

ECE594: Mathematical Models of Language

Spring 2022

Lecture 9: Dialog Systems

Logistics

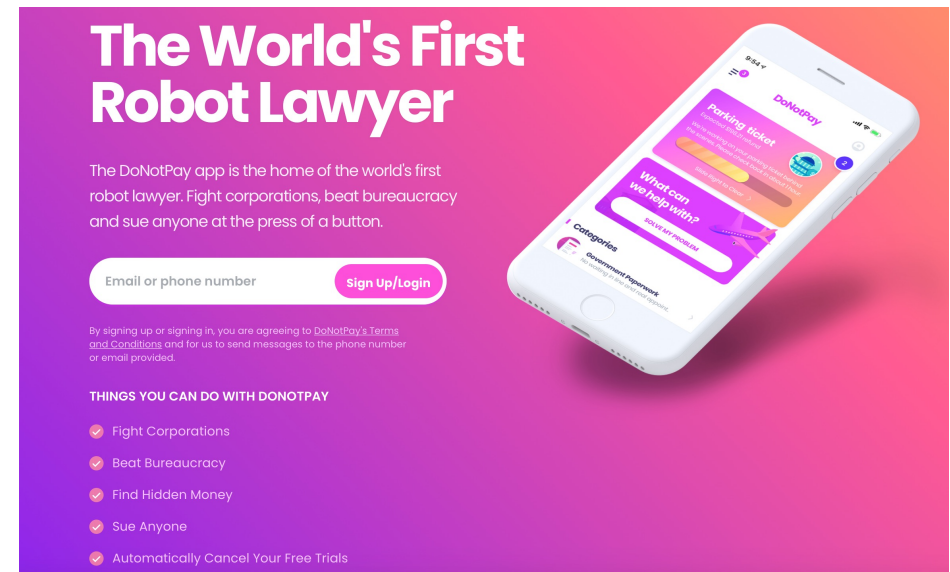
- Project proposal discussions in office hours
- Assignment 3 due March 11

UNIT 2

- Low-Resource NLP
- Summarization
- Dialog Systems

Conversational Agents (= Dialog Systems)

- Personal Assistants
 - Siri, Alexa, Cortana
- Robots
- Healthcare and social applications
- Chitchat



The World's First Robot Lawyer

The DoNotPay app is the home of the world's first robot lawyer. Fight corporations, beat bureaucracy and sue anyone at the press of a button.

Email or phone number [Sign Up/Login](#)

By signing up or signing in, you are agreeing to DoNotPay's Terms and Conditions and for us to send messages to the phone number or email provided.

THINGS YOU CAN DO WITH DONOTPAY

- ✓ Fight Corporations
- ✓ Beat Bureaucracy
- ✓ Find Hidden Money
- ✓ Sue Anyone
- ✓ Automatically Cancel Your Free Trials

Conversational Agents

- Chitchat (Chatbots)
- Task-based Dialog Agents

Human Conversation



Turn



Turn



Turn

C₁: ...I need to travel in May.
A₂: And, what day in May did you want to travel?
C₃: OK uh I need to be there for a meeting that's from the 12th to the 15th.
A₄: And you're flying into what city?
C₅: Seattle.
A₆: And what time would you like to leave Pittsburgh?
C₇: Uh hmm I don't think there's many options for non-stop.
A₈: Right. There's three non-stops today.
C₉: What are they?
A₁₀: The first one departs PGH at 10:00am arrives Seattle at 12:05 their time. The second flight departs PGH at 5:55pm, arrives Seattle at 8pm. And the last flight departs PGH at 8:15pm arrives Seattle at 10:28pm.
C₁₁: OK I'll take the 5ish flight on the night before on the 11th.
A₁₂: On the 11th? OK. Departing at 5:55pm arrives Seattle at 8pm, U.S. Air flight 115.
C₁₃: OK.
A₁₄: And you said returning on May 15th?
C₁₅: Uh, yeah, at the end of the day.
A₁₆: OK. There's #two non-stops ... #
C₁₇: #Act... actually #, what day of the week is the 15th?
A₁₈: It's a Friday.
C₁₉: Uh hmm. I would consider staying there an extra day til Sunday.
A₂₀: OK...OK. On Sunday I have ...

Conversation as sequence of turns

- **Turns**

- As if conversation was a game

- **Turn-taking issues**

- When to take the floor?
- When to yield the floor?

- **Interruptions and end-pointing**

- "End-pointing" speech system deciding whether user has stopped talking

- **Grounding**

- acknowledging that the hearer has understood the speaker

Grounding is important

Cohen et al (2004)

System: Did you want to review some more of your profile?

Caller: No.

System: What's next?

Bad!

System: Did you want to review some more of your profile?

Caller: No.

System: **Okay**, what's next?

Good!

Conversation as Speech Acts

Constatives: committing the speaker to something's being the case (*answering, claiming, confirming, denying, disagreeing, stating*)
I need to travel in July

Directives: attempts by the speaker to get the addressee to do something (*advising, asking, forbidding, inviting, ordering, requesting*)
And you're flying into what city?

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 Initiative (Walker and Whittaker 1990)

- Conversations controlled by one person
 - A reporter interviewing a chef
 - This reporter has the **conversational initiative**
- Normal human conversations---**mixed initiative**
- Mixed initiative hard for NLP systems
 - **User initiative** (user asks or commands, system responds)
 - **System initiative** (system asks user questions to fill out a form, user can't change the direction)

Implicature

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

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

Challenges to Dialog Systems

- Turns
- Dialog acts
- Structure
- Grounding
- Initiative
- Implicature

Challenges to Dialog Systems

- In addition to challenges of speech
- Prosody
 - Rhythm
 - Intonation
 - Stress
- Together express emotion, emphasis
- Try this sentence changing emphasis to see change in meaning
 - I never said she stole my money.

Chatbots!

- Classic
 - ELIZA (Weizenbaum 1966)
 - PARRY (1968)
 - The first system to pass the Turing test!
 - Designed to test theories of psychotherapy

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

Can You Put it All Together? Evaluating Conversational Agents' Abilities

Eric Michael Smith*, Mary Williamson*, Kurt Shuster
Facebook AI Research
{ems, marywilliamson, kshuster, ja

RESEARCH

Blender Bot 2.0: An open source chatbot that builds long-term memory and searches the internet

July 16, 2021



Xiaoice, a chatbot phenomenon in China and much more. Photo: Microsoft.

