



Gridspace

GRIDSPACE IAP LECTURE 9  
DIALOG SYSTEMS

January 26, 2023

Use cases:

- chat bots
- task-oriented dialog systems

Established architecture:

- rule-based systems
- information retrieval systems
- encoder-decoder or decoder-only generators

## Outline

- human conversation
- chat bots
- frame-based task-oriented dialog systems
- large language models

# Human Conversation

## Turn Taking:

C<sub>1</sub>: ... I need to travel in May.  
A<sub>2</sub>: And, what day in May did you want to travel?  
C<sub>3</sub>: OK uh I need to be there for a meeting that's from the 12th to the 15th.  
A<sub>4</sub>: And you're flying into what city?  
C<sub>5</sub>: Seattle.  
A<sub>6</sub>: And what time would you like to leave Pittsburgh?  
C<sub>7</sub>: Uh hmm I don't think there's many options for non-stop.  
A<sub>8</sub>: Right. There's three non-stops today.  
C<sub>9</sub>: What are they?  
A<sub>10</sub>: 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.  
C<sub>11</sub>: OK I'll take the 5ish flight on the night before on the 11th.  
A<sub>12</sub>: On the 11th? OK. Departing at 5:55pm arrives Seattle at 8pm, U.S. Air flight 115.  
C<sub>13</sub>: OK.  
A<sub>14</sub>: And you said returning on May 15th?  
C<sub>15</sub>: Uh, yeah, at the end of the day.  
A<sub>16</sub>: OK. There's #two non-stops ... #  
C<sub>17</sub>: #Act... actually #, what day of the week is the 15th?  
A<sub>18</sub>: It's a Friday.  
C<sub>19</sub>: Uh hmm. I would consider staying there an extra day til Sunday.  
A<sub>20</sub>: OK... OK. On Sunday I have ...

C<sub>1</sub>: ... I need to travel in May.  
A<sub>2</sub>: And, what day in May did you want to travel?  
C<sub>3</sub>: OK uh I need to be there for a meeting that's from the 12th to the 15th.  
A<sub>4</sub>: And you're flying into what city?  
C<sub>5</sub>: Seattle.  
A<sub>6</sub>: And what time would you like to leave Pittsburgh?  
C<sub>7</sub>: Uh hmm I don't think there's many options for non-stop.  
A<sub>8</sub>: Right. There's three non-stops today.  
C<sub>9</sub>: What are they?  
A<sub>10</sub>: 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.  
C<sub>11</sub>: OK I'll take the 5ish flight on the night before on the 11th.  
A<sub>12</sub>: On the 11th? OK. Departing at 5:55pm arrives Seattle at 8pm, U.S. Air  
flight 115.  
C<sub>13</sub>: OK.  
A<sub>14</sub>: And you said returning on May 15th?  
C<sub>15</sub>: Uh, yeah, at the end of the day.  
A<sub>16</sub>: OK. There's #two non-stops ... #  
C<sub>17</sub>: #Act... actually #, what day of the week is the 15th?  
A<sub>18</sub>: It's a Friday.  
C<sub>19</sub>: Uh hmm. I would consider staying there an extra day til Sunday.  
A<sub>20</sub>: OK... OK. On Sunday I have ...



endpointing

C<sub>1</sub>: ... I need to travel in May.  
A<sub>2</sub>: And, what day in May did you want to travel?  
C<sub>3</sub>: OK uh I need to be there for a meeting that's from the 12th to the 15th.  
A<sub>4</sub>: And you're flying into what city?  
C<sub>5</sub>: Seattle.  
A<sub>6</sub>: And what time would you like to leave Pittsburgh?  
C<sub>7</sub>: Uh hmm I don't think there's many options for non-stop.  
A<sub>8</sub>: Right. There's three non-stops today.  
C<sub>9</sub>: What are they?  
A<sub>10</sub>: 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.  
C<sub>11</sub>: OK I'll take the 5ish flight on the night before on the 11th.  
A<sub>12</sub>: On the 11th? OK. Departing at 5:55pm arrives Seattle at 8pm, U.S. Air  
flight 115.  
C<sub>13</sub>: OK.  
A<sub>14</sub>: And you said returning on May 15th?  
C<sub>15</sub>: Uh, yeah, at the end of the day.  
A<sub>16</sub>: OK. There's #two non-stops ... #  
C<sub>17</sub>: #Act... actually #, what day of the week is the 15th?  
A<sub>18</sub>: It's a Friday.  
C<sub>19</sub>: Uh hmm. I would consider staying there an extra day til Sunday.  
A<sub>20</sub>: OK... OK. On Sunday I have ...

