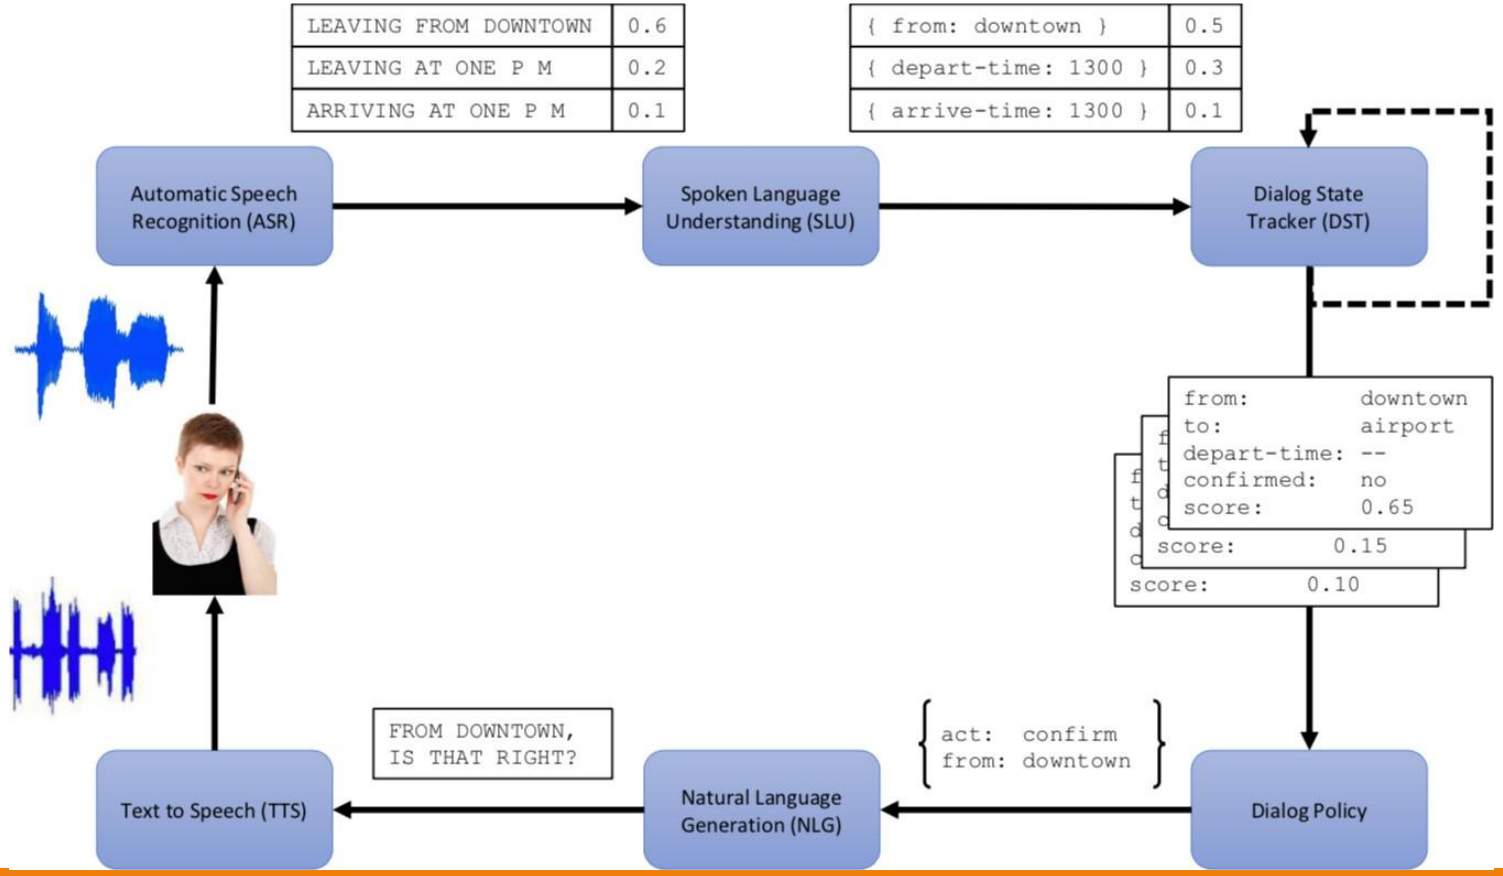

Williams et al. (2016)

NLU Component:

- Extract slot fillers using machine learning rather than rules

Dialogue State Tracker:

- Maintains current state of dialogue, user's most recent dialogue act

Dialogue policy:

- Decides what the system should do or say next
- When to answer user's questions, when to make a suggestion

Natural Language Generation Component:

- Condition on exact context to produce turns that seem much more natural

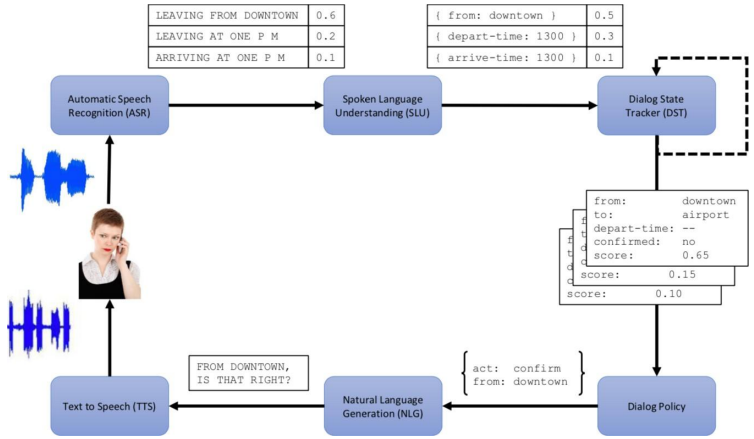

Dialogue State Tracking

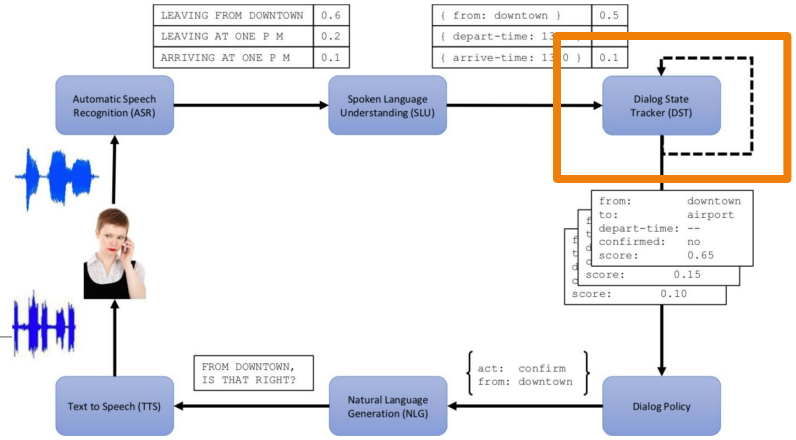

- Keep track of
  - Current state of the frame (the fillers of each slot)
  - User's most recent dialogue act

User:	I'm looking for a cheaper restaurant
	`inform(price=cheap)`
System:	Sure. What kind - and where?
User:	Thai food, somewhere downtown
	`inform(price=cheap, food=Thai, area=centre)`
System:	The House serves cheap Thai food
User:	Where is it?
	`inform(price=cheap, food=Thai, area=centre); request(address)`
System:	The House is at 106 Regent Street

Sample output of a dialogue state tracker after each turn

Dialogue Acts

Combining idea of speech acts and grounding into a single representation

	Utterance	Dialogue act
U:	Hi, I am looking for somewhere to eat.	`hello(task = find,type=restaurant)`
S:	You are looking for a restaurant. What type of food do you like?	`confreq(type = restaurant, food)`
U:	I'd like an Italian somewhere near the museum.	`inform(food = Italian, near=museum)`
S:	Roma is a nice Italian restaurant near the museum.	`inform(name = "Roma", type = restaurant, food = Italian, near = museum)`
U:	Is it reasonably priced?	`confirm(pricerange = moderate)`
S:	Yes, Roma is in the moderate price range.	`affirm(name = "Roma", pricerange = moderate)`
U:	What is the phone number?	`request(phone)`
S:	The number of Roma is 385456.	`inform(name = "Roma", phone = "385456")`
U:	Ok, thank you goodbye.	`bye()`

Sample dialogue from the recommender system of Young et al. (2010)

Dialogue Policy

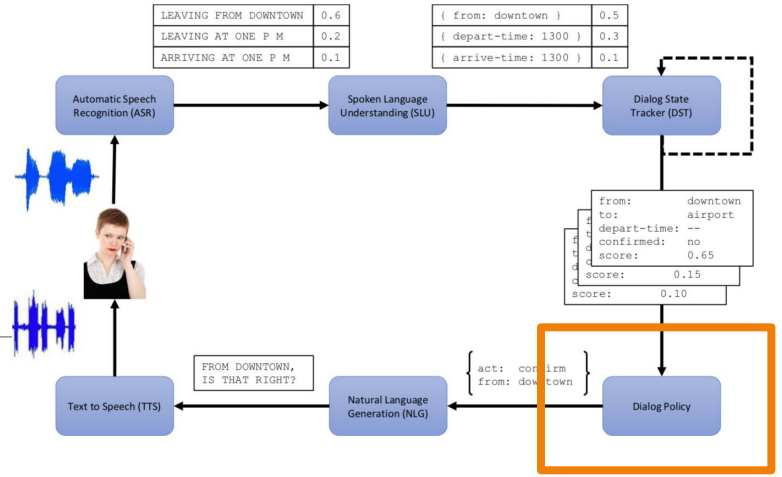

- What action the system should take next
- What dialogue act to generate
- Predict which action A_i to take

$$\hat{A}_i = \underset{A_i \in A}{\operatorname{argmax}} P(A_i | (A_1, U_1, \ldots, A_{i-1}, U_{i-1}))$$

A = Dialogue Acts from System; U = Dialogue Acts from User

- Simplification: Condition just on the current dialogue state

$$\hat{A}_i = \underset{A_i \in A}{\operatorname{argmax}} P(A_i | (Frame_{i-1}, A_{i-1}, U_{i-1}))$$

Think-Pair-Share

How would you know if your dialog policy* is good?

*What action the system should take next