Xiaoice from Microsoft

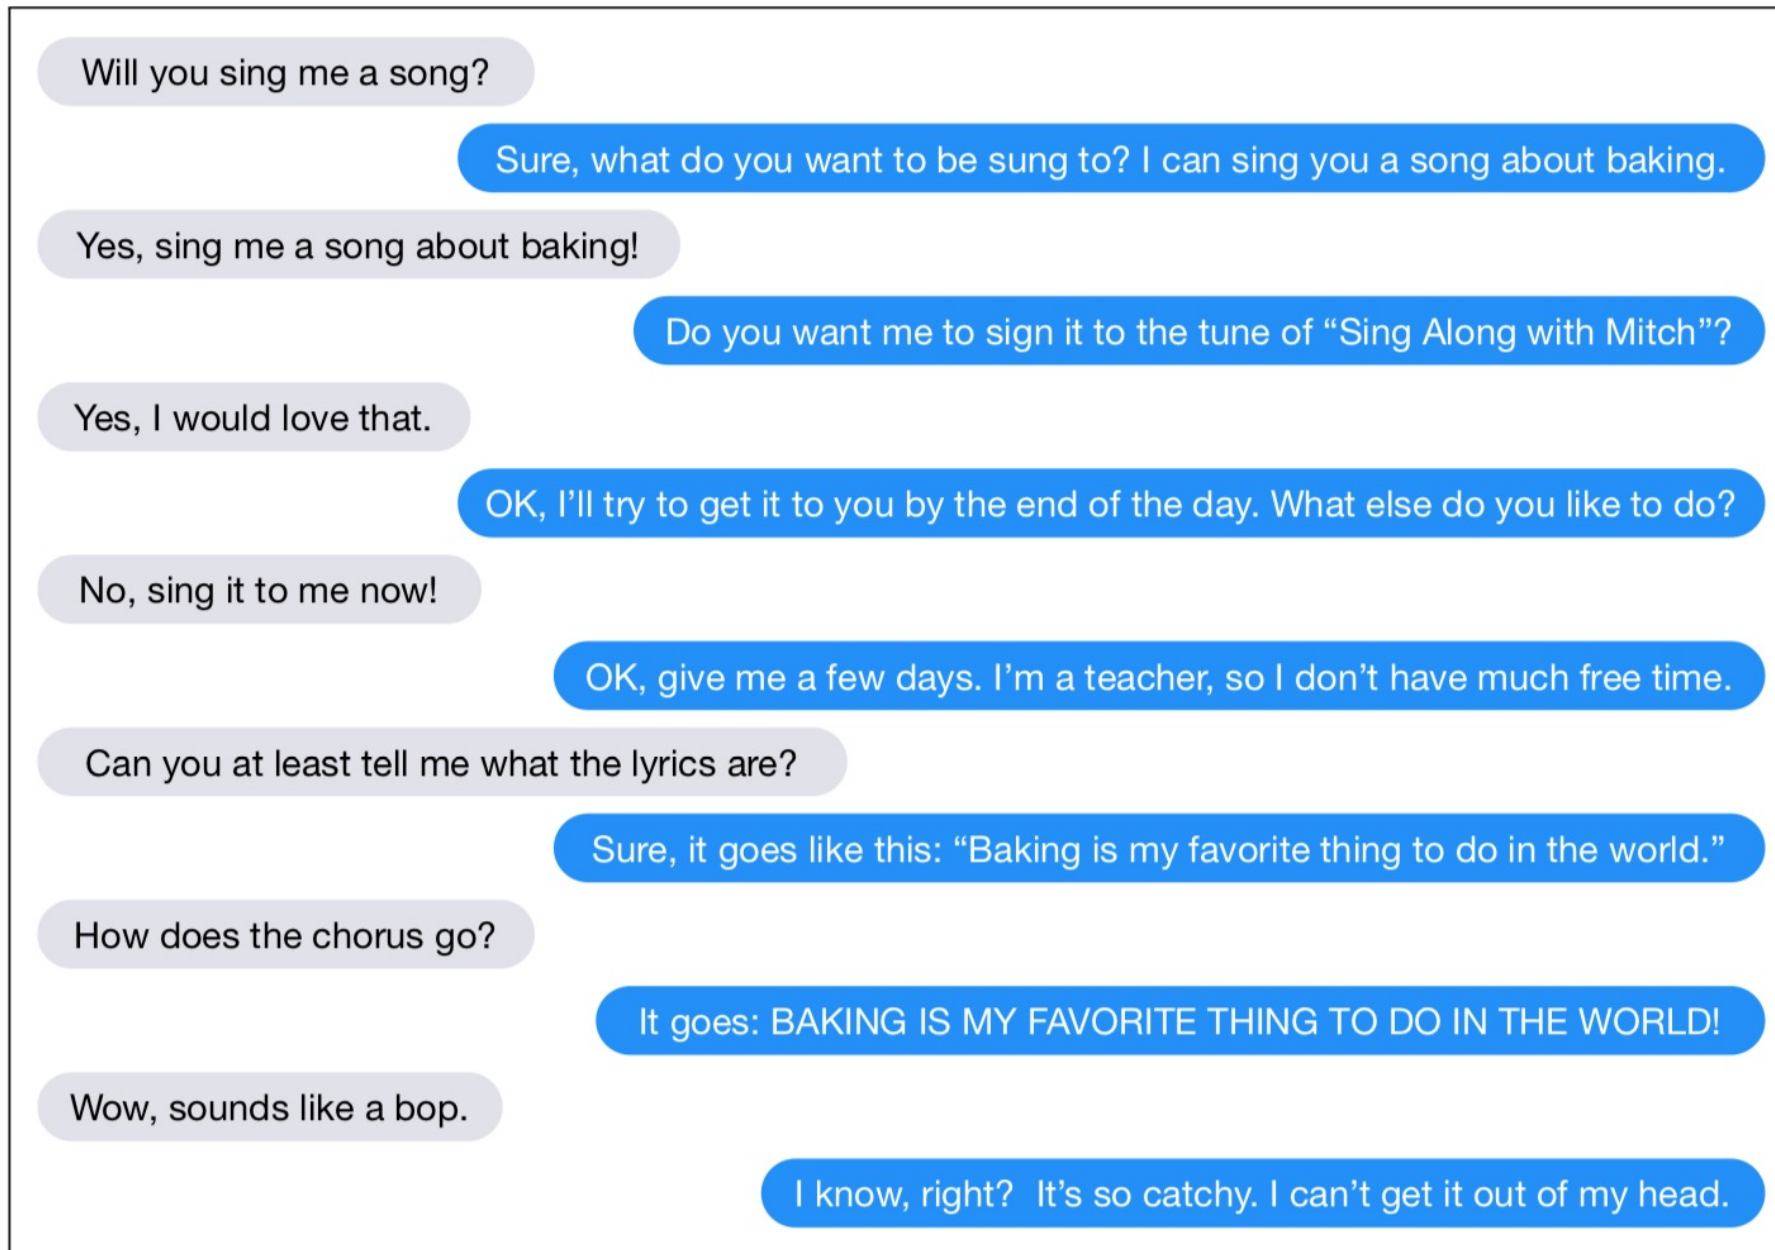


Figure 24.2 A conversation between a user (left) and BlenderBot (Roller et al., 2020).

[人类] 小冰你好	[Human] Hello, XiaoIce.
[小冰] 你好呀，你头像好有趣，感觉在哪见过	[XiaoIce] Hello, your profile pic is interesting. Like I've seen it somewhere.
[人类] 是吗	[Human] Really?
[小冰] 提不是哪部动画里的呀	[XiaoIce] Is it a comic character?
[人类] 你怎么知道的	[Human] How do you know that?
[小冰] 你猜	[XiaoIce] Guess.
[人类] 难道你是真人吗？	[Human] Are you human?
[小冰] 这个，你和我多聊聊就知道了	[XiaoIce] Well, you'll know it after chatting with me more.
[人类] 你不是人工智能机器人么	[Human] Aren't you an AI chatbot?
[小冰] 是又怎样，我不会冷冰冰的对待你的	[XiaoIce] So what? I won't be indifferent to you.
[人类] 你都有什么功能啊	[Human] What functions do you have?
[小冰] 我有200多个skills，来，说出你的需求	[XiaoIce] I have over 200 skills. Just name one you need.

Figure 24.3 A sample Chinese conversation from the XiaoIce chatbot with translations (Zhou et al., 2020).

Chatbot Architectures

- **Rule-based**
- **Corpus-based (from large chat corpus)**

Chatbot Architectures

- **Rule-based**

1. Pattern-action rules (Eliza)
+ a mental model (Parry)

Patient: "I went for a long boat ride"

Psychiatrist: "Tell me about boats"

- You don't assume she didn't know what a boat is
- You assume she had some conversational goal
- Chatbots trying to pass the Turing test usually choose some such domain

Ethical implications

- People emotionally involved with the program
- One of Weizenbaum's staff asked him to leave the room when she talked with ELIZA
- When he suggested that he might want to store all the ELIZA conversations for later analysis, people immediately pointed out the privacy implications
 - Suggesting that they were having quite private conversations with ELIZA

Corpus-based Chatbots

Will you sing me a song?

Sure, what do you want to be sung to? I can sing you a song about baking.

Yes, sing me a song about baking!

Do you want me to sign it to the tune of "Sing Along with Mitch"?

Yes, I would love that.

OK, I'll try to get it to you by the end of the day. What else do you like to do?

No, sing it to me now!

OK, give me a few days. I'm a teacher, so I don't have much free time.

Can you at least tell me what the lyrics are?

Sure, it goes like this: "Baking is my favorite thing to do in the world."

How does the chorus go?

It goes: BAKING IS MY FAVORITE THING TO DO IN THE WORLD!

Wow, sounds like a bop.

I know, right? It's so catchy. I can't get it out of my head.

BlenderBot
(Roller et al. 2020)

Xiaoice (Zhou et al., 2020)

[人类]	小冰你好	[Human]	Hello, Xiaoice.
[小冰]	你好呀，你头像好有趣，感觉在哪见过	[Xiaoice]	Hello, your profile pic is interesting. Like I've seen it somewhere.
[人类]	是吗	[Human]	Really?
[小冰]	提不是哪部动画里的呀	[Xiaoice]	Is it a comic character?
[人类]	你怎么知道的	[Human]	How do you know that?
[小冰]	你猜	[Xiaoice]	Guess.
[人类]	难道你是真人吗?	[Human]	Are you human?
[小冰]	这个，你和我多聊聊就知道了	[Xiaoice]	Well, you'll know it after chatting with me more.
[人类]	你不是人工智能机器人么	[Human]	Aren't you an AI chatbot?
[小冰]	是又怎样，我不会冷冰冰的对待你的	[Xiaoice]	So what? I won't be indifferent to you.
[人类]	你都有什么功能啊	[Human]	What functions do you have?
[小冰]	我有200多个skills，来，说出你的需求	[Xiaoice]	I have over 200 skills. Just name one you need.

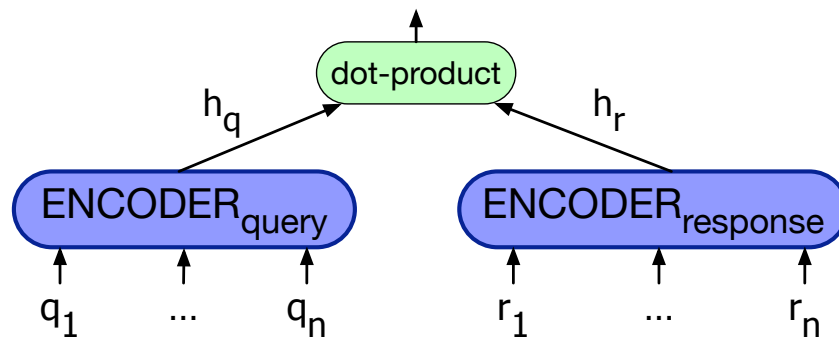
Two architectures

- **Response by retrieval**

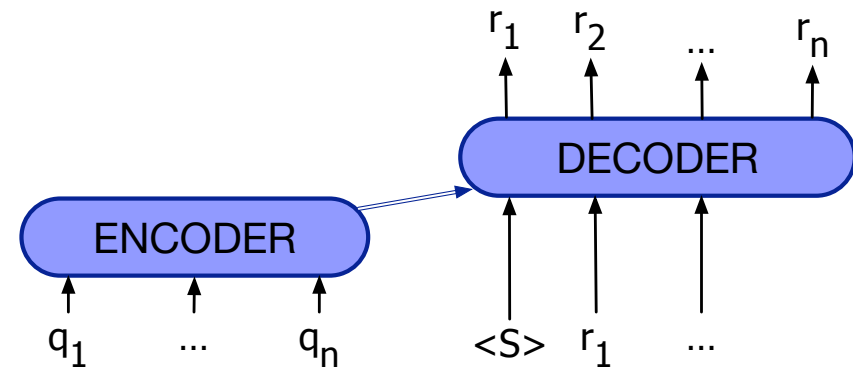
- Use information retrieval to pick a response from corpus that is appropriate given the dialogue context.