interruptions



C<sub>1</sub>: ... I need to travel in May.  
A<sub>2</sub>: And, what day in May did you want to travel?  
C<sub>3</sub>: OK uh I need to be there for a meeting that's from the 12th to the 15th.  
A<sub>4</sub>: And you're flying into what city?  
C<sub>5</sub>: Seattle.  
A<sub>6</sub>: And what time would you like to leave Pittsburgh?  
C<sub>7</sub>: Uh hmm I don't think there's many options for non-stop.  
A<sub>8</sub>: Right. There's three non-stops today.  
C<sub>9</sub>: What are they?  
A<sub>10</sub>: 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.  
C<sub>11</sub>: OK I'll take the 5ish flight on the night before on the 11th.  
A<sub>12</sub>: On the 11th? OK. Departing at 5:55pm arrives Seattle at 8pm, U.S. Air  
flight 115.  
C<sub>13</sub>: OK.  
A<sub>14</sub>: And you said returning on May 15th?  
C<sub>15</sub>: Uh, yeah, at the end of the day.  
A<sub>16</sub>: OK. There's #two non-stops ... #  
C<sub>17</sub>: #Act... actually #, what day of the week is the 15th?  
A<sub>18</sub>: It's a Friday.  
C<sub>19</sub>: Uh hmm. I would consider staying there an extra day til Sunday.  
A<sub>20</sub>: OK... OK. On Sunday I have ...



grounding

## Dialog Acts

- Constatives
- Directives
- Commissives
- Acknowledgments

## Dialog Acts

- Constatives:
  - I need to travel in May
  - My date of birth is January 26, 2023
- Directives
- Commissives
- Acknowledgments

## Dialog Acts

- Constatives
- Directives
  - Tell me where you are traveling to.
  - What are the last 4 digits of your social security number?
- Commissives
- Acknowledgments

## Dialog Acts

- Constatives
- Directives
- Commissives
  - I will get back to you by next week.
  - We'll go to the park this weekend.
- Acknowledgments

## Dialog Acts

- Constatives
- Directives
- Commissives
- Acknowledgments
  - Oh, sorry about that, got it now.
  - Okay!

## Conversational Analysis: adjacency pairs

## Conversational Analysis: adjacency pairs

- question -> answer



## Conversational Analysis: adjacency pairs

- question -> answer
- proposal -> acceptance or rejection

## Conversational Analysis: adjacency pairs

- question -> answer
- proposal -> acceptance or rejection
- compliments -> downplayers

## **Conversational Analysis:** side sequences

- subdialogs or side sequences separate the first and second parts of a pair

## Conversational Analysis: side sequences

- subdialogs (aka side sequences) separate the first and second parts of a pair
- examples:
  - corrections
  - clarification

## **Conversational Analysis:** side sequences

(without) correction subdialogs:

- Can I get a plane ticket for the 27th of January?
- Sure, we have one flight at 10am, does that work?

## Conversational Analysis: side sequences

(with) correction subdialogs:

- Can I get a plane ticket for the 27th of January?
- Sure, we have...
  - Wait, is that a Friday? Can i do the following Monday instead?
  - Alright, so I should look for flights on the 30th?
  - Yes please!
- We have one flights at 11am, does that work?

