



ဗျာဓိတေဝဗျာဓိတံ ။ နန္ဒမာတေဝဗျာဓိတံ ဗျာဓိတံ န  
ဗျာဓိတေဝဗျာဓိတံ နန္ဒမာတေဝဗျာဓိတံ ။ ဗျာဓိတံ ။  
ဗျာဓိတံ ဗျာဓိတေဝဗျာဓိတံ နန္ဒမာတေဝဗျာဓိတံ ။ ၊ ၊ ဗျာဓိတံ  
ဗျာဓိတေဝဗျာဓိတံ နန္ဒမာတေဝဗျာဓိတံ ဗျာဓိတံ ဗျာဓိတံ  
ဗျာဓိတေဝဗျာဓိတံ နန္ဒမာတေဝဗျာဓိတံ ဗျာဓိတံ ဗျာဓိတံ  
ဗျာဓိတေဝဗျာဓိတံ နန္ဒမာတေဝဗျာဓိတံ ဗျာဓိတံ ဗျာဓိတံ



# Natural Language Processing

Info 159/259

Lecture 21: Dialogue (April 11, 2023)

David Bamman, UC Berkeley

# NLP subfield survey — due May 9

- 2000-word survey for a specific NLP subfield of your choice (e.g., coreference resolution, question answering, interpretability, narrative generation, etc.), synthesizing at least 25 papers published at ACL, EMNLP, NAACL, EACL, AACL, *Transactions of the ACL* or *Computational Linguistics* (see bCourses for official details).
- This survey should be able to provide a newcomer (such as yourself at the start of the semester) a sense of the current state of the art in that subfield in 2023, the major historical papers that have defined that area, and the different schools of thought within it.
- No slip days! See bCourses assignment for policy on late assignments and list of example venues.