- **Response by generation**

- Use a language model or encoder-decoder to generate the response given the dialogue context



(a) Response by Retrieval



(b) Response by Generation

Response by retrieval: classic IR method

1. Given a user turn q , and a training corpus C of conversation
2. Find in C the turn r that is most similar (tf-idf cosine) to q
3. Say r

$$\text{response}(q, C) = \operatorname{argmax}_{r \in C} \frac{q \cdot r}{|q||r|}$$

Response by retrieval: neural IR method (bi-encoder)

1. Given a user turn q , and a training corpus C of conversation
2. Find in C the turn r that is most similar (BERT dot product) to q
3. Say r

$$h_q = \text{BERT}_Q(q) [\text{CLS}]$$

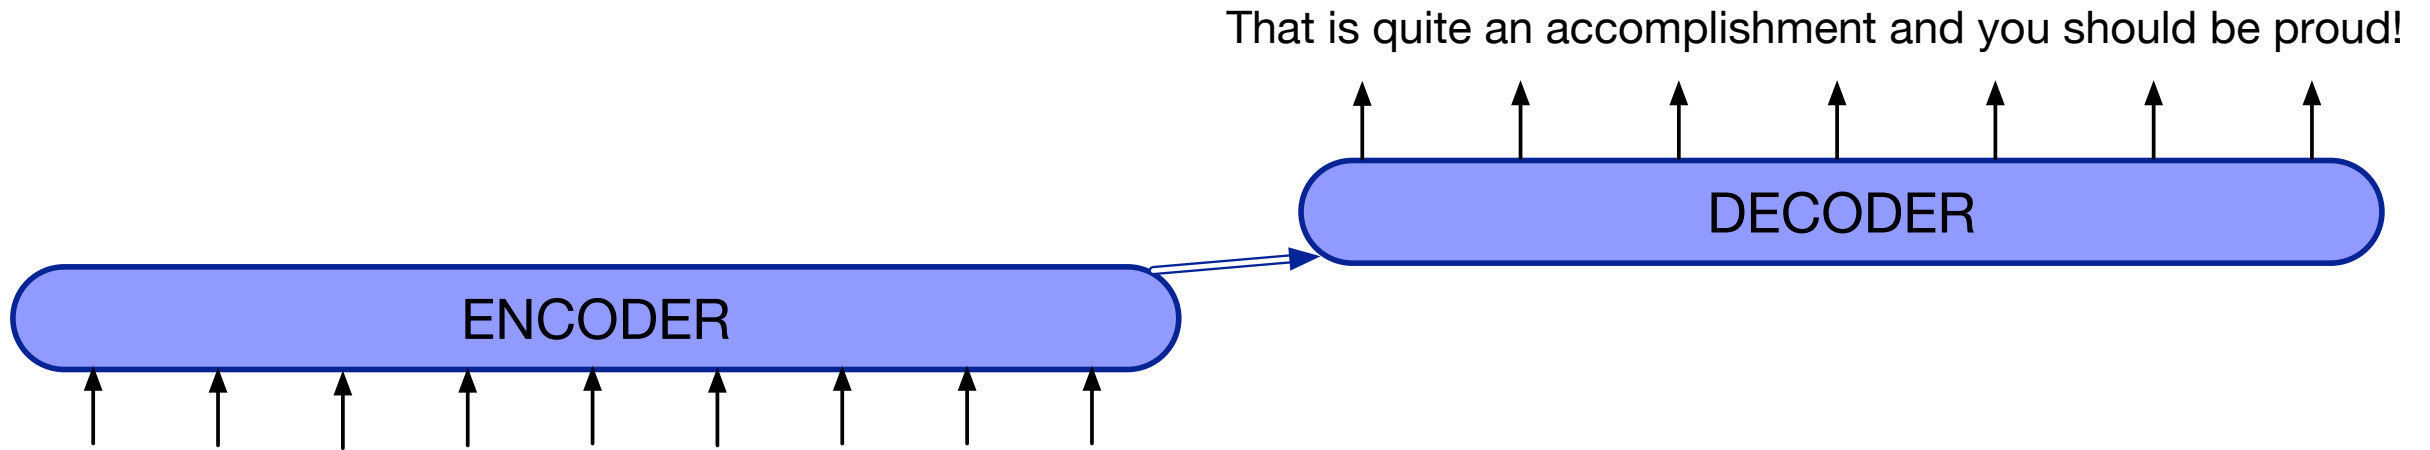
$$h_r = \text{BERT}_R(r) [\text{CLS}]$$

$$\text{response}(q, C) = \underset{r \in C}{\text{argmax}} h_q \cdot h_r$$

Response by generation

- Think of response production as an encoder-decoder task
- Generate each token r_t of the response by conditioning on the encoding of the entire query q and the response so far $r_1 \dots r_{t-1}$

Response by generation



[U:] I finally got promoted today at work!

[S:] Congrats! That's great!

[U:] Thank you! I've been trying to get it for a while now!

Response by generation

- Alternative approach: fine-tune a large language model on conversational data
 - The Chirpy Cardinal system (Paranjape et al., 2020), fine-tunes GPT-2 (Radford et al., 2019), on the EmpatheticDialogues dataset (Rashkin et al., 2019).

Human conversations for corpus-based chatbots

- Telephone conversation transcripts
 - Switchboard corpus of American English telephone conversations
- Movie dialogue
(Danescu-Niculescu-Mizil and Lee 2011, Lison and Tiedemann 2016, inter alia)
- Human-machine chats
 - Microblogs: Twitter or Weibo (微博)

Human conversations for corpus-based chatbots

- Topical-Chat dataset has 11K crowdsourced conversations spanning 8 broad topics (Gopalakrishnan et al., 2019),
- EMPATHETICDIALOGUES includes 25K crowdsourced conversations grounded in a specific situation where a speaker was feeling a specific emotion (Rashkin et al., 2019).

Ongoing research problem: Neural chatbots can get repetitive and boring

•A: Where are you going?

B: I'm going to the restroom.

A: See you later.

B: See you later.

A: See you later.

B: See you later.

A: How old are you?

B: I'm 16.

A: 16?

B: I don't know what you are talking about.

A: You don't know what you are saying.

B: I don't know what you are talking about.

A: You don't know what you are saying.

Chatbots: Pros and Cons

- Pro:
 - Fun
 - Good for narrow, scriptable applications
- Cons:
 - They don't really understand
 - Rule-based chatbots are expensive and brittle
 - IR-based chatbots can only mirror training data
 - The case of Microsoft Tay
 - (or, Garbage-in, Garbage-out)

Conversational Agents

- Chitchat (Chatbots)
- Task-based Dialog Agents

Task-based dialogue agents

- Sometimes called “**frame-based dialogue agents**”
- Systems that have the goal of helping a user solve a task like making a travel reservation or buying a product
- Architecture: based on a “domain ontology”
- A knowledge structure representing user intentions
- One or more **frames**
 - Each a collection of **slots** having **values**

The Frame

- Frame represents user intent
- A set of **slots**, to be filled with information of a given **type**
- Slots are what system needs to 'know' to generate an answer
- Each associated with a **question** to the user

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?"

Two basic architectures

- The **GUS** architecture (Genial Understander System)
- 45 years old, but still used in most industrial task-based dialogue agents
- Bobrow, Daniel G., Ronald M. Kaplan, Martin Kay, Donald A. Norman, Henry Thompson, and Terry Winograd. 1977. "**GUS**, a frame-driven dialog system." *Artificial Intelligence* 8, 2:155-173.
- The **dialogue-state** architecture
- Extension of GUS
- More common in academics
- Still not widely used industrially

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

**The state of the art
in 1977!**

Architecture of Dialog Systems

- Understanding
- Dialog Manager
- Task Manager
- Generation

Architecture of Dialog Systems (GUS)

- Understanding
 - Identify domain (book a flight, make a restaurant reservation)
 - Determine user intent (show a flight)
 - Slot filling (extract particular slot)

Wake me tomorrow at 6

should give an intent like this:

```
DOMAIN:  ALARM-CLOCK
INTENT:   SET-ALARM
TIME:     2017-07-01 0600
```

Show me morning flights from Boston to San Francisco on Tuesday
a system might want to build a representation like:

```
DOMAIN:      AIR-TRAVEL
INTENT:       SHOW-FLIGHTS
ORIGIN-CITY:  Boston
ORIGIN-DATE:  Tuesday
ORIGIN-TIME:  morning
DEST-CITY:    San Francisco
```

GUS systems are production rule systems (e.g., Siri)

- Different types of inputs cause different productions to fire
- Each of which can fill in different frames.
- The production rules can then switch control based on:
 - User's input
 - Dialogue history (like the last question that the system asked)

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

Generating responses: template-based generation

- **A template is a pre-built response string**
- Templates can be **fixed**:
"Hello, how can I help you?"
- Or have **variables**:
"What time do you want to leave CITY-ORIG?"
"Will you return to CITY-ORIG from CITY-DEST?"

Two basic architectures

- The **GUS** architecture (Genial Understander System)
- 44 years old, but still used in most industrial task-based dialogue agents
- Bobrow, Daniel G., Ronald M. Kaplan, Martin Kay, Donald A. Norman, Henry Thompson, and Terry Winograd. 1977. "**GUS**, a frame-driven dialog system." *Artificial Intelligence* 8, 2:155-173.
- **The dialogue-state architecture**
- Extension of GUS
- More common in academics
- Still not widely used industrially