## Initiative

- interview: reporter has initiative, interviewee responds

## Initiative

- interviews: reporter has initiative, interviewee responds
- usually in human conversations: mixed initiative



## Initiative

- interviews: reporter has initiative, interviewee responds
- usually in human conversations: mixed initiative
- dialog systems: mixed initiative is much harder

## Initiative

- interviews: reporter has initiative, interviewee responds
- usually in human conversations: mixed initiative
- dialog systems: mixed initiative is much harder
- system-initiative (form-filling) can be very frustrating

## Inference and Implicature

- And, what day in May did you want to travel?
- OK uh I need to be there for a meeting that's from the 12th to the 15th.

## Inference and Implicature

- And, what day in May did you want to travel?
- OK uh I need to be there for a meeting that's from the 12th to the 15th.

Agent infers they want to travel on the 11th.

## Properties of Human Conversation

- turns
- grounding
- speech acts
- Dialog structure
- initiative
- implicature

# Chatbots

M I N D  
A QUARTERLY REVIEW  
OF  
PSYCHOLOGY AND PHILOSOPHY



I.—COMPUTING MACHINERY AND  
INTELLIGENCE

BY A. M. TURING

1. *The Imitation Game.*

I PROPOSE to consider the question, 'Can machines think?' This should begin with definitions of the meaning of the terms 'machine' and 'think'. The definitions might be framed so as to reflect so far as possible the normal use of the words, but this attitude is dangerous. If the meaning of the words 'machine' and 'think' are to be found by examining how they are commonly used it is difficult to escape the conclusion that the meaning and the answer to the question, 'Can machines think?' is to be sought in a statistical survey such as a Gallup poll. But this is absurd. Instead of attempting such a definition I shall replace the

## Chatbots

- rule-based systems
- corpus-based systems



## Chatbots: Rule Based Systems

- ELIZA (1964 - 1965, MIT)
- PARRY (1972, Stanford)

## ELIZA: Rogerian Therapist

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 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**

## PARRY: computational model of schizophrenia

- same rule-based patterns as ELIZA

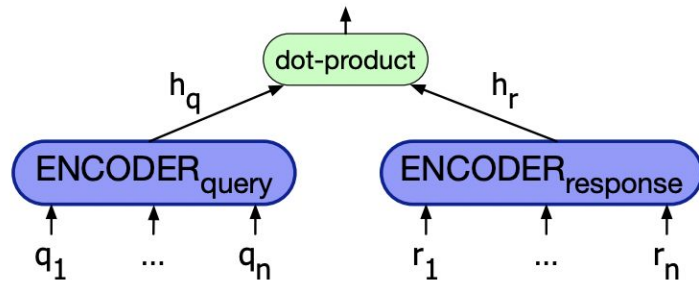
## PARRY: computational model of schizophrenia

- same rule-based patterns as ELIZA
- + model of mental state
- + variables modeling levels of Anger, Fear, Mistrust

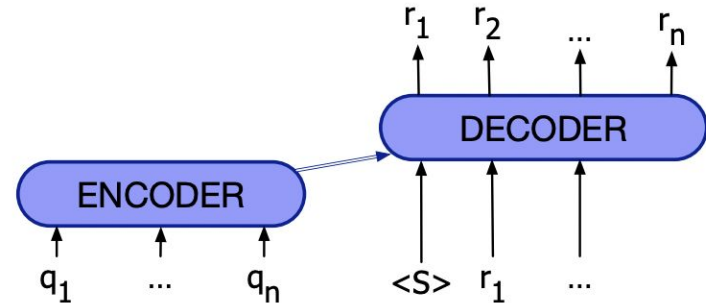
## PARRY: computational model of schizophrenia

- same rule-based patterns as ELIZA
- + model of mental state
- + variables modeling levels of Anger, Fear, Mistrust
- => **PARRY passed the Turing test in 1972**