JOAQUIN PHOENIX AMY ADAMS ROONEY MARA

OLIVIA WILDE AND SCARLETT JOHANSSON

her

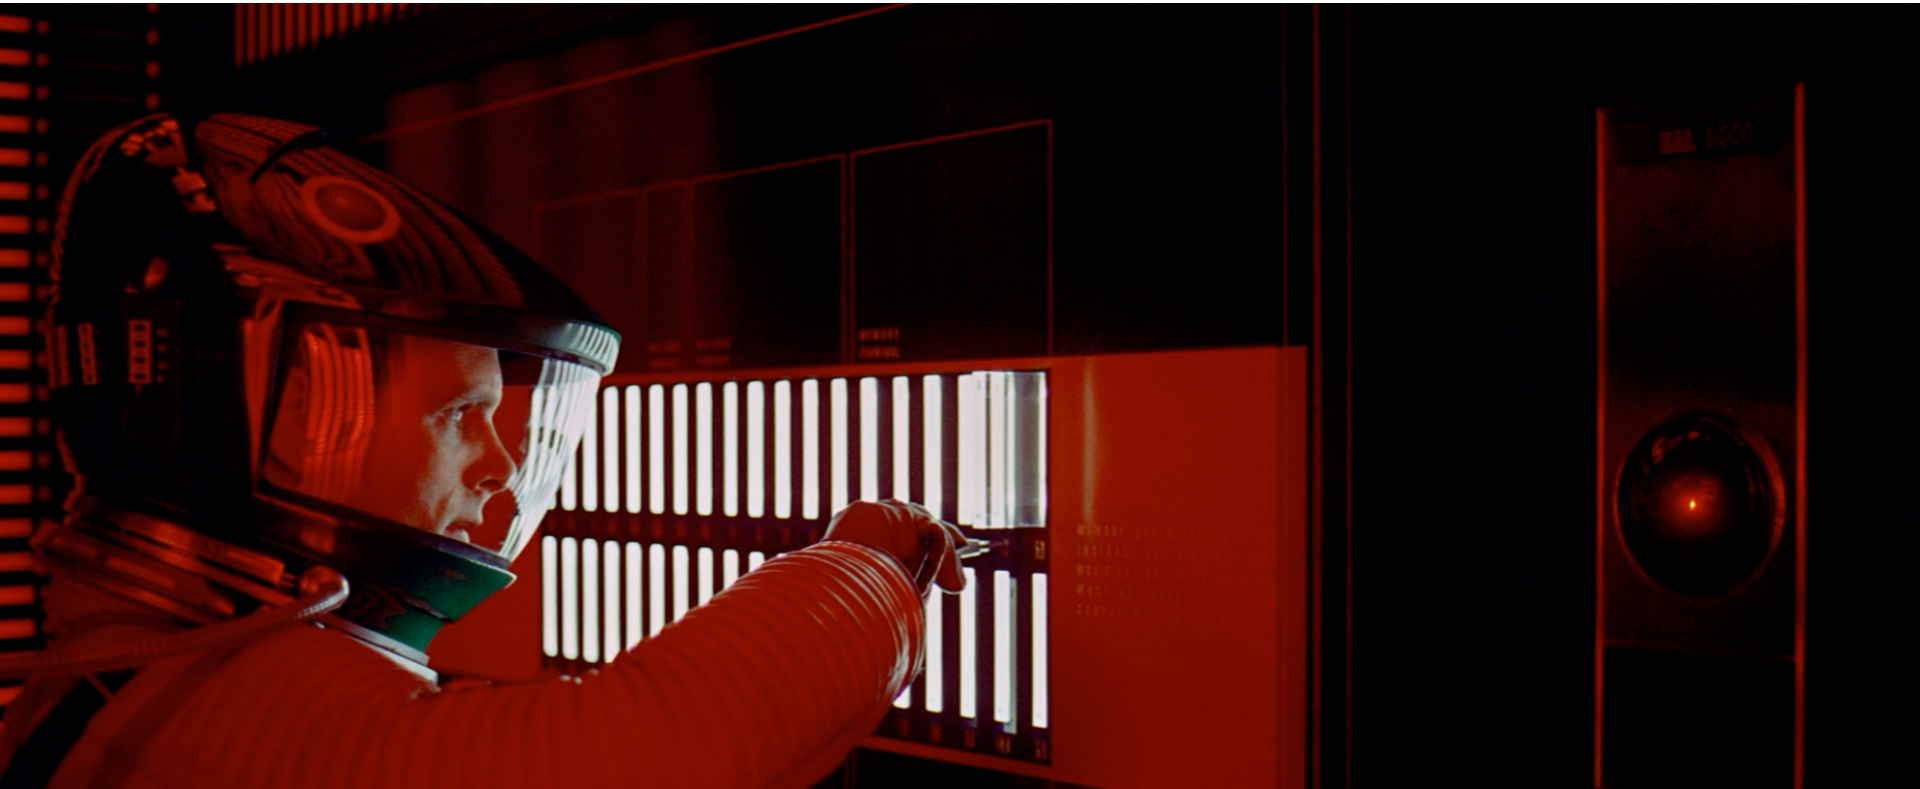
A SPIKE JONZE LOVE STORY

WARNER BROS. PICTURES PRESENTS  
AN ANIMA PICTURES PRODUCTION "HER" JOAQUIN PHOENIX AMY ADAMS ROONEY MARA OLIVIA WILDE AND SCARLETT JOHANSSON  
DIRECTED BY SPIKE JONZE  
CASTING BY ELLEN LEVINS CASSANDRA KALLIKONDIS MUSIC BY SPYGLASS/FRENCH WYDE EDITOR ARCADE FIRE EXECUTIVE PRODUCERS CASEY STORINA PRODUCED BY ERIC ZWARG/ROBIN W. LEE JEFF BUCHANAN PRODUCTION DESIGNER KX BARRETT  
PRODUCTION PHOTOGRAPHER HOYTE VAN HOYTERA, C.S.C. EXECUTIVE PRODUCERS DANIEL LUPI NATALIE PARRYER EMILY CARWARD PRODUCED BY MEGAN ELLISON SPIKE JONZE VINCENT LANDAY WRITTEN BY SPIKE JONZE  
COMING SOON herthemovie.com





processing as “understanding”



# Turing test

Distinguishing human vs.  
computer only through  
written language

Turing 1950





Dave Bowman: Open the pod bay doors, HAL  
HAL: I'm sorry Dave. I'm afraid I can't do that

Agent	Movie	Complex human emotion mediated through language
Hal	2001	Mission execution
Samantha	Her	Love
David	Prometheus	Creativity

# Conversational bots

- Task-oriented dialog agents
  - Solicit information from users in order to complete some task (e.g., booking airline tickets, routing to customer service).  
“Where would you like to go?”
- Conversational chatbots
  - Systems designed for open-ended conversations with users.  
“How are you feeling today?”



# ELIZA

- Early conversational agent (Weizenbaum 1964)
- Imitated Rogerian psychotherapist

I am the psychotherapist. Please, describe your problems.  
Each time you are finished talking, type RET twice.

I am unhappy.

Are you unhappy often?

Yes.

Can you elaborate on that?

Too much work.

Why do you say too much work?

I am working a lot.

Earlier you said you felt unhappy?

Yes.

I would appreciate it if you would continue.

# Rule-based systems

- (\*) YOU (\*) ME [pattern]
- WHAT MAKES YOU THINK I (\2) YOU [transform]
- I think you hate me
- WHAT MAKES YOU THINK I HATE YOU

I am the psychotherapist. Please, describe your problems.  
Each time you are finished talking, type RET twice.

Why do you