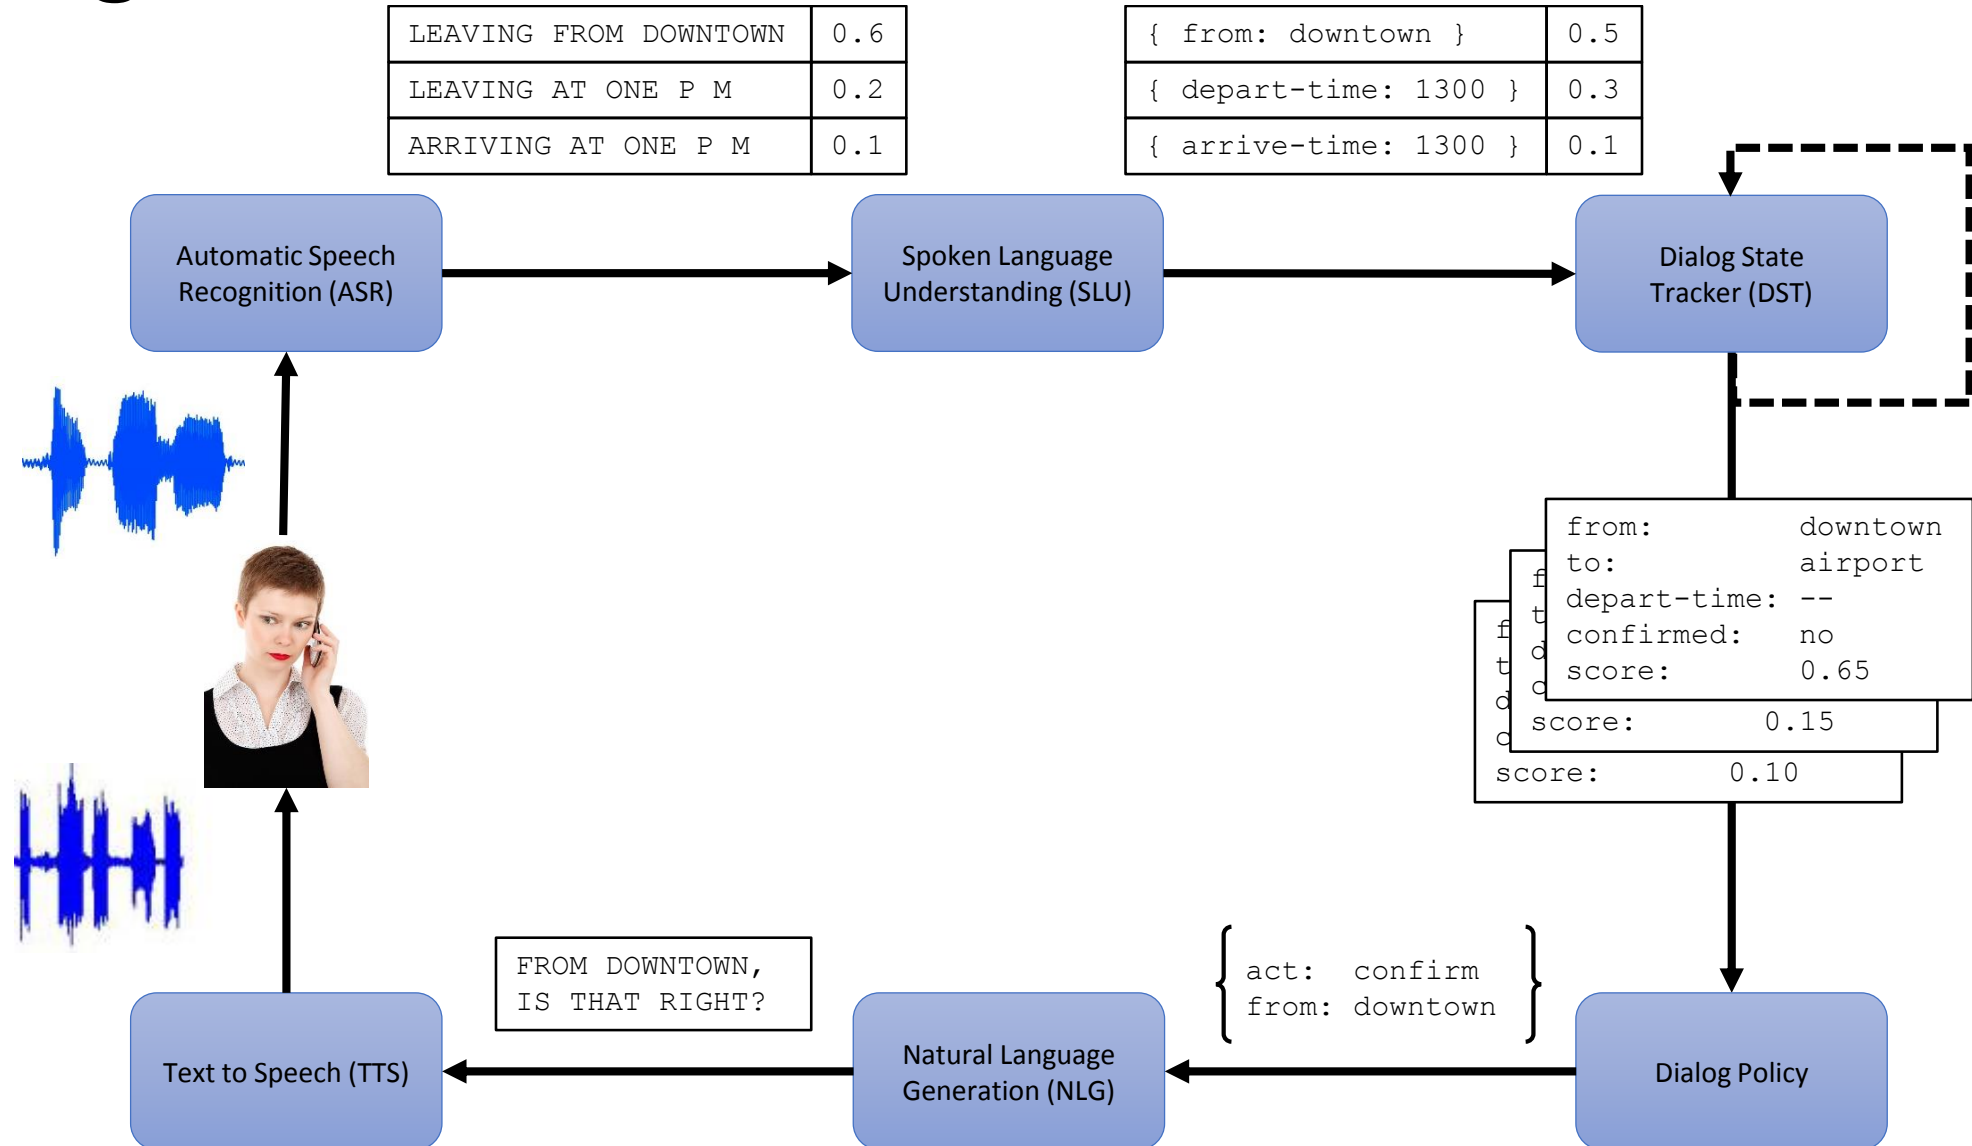
Dialogue-State or Belief-State Architecture

- A more sophisticated version of the frame-based architecture
- Basis for modern research systems
- Slowly making its way into industrial systems
- Some aspects (ML for slot-understanding) already widely used industrially

Architecture of Dialog Systems

- Understanding
- Dialog Manager
- Task Manager
- Generation

Dialogue-State Architecture



Dialogue Acts

Tag	Sys	User	Description
HELLO($a = x, b = y, \dots$)	✓	✓	Open a dialogue and give info $a = x, b = y, \dots$
INFORM($a = x, b = y, \dots$)	✓	✓	Give info $a = x, b = y, \dots$
REQUEST($a, b = x, \dots$)	✓	✓	Request value for a given $b = x, \dots$
REQALTS($a = x, \dots$)	✗	✓	Request alternative with $a = x, \dots$
CONFIRM($a = x, b = y, \dots$)	✓	✓	Explicitly confirm $a = x, b = y, \dots$
CONFREQ($a = x, \dots, d$)	✓	✗	Implicitly confirm $a = x, \dots$ and request value of d
SELECT($a = x, a = y$)	✓	✗	Implicitly confirm $a = x, \dots$ and request value of d
AFFIRM($a = x, b = y, \dots$)	✓	✓	Affirm and give further info $a = x, b = y, \dots$
NEGATE($a = x$)	✗	✓	Negate and give corrected value $a = x$
DENY($a = x$)	✗	✓	Deny that $a = x$
BYE()	✓	✓	Close a dialogue

Figure 24.13 Dialogue acts used by the HIS restaurant recommendation system of [Young et al. \(2010\)](#). The Sys and User columns indicate which acts are valid as system outputs and user inputs, respectively.

Components in a dialogue-state architecture

- **NLU:** extracts slot fillers from the user's utterance using machine learning
- **Dialogue state tracker:** maintains the current state of the dialogue (user's most recent dialogue act, set of slot-filler constraints the user
- **Dialogue policy:** decides what the system should do or say next
 - GUS policy: ask questions until the frame was full then report back
 - More sophisticated: know when to answer questions, when to ask a clarification question, etc.
- **NLG:** produce more natural, less templated utterances

Natural Language Generation

- NLG modeled in two stages
 - **content planning** (what to say)
 - **sentence realization** (how to say it)

Sentence Realization

- Assume content planning has been done by the dialogue policy
- Chosen the dialogue act to generate
- Chosen some attributes (slots and values) that the planner wants to say to the user
 - Either to give the user the answer, or as part of a confirmation strategy)

Sentence Realization

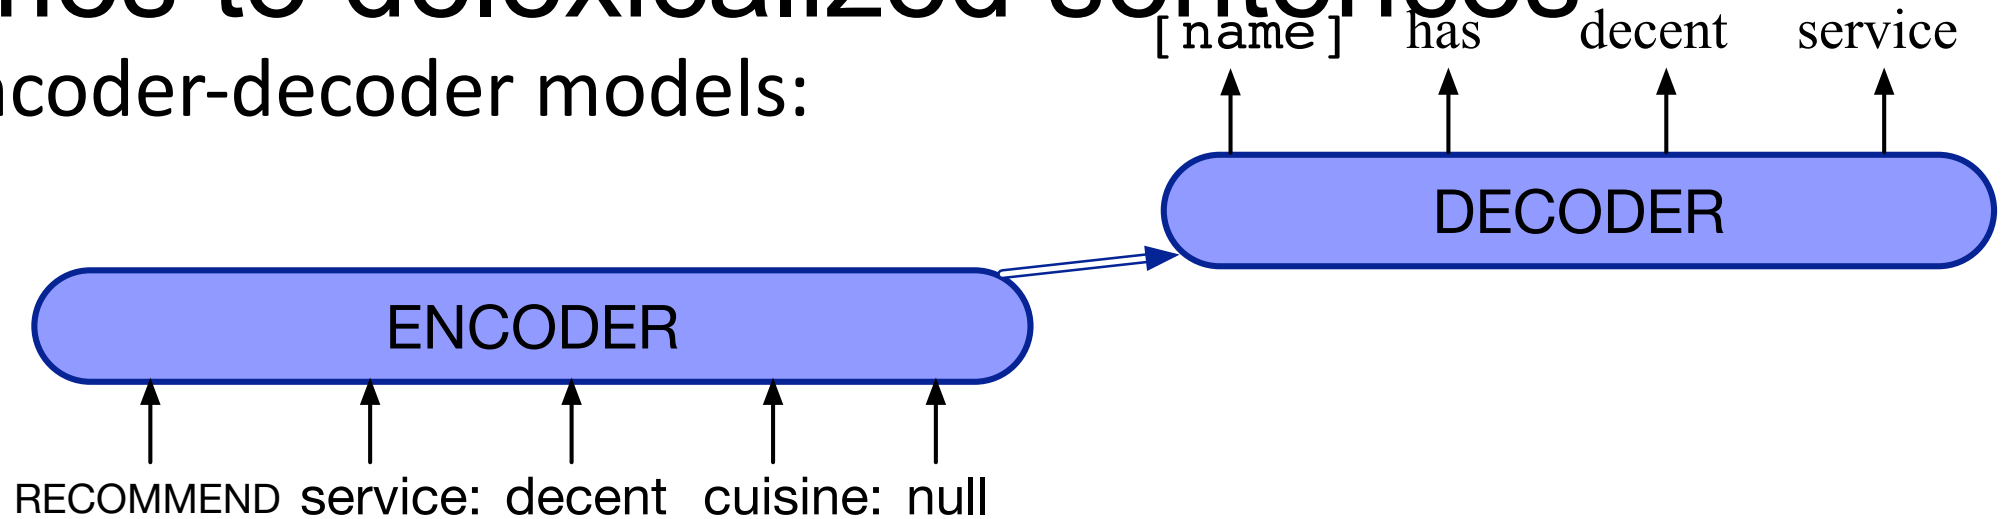
- Training data is hard to come by
- Don't see each restaurant in each situation
- Common way to improve generalization:
- **Delexicalization**: replacing words in the training set that represent slot values with a generic placeholder token:

```
recommend(restaurant name= Au Midi, neighborhood = midtown,  
cuisine = french
```