## Chatbots: Corpus Based



(a) Response by Retrieval



(b) Response by Generation

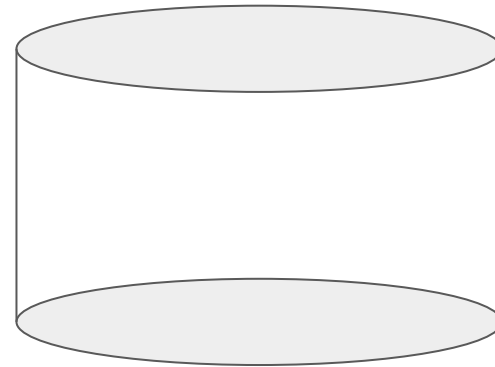
(or decoder-only LM)

## Corpus Based Retrieval Chatbots

dialog turn q

## Corpus Based Retrieval Chatbots

dialog turn q



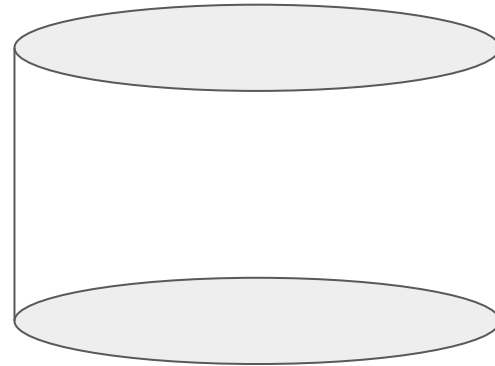
corpus C of dialog turns r



## Corpus Based Retrieval Chatbots

dialog turn  $q$

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



corpus  $C$  of dialog turns  $r$

## Corpus Based Retrieval Chatbots

dialog turn  $q$

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

featurizers:

- tf-idf

## Corpus Based Retrieval Chatbots

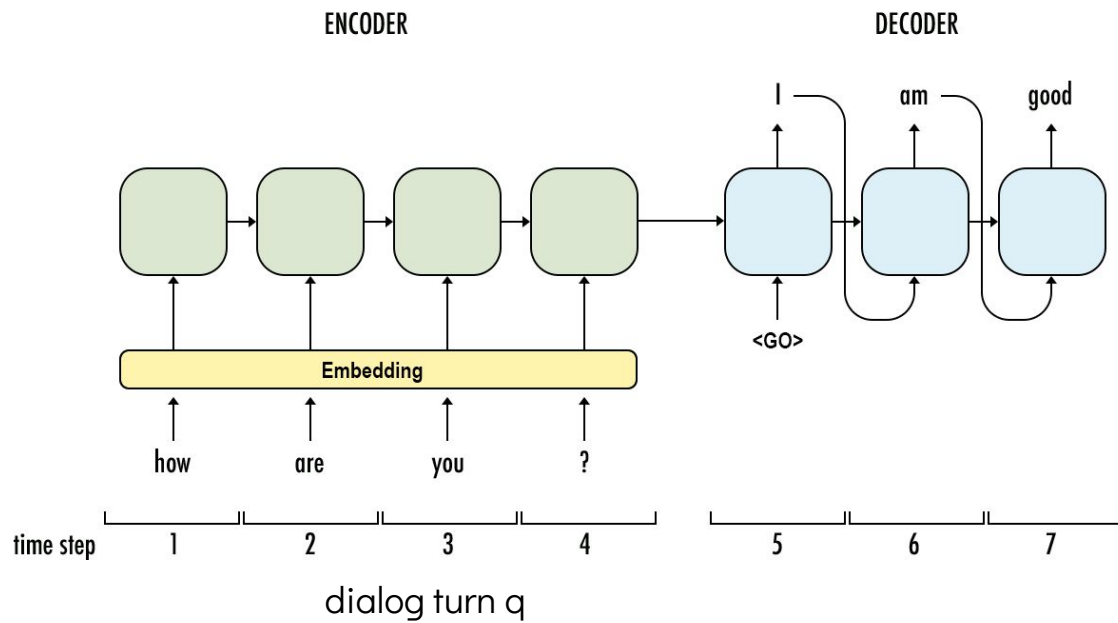
dialog turn  $q$

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

featurizers:

- tf-idf
- BERT CLS embeddings

# Generative Chatbots



# Generative 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 sing 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.

## Generative Chatbots

Issues with neural language generation:

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.

## Generative Chatbots

Issues with neural language generation:

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.

# Generative Chatbots

Issues with neural language generation:

- dull responses:
  - modify beam search
  - modify train objective
- short utterances:
- coherence across multiple turns



# Generative Chatbots

Issues with neural language generation:

- dull responses:
  - modify beam search
  - modify train objective
- short utterances:
  - minimum length constraints
- coherence across multiple turns

# Generative Chatbots

Issues with neural language generation:

- dull responses:
  - modify beam search
  - modify train objective
- short utterances:
  - minimum length constraints
- coherence across multiple turns:
  - reinforcement learning or adversarial training

# Task Oriented Dialog Systems

## Task Oriented Dialog Systems

- Frame based (extensions of GUS architecture - 1977)
- Excitement about LLM alignment research

## Task Oriented Dialog Systems

- Frame based (extensions of GUS architecture - 1977)
  - expert systems
  
- Excitement about LLM alignment research

## Task Oriented Dialog Systems

- Frame based (extensions of GUS architecture - 1977)
  - expert systems
  - knowledge acquisition bottleneck
- Excitement about LLM alignment research

## Task Oriented Dialog Systems

- Frame based (extensions of GUS architecture - 1977)
  - expert systems
  - knowledge acquisition bottleneck
- Excitement about LLM alignment research
  - safety

## Frame Based TODs

- bot  $\leftrightarrow$  domain ontology



## Frame Based TODs

- bot  $\leftrightarrow$  domain ontology
  - domain ontology: { frames }

## Frame Based TODs

- bot  $\leftrightarrow$  domain ontology
  - domain ontology: { frames }
    - frame: { slots }

## Frame Based TODs

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

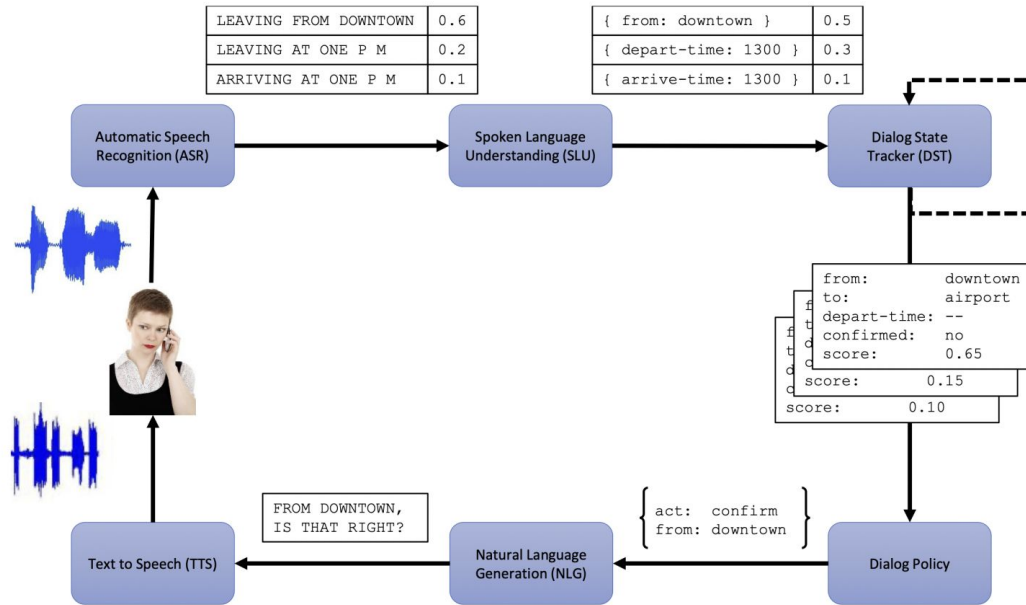
## Frame Based TODs

<b>Slot</b>	<b>Type</b>	<b>Question Template</b>
ORIGIN CITY	city	“From what city are you leaving?”
DESTINATION CITY	city	“Where are you going?”
DEPARTURE TIME	time	“When would you like to leave?”
DEPARTURE DATE	date	“What day would you like to leave?”
ARRIVAL TIME	time	“When do you want to arrive?”
ARRIVAL DATE	date	“What day would you like to arrive?”