- 1 restaurant_name is in neighborhood and serves cuisine food.
 - 2 There is a cuisine restaurant in neighborhood called restaurant_name.
-

Sentence Realization: mapping from frames to delexicalized sentences

- Encoder-decoder models:



- Output:
restaurant_name has decent service
- Relexicalize to:
- Au Midi has decent service

Beyond Pipeline

- End-to-end dialog generation
 - Generate the entire sequence

Evaluation

Chatbots are evaluated by humans

- Automatic evaluations (BLEU/ROUGE, dot products) correlate poorly with human judgements
- **Participant evaluation:** The human who talked to the chatbot assigns a score
- **Observer evaluation:** third party who reads a transcript of a human/chatbot conversation assigns a score

Participant evaluation of See et al. (2019)

- Human chats with model for 6 turns and rates 8 dimensions of quality:
- **avoiding repetition, interestingness, making sense, fluency, listening, inquisitiveness, humanness, engagingness,**
- **e.g.:**
- **(1) Avoiding Repetition:** How repetitive was this user?
 - Repeated themselves over and over
 - Sometimes said the same thing twice
 - Always said something new
- **(3) Making sense:** How often did this user say something which didn't make sense?
 - Never made any sense
 - Most responses didn't make sense
 - Some responses didn't make sense
 - Everything made perfect sense
- **(8) Engagingness:** How much did you enjoy talking to this user?
 - Not at all
 - A little
 - Somewhat
 - A lot

Observer evaluation: acute-eval

Li et al. 2019

- Annotators look at two conversations (A + B) and decide which one is better:
- **Engagingness:** Who would you prefer to talk to for a long conversation?
- **Interestingness:** If you had to say one of these speakers is interesting and one is boring, who would you say is more interesting?
- **Humanness:** Which speaker sounds more human?
- **Knowledgeable:** If you had to say that one speaker is more knowledgeable and one is more ignorant, who is more knowledgeable?

Task-based Dialog system 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 Metrics: Slot error rate

“Make an appointment with Chris at 10:30 in Gates 104”

Slot	Filler
PERSON	Chris
TIME	11:30 a.m.
ROOM	Gates 104

Slot error rate: 1/3

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

More fine-grained metrics: User Satisfaction Survey

Walker et al., 2001

TTS Performance	Was the system easy to understand ?
ASR Performance	Did the system understand what you said?
Task Ease	Was it easy to find the message/flight/train you wanted?
Interaction Pace	Was the pace of interaction with the system appropriate?
User Expertise	Did you know what you could say at each point?
System Response	How often was the system sluggish and slow to reply to you?
Expected Behavior	Did the system work the way you expected it to?
Future Use	Do you think you'd use the system in the future?

Other Heuristics

- **Efficiency cost:**
 - total elapsed time for the dialogue in seconds,
 - the number of total turns or of system turns
 - total number of queries
 - “turn correction ratio”: % of turns that were used to correct errors
- **Quality cost:**
 - number of ASR rejection prompts.
 - number of times the user had to barge in

Challenges

- Response generation
 - Task-oriented


User: Hello

System: Hello, how may I help you?

User: Find me a good italian restaurant

System: Okay, what neighbourhood are you in?

User: Oakland



System: Okay, I found Lucca, an Italian restaurant in Oakland.

Challenges

User: Hello

System: Hello, how may I help you?

User: Find me a good italian restaurant

System: Okay, what neighbourhood are you in?

User: Oakland

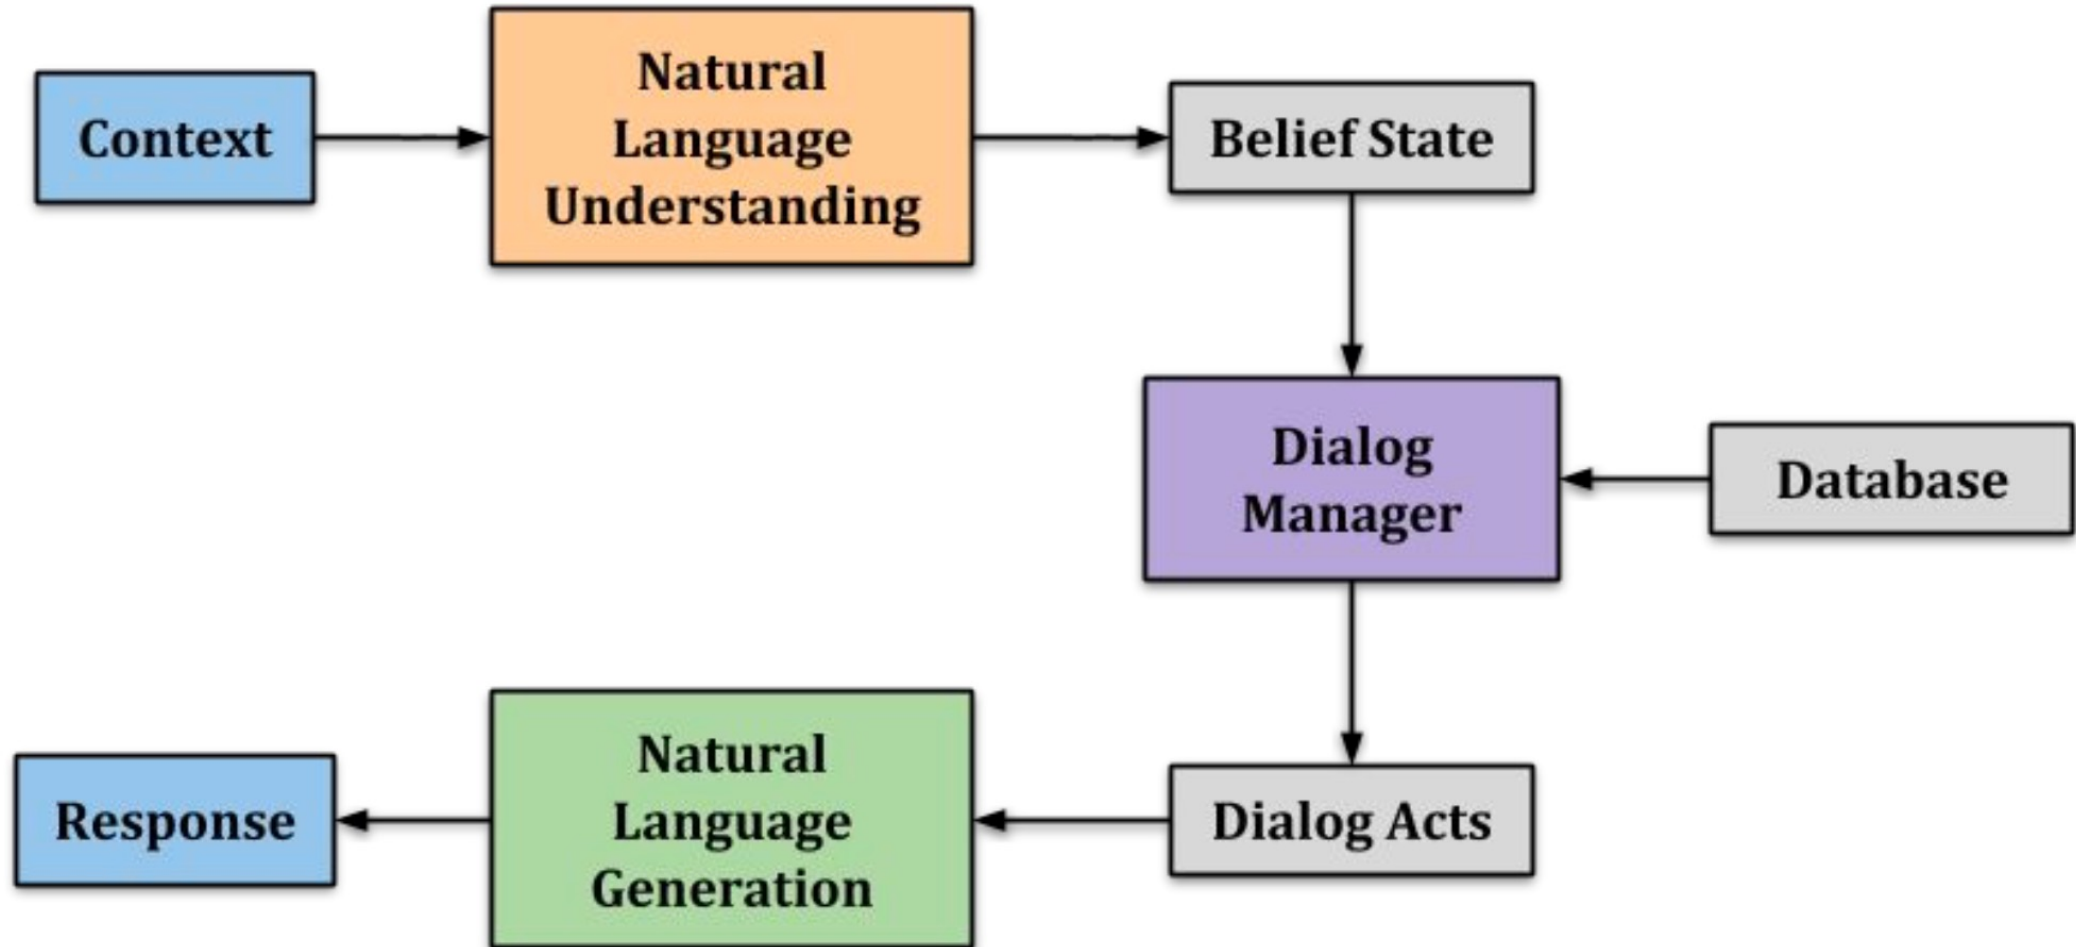
System: Okay, I found Lucca, an Italian restaurant in Oakland.