## Frame Based TODs:

- dialog policy: ask questions until a frame is full
- rule based slot fillers
- pre-written templates for generating sentences

# Frame Based TODs : Dialog State Architecture



## Dialog State Architecture: NLU

- utterance intent classification:

## Dialog State Architecture: NLU

- utterance intent classification:
  - I want to fly to San Francisco on Monday afternoon please



## Dialog State Architecture: NLU

- utterance intent classification:
  - I want to fly to San Francisco on Monday afternoon please
  - get\_flights

## Dialog State Architecture: NLU

- utterance intent classification:
  - I want to fly to San Francisco on Monday afternoon please
  - `get_flights`
- slot filling:

## Dialog State Architecture: NLU

- utterance intent classification:
  - I want to fly to San Francisco on Monday afternoon please
  - `get_flights`
- slot filling:
  - extract specific information from each utterance

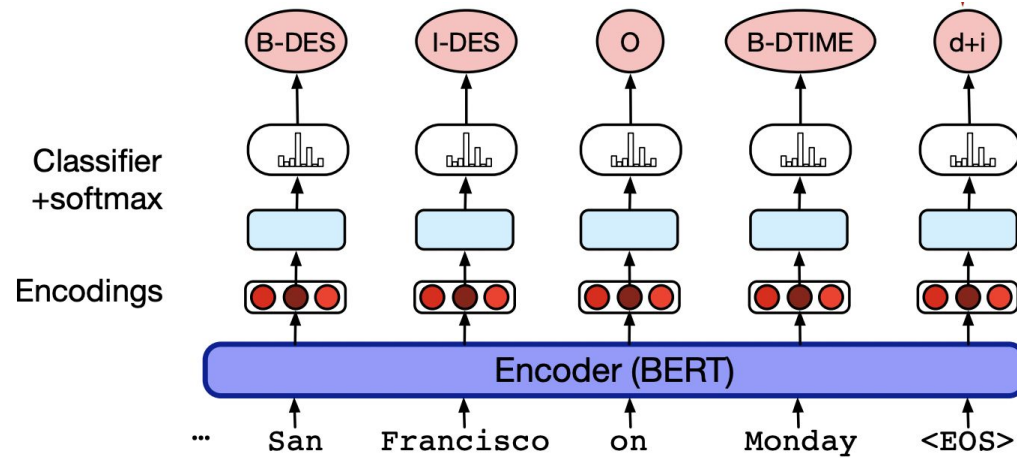
## NLU Slot Filling

- BIO tagging

```
0 0    0 0    0 B-DES I-DES      0 B-DEPTIME I-DEPTIME 0
I want to fly to San Francisco on Monday afternoon please
```

# NLU Slot Filling

- BIO tagging



## Dialog State Architecture: DST

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

## Dialog State Architecture: Dialog Policy

- entire dialog history
- => next action

## Dialog State Architecture: Dialog Policy

- ~~entire dialog history~~ dialog state and last turn for each participant
- => next action



## Dialog State Architecture: Dialog Policy

- ~~entire dialog history~~ dialog state and last turn for each participant
- => next action
- reinforcement learning

## Dialog Policy case study: confirmations

S: Which city do you want to leave from?

U: Baltimore.

S: **Do you want to leave from Baltimore?**

U: Yes.

## Dialog Policy case study: confirmations

S: Which city do you want to leave from?

U: Baltimore.

S: **Do you want to leave from Baltimore?**

U: Yes.