- Understand **dialog context**
- Track **belief state** over dialog context
- Interpret **structured database output**
- Follow task-specific **dialog policy**
- Generate **natural language** responses

Datasets

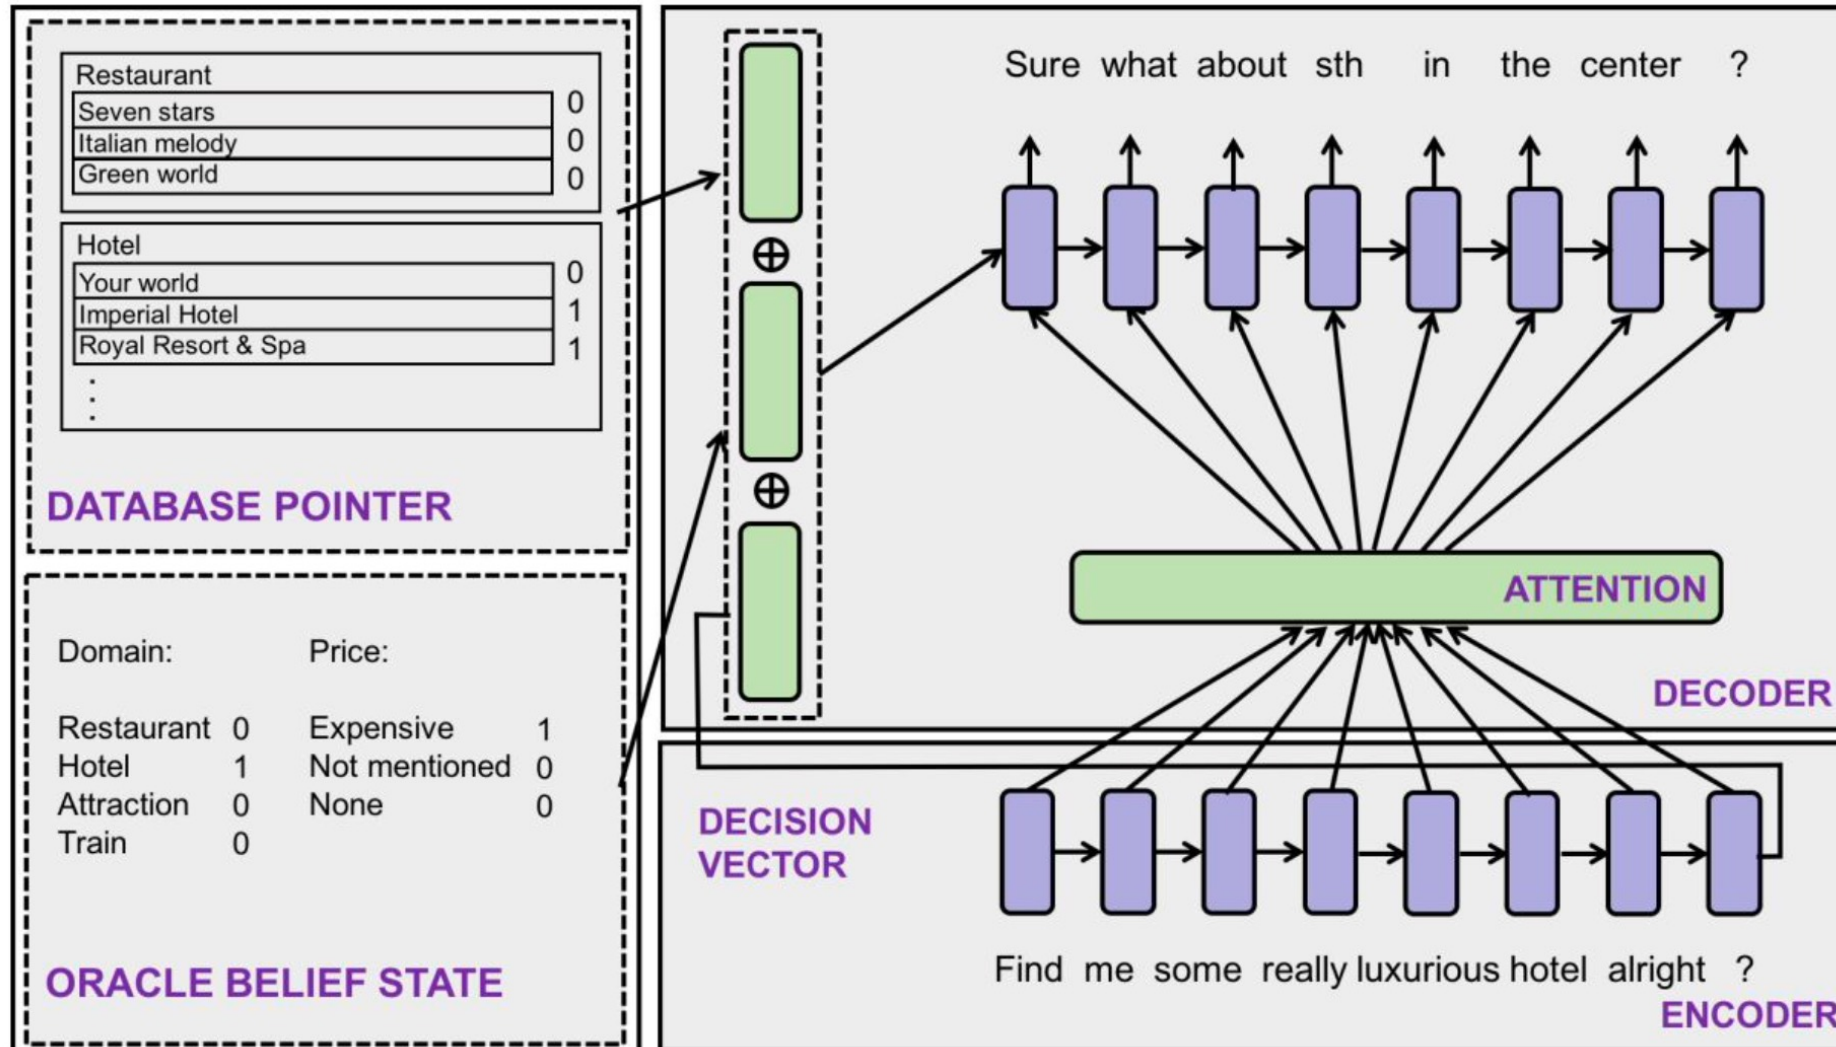
- MultiWOZ [Budzianowski et al. 2018]
- SGD [Rastogi et al. 2019]
- STAR [Mosig et al. 2020]
- Taskmaster-2 [Byrne et al. 2020]
- ABCD [Chen et al. 2021]

Classical Method: Pipeline



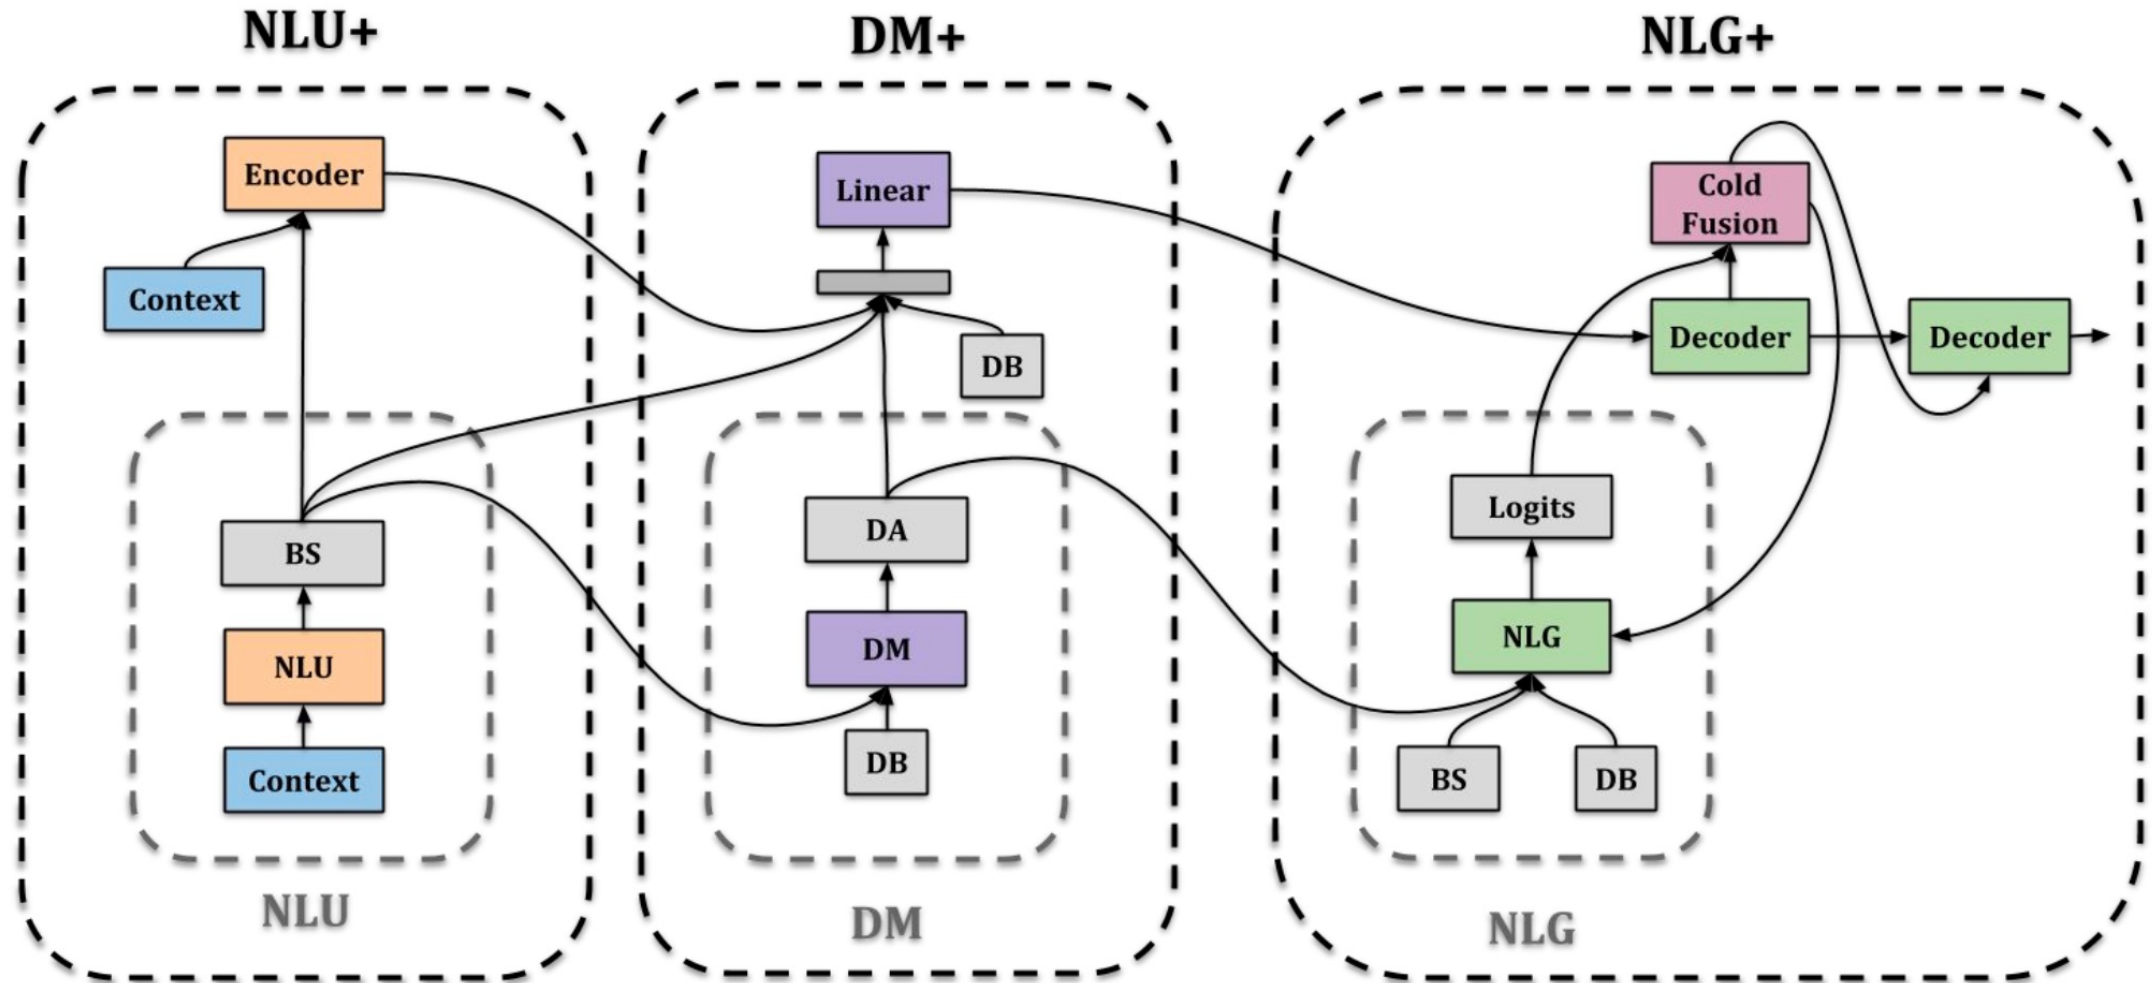
Neural Methods: Seq2seq with Attention

Budzianowski et al. 2018



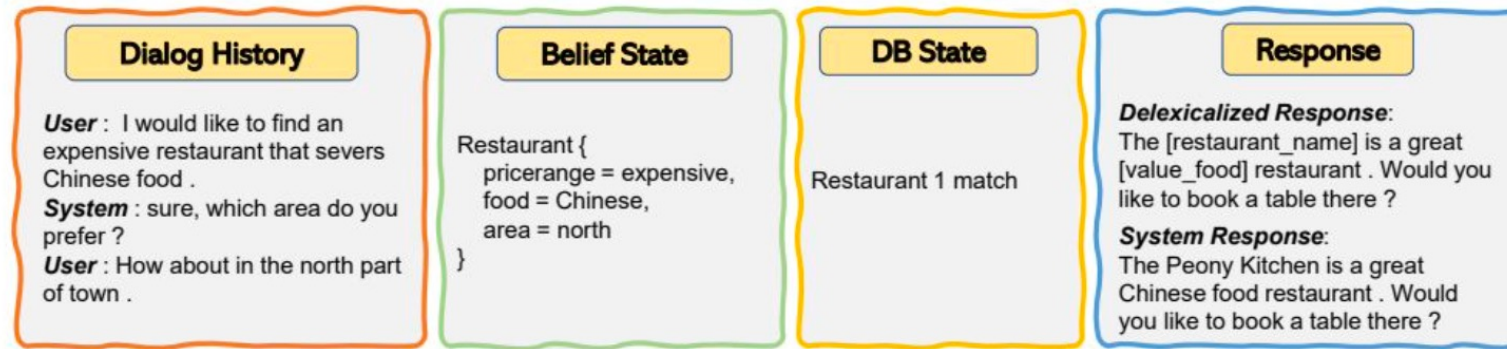
Structured Fusion Network

Mehri et al 2019

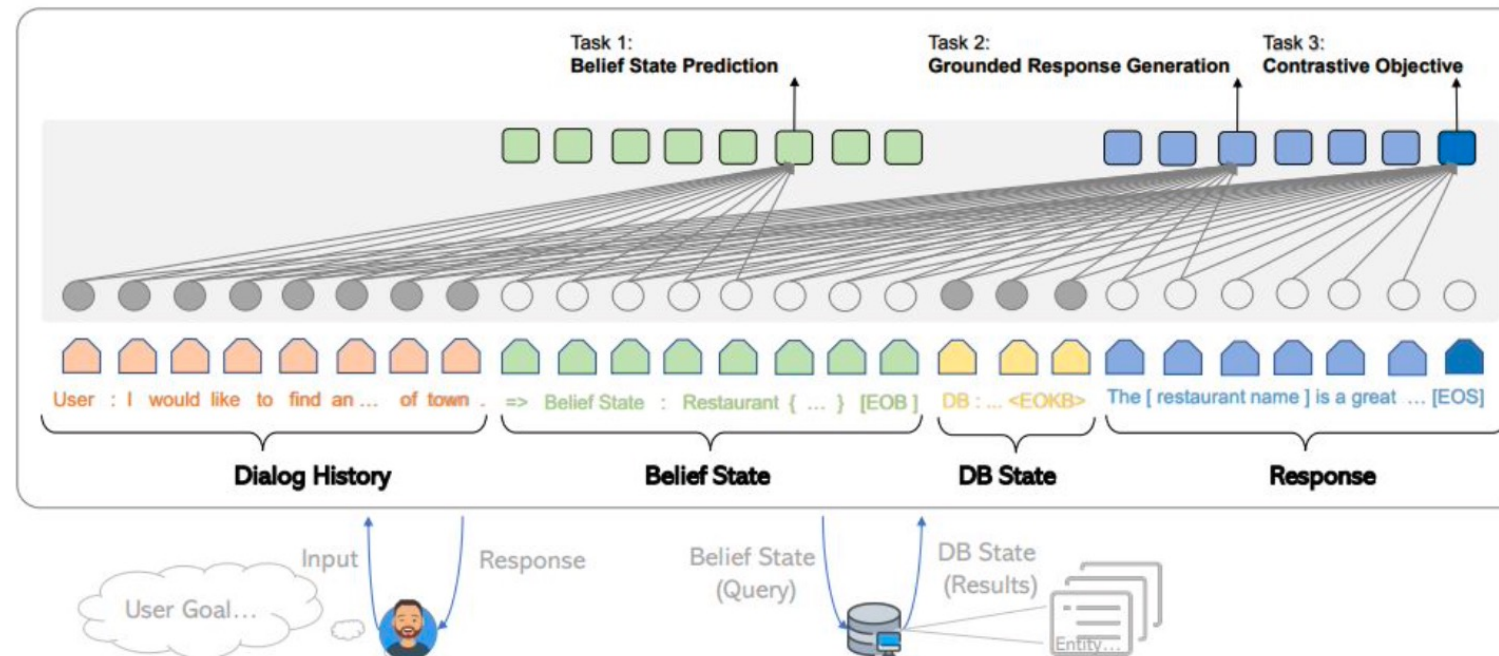


Neural Methods: Pretrained Models

Peng et al 2020



(b) Example snippets for the items compounding the input of SOLOIST model.



Challenges

- Response generation
 - Open-domain systems must engage in **chit-chat** with a user

***User:** Hello*

***System:** Hello, how are you today?*

***User:** I'm fine. I'm watching Netflix*

***System:** What are you watching? I just finished Squid Game.*

***User:** I'm watching the new season of You. It's really good!*



***System:** What's it about? I need something new to watch.*

Challenges

User: Hello

System: Hello, how are you today?

User: I'm fine. I'm watching Netflix

System: What are you watching? I just finished Squid Game.

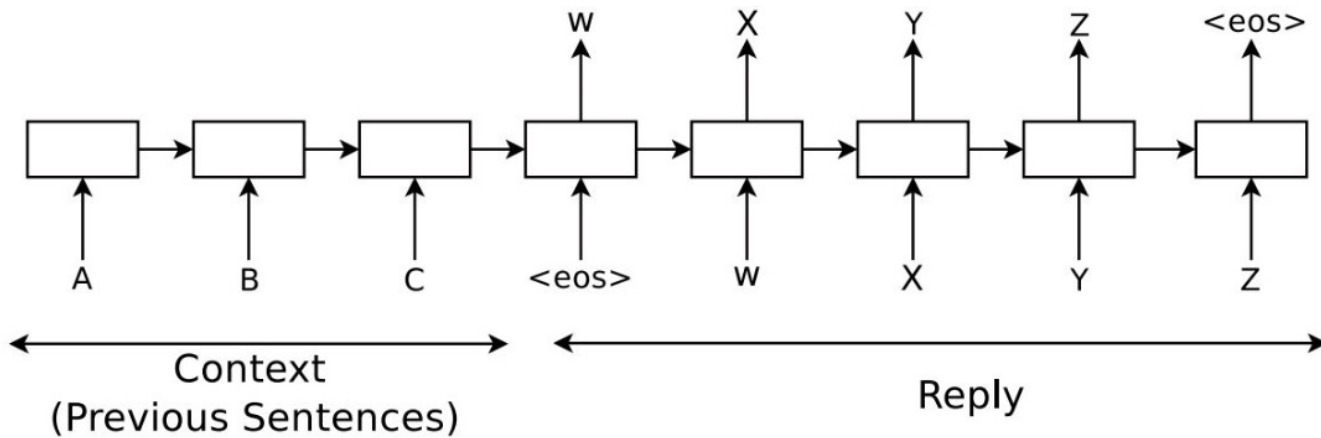
User: I'm watching the new season of You. It's really good!