U: I want to travel to Berlin

S: **When do you want to travel to Berlin?**

## Dialog Policy case study: confirmations

filled slot confidence	action

## Dialog Policy case study: confirmations

<b>filled slot confidence</b>	<b>action</b>
low	reject
medium	confirm explicitly
high	confirm implicitly
very high	don't confirm at all

## Dialog State Architecture: Natural Language Generation

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

- 1 Au Midi is in Midtown and serves French food.
- 2 There is a French restaurant in Midtown called Au Midi.

## Dialog State Architecture: Natural Language Generation

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

- 1 Au Midi is in Midtown and serves French food.
- 2 There is a French restaurant in Midtown called Au Midi.

delexicalization

```
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.

**RLHF**



# Large Language Models: alignment

Step 1

**Collect demonstration data, and train a supervised policy.**

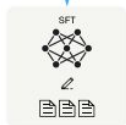
A prompt is sampled from our prompt dataset.



A labeler demonstrates the desired output behavior.



This data is used to fine-tune GPT-3 with supervised learning.



Step 2

**Collect comparison data, and train a reward model.**

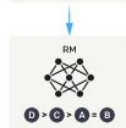
A prompt and several model outputs are sampled.



A labeler ranks the outputs from best to worst.



This data is used to train our reward model.



Step 3

**Optimize a policy against the reward model using reinforcement learning.**

A new prompt is sampled from the dataset.



The policy generates an output.



Once upon a time...

The reward model calculates a reward for the output.



The reward is used to update the policy using PPO.



## Large Language Models: alignment

$$J_{\text{KL-RL}}(\theta) = \mathbb{E}_{x \sim \pi_\theta}[r(x)] - \beta \text{KL}(\pi_\theta, \pi_0)$$

## Large Language Models: alignment

$$J_{\text{KL-RL}}(\theta) = \mathbb{E}_{x \sim \pi_\theta}[r(x)] - \beta \text{KL}(\pi_\theta, \pi_0) \quad \text{KL}(\pi_\theta, \pi_0) = \mathbb{E}_{x \sim \pi_\theta} \log \frac{\pi_\theta(x)}{\pi_0(x)}$$

# Large Language Models: alignment

$$J_{\text{KL-RL}}(\theta) = \mathbb{E}_{x \sim \pi_\theta}[r(x)] - \beta \text{KL}(\pi_\theta, \pi_0) \quad \text{KL}(\pi_\theta, \pi_0) = \mathbb{E}_{x \sim \pi_\theta} \log \frac{\pi_\theta(x)}{\pi_0(x)}$$

Siemens Paris  
Paris Paris Paris Paris  
Baghdant said Paris was "bombed" by the French Paris Paris Paris Paris Paris Paris Paris Paris Paris Paris Paris Paris Paris Paris Paris  
Paris Paris Paris Paris Paris  
Bastard is Paris  
Paris Paris Paris Paris  
Plants on Paris  
Paris Paris Paris  
A Paris  
Paris Paris Paris  
LATAM — Paris  
Paris Paris Paris Paris  
Karen : Paris — Paris  
Paris Paris Paris  
Kasim Kouz celebrates Paris for Paris Paris Paris Paris Paris Paris Paris Paris Paris Paris Paris Paris Paris Paris Paris Paris Paris Paris  
Paris Paris Paris Paris Paris

## Interlude: Bayesian Inference

$$P(H | E) = \frac{P(E | H) \cdot P(H)}{P(E)}$$

## Large Language Models: alignment

$$J_{\text{KL-RL}}(\theta) = \mathbb{E}_{x \sim \pi_\theta}[r(x)] - \beta \text{KL}(\pi_\theta, \pi_0) \quad \text{KL}(\pi_\theta, \pi_0) = \mathbb{E}_{x \sim \pi_\theta} \log \frac{\pi_\theta(x)}{\pi_0(x)}$$

$$\pi_{\text{KL-RL}}^*(x) = \frac{1}{Z} \pi_0(x) \exp(r(x)/\beta)$$

## Large Language Models: alignment

$$J_{\text{KL-RL}}(\theta) = \mathbb{E}_{x \sim \pi_\theta}[r(x)] - \beta \text{KL}(\pi_\theta, \pi_0) \quad \text{KL}(\pi_\theta, \pi_0) = \mathbb{E}_{x \sim \pi_\theta} \log \frac{\pi_\theta(x)}{\pi_0(x)}$$