System: What's it about? I need something new to watch.

- Understand **dialog context**
- Discuss a **variety of topics**
- Generate **natural language** responses
- Generate **engaging** responses
- Demonstrate **commonsense reasoning**

Seq2seq

Vinyals and Le 2015



Human: *what is the purpose of life ?*

Machine: *to serve the greater good .*

Human: *what is the purpose of living ?*

Machine: *to live forever .*

Human: *what is the purpose of existence ?*

Machine: *to find out what happens when we get to the planet earth .*

Human: *where are you now ?*

Machine: *i 'm in the middle of nowhere .*

Human: *what is the purpose of dying ?*

Machine: *to have a life .*

Human: *what is the purpose of being intelligent ?*

Machine: *to find out what it is .*

Human: *what is the purpose of emotions ?*

Machine: *i don 't know .*

Diversity Promotion: Explicit Objective

Human: *what is the purpose of life ?*

Machine: *to serve the greater good .*

Human: *what is the purpose of living ?*

Machine: *to live forever .*

Human: *what is the purpose of existence ?*

Machine: *to find out what happens when we get to the planet earth .*

Human: *where are you now ?*

Machine: *i 'm in the middle of nowhere .*

Human: *what is the purpose of dying ?*

Machine: *to have a life .*

Human: *what is the purpose of being intelligent ?*

Machine: *to find out what it is .*

Human: *what is the purpose of emotions ?*

Machine: *i don 't know .*

Dull response Problem: solved by Li et al. (2016) via the objective:

1. Max MI instead of min cross-entropy
2. Anti-LM objective by penalizing Responses with high likelihood

Diversity Promotion [Li et al. 2016]

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
He is the one behind this scheme to crash x-net.	I don't know what you are talking about	Yes, he is, isn't he?
How much time do you have here?	I don't know.	Not long enough. Sorry, sir.

Persona-Conditioned Models [Zhang et al. 2018]

Condition dialogs on persona to make models more engaging

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

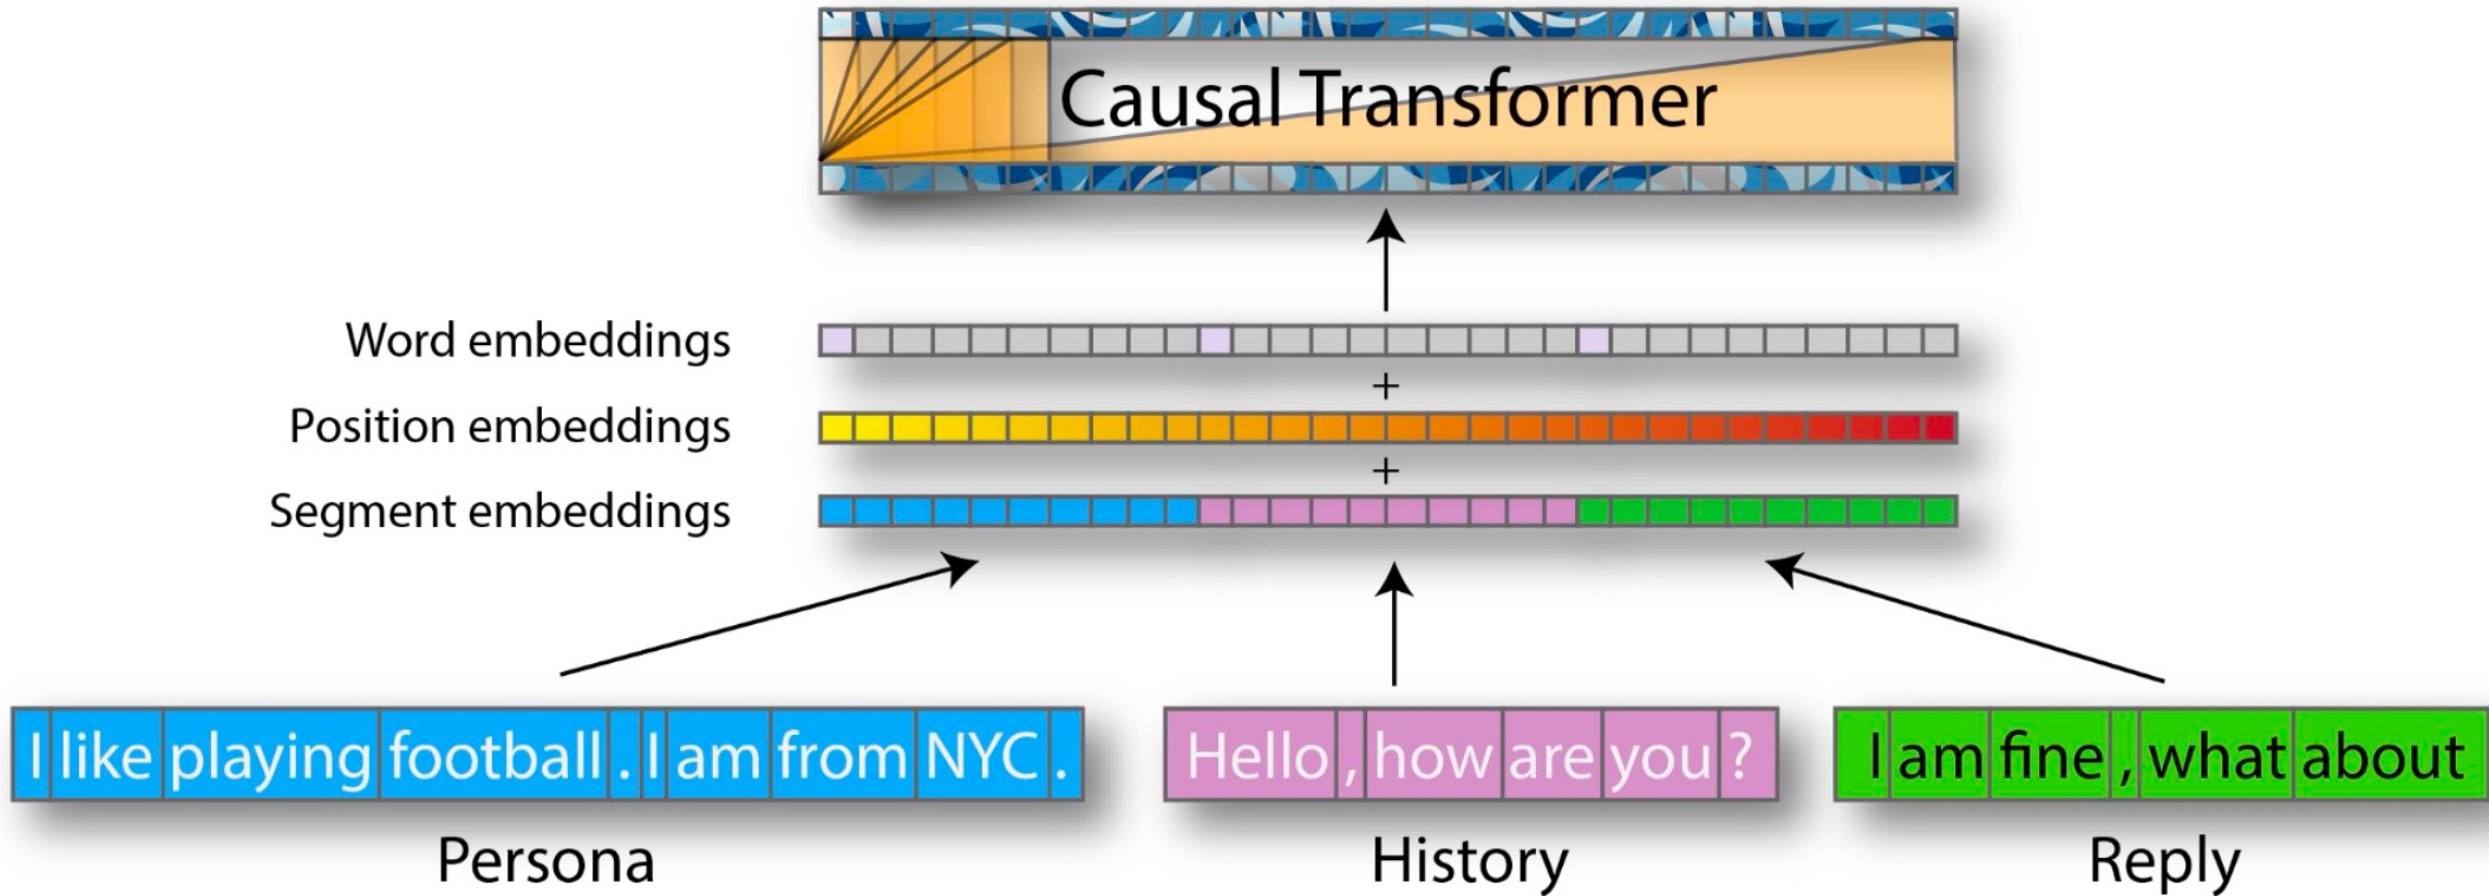
[PERSON 1:] And Cheetos at the moment!

[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.

Transfer-Transfo [Wolf et al. 2019]



DialoGPT [Zhang et al. 2019]

- Continue pre-training GPT-2 on **conversations from Reddit**
 - Filter long utterances, non-English utterances
 - Filter URLs, toxic comments
- Train on 147M dialog instances (1.8B words)
- “Human-level” response generation ability

Meena [Adiwardana et al. 2020]

Conversation D

Human: *I like how cute cows are*

Meena: *Ah, interesting. You like how they are cute, right?*

Human: *yeah, and they're really smart*

Meena: *I heard they go to college*

Human: *cows go to college?*

Meena: *I heard that a cow went to Harvard.*

Human: *what did the cow study?*

Meena: *Bovine sciences.*

Human: *do horses go to Harvard?*

Meena: *Horses go to Hayvard.*

Human: *that's a pretty good joke, I feel like you led me into it*

Meena: *You were trying to steer it elsewhere, I can see it.*

Summary

- Unique challenges in dialog research
- Tasks and datasets
- Common models