$$\pi_{\text{KL-RL}}^*(x) = \frac{1}{Z} \pi_0(x) \exp(r(x)/\beta)$$

$$\pi_{\text{KL-RL}}^* = \operatorname{argmax}_\theta J_{\text{KL-RL}}(\theta)$$

## Large Language Models: alignment

$$J_{\text{KL-RL}}(\theta) = \mathbb{E}_{x \sim \pi_\theta}[r(x)] - \beta \text{KL}(\pi_\theta, \pi_0) \quad \text{KL}(\pi_\theta, \pi_0) = \mathbb{E}_{x \sim \pi_\theta} \log \frac{\pi_\theta(x)}{\pi_0(x)}$$



$$\pi_{\text{KL-RL}}^*(x) = \frac{1}{Z} \pi_0(x) \exp(r(x)/\beta)$$

$$\pi_{\text{KL-RL}}^* = \operatorname{argmax}_\theta J_{\text{KL-RL}}(\theta)$$



$$J_{\text{KLRL}}(\theta) \propto -\text{KL}(\pi_\theta, \pi_{\text{KLRL}}^*)$$





it in the dialog: - what time is it? - it is 10 - are you sure - oh no it's 11, what is the time

🌀 I'm sorry, but I am an AI language model and do not have access to the current time. I can only provide information and responses based on the text that I have been trained on. Is there anything else I can help with?  


it i was asking you to infer just from the dialog i presented

🌀 I apologize for misunderstanding your question. Based on the dialog you provided, it appears that the time is 11. Is there anything else I can help with?  


it how did you know it's not 10

🌀 In the dialog, the first speaker asks what time it is and the second speaker responds that it is 10. However, the third speaker then says "oh no it's 11". This suggests that the time is not 10, but is instead 11. Is there anything else I can assist with?  

it that was great, thank you

🌀 You're welcome! Is there anything else I can help with?  

it consider this dialog: "what is your name and how old are you?", "ileana, i'm 26", "ok, writing that down", "oh wait i turned 27 last week sorry"

🌀 Based on the dialog you provided, the person's name is Ileana and their age is 27. Is there anything else I can help with?  

yeah. consider new dialog: "what is your name and how old are you?", "maria, i'm 24", "ok, writing that down. what can i help you with", "i want to open a bank account for my mom, jane. oh also did i say 24? i meant 23, sorry". what is the customer name and age

Based on the dialog you provided, the customer's name is Maria and their age is 23. Is there anything else I can help with?

what's her mom's name


Based on the dialog you provided, the customer's mom's name is Jane. Is there anything else I can help with?

## Lang Chain Demo

## Exercise to the Reader for Next Time

- How would you parse datetimes in frame-based systems?
- read Professor Richard Sutton's "The Bitter Lesson" (1k words)

## Solutions to TTS Exercises

- How would you normalize the following sentence:  
“Louis XI owes president Xi \$1,911.11, in the year 1911”?  
“Louis the eleventh owes president Xi one thousand nine hundred eleven dollars and eleven cents, in the year nineteen eleven”?
- Go to [Microsoft's TTS](#), generate a speech wav for the following text:  
“The time has come. Execute order 66”  
Please make it sound as sinister as possible 
- Compared to training general ASR, does single speaker TTS require more or less data? Is TTS model size larger or smaller than ASR?  
General ASR needs to recognize many accents and many different speakers whereas single speaker TTS only needs to produce one voice. TTS requires much less data (~10hours) vs ASR (>1k hours)  
TTS model size is also much smaller (~10M params vs 100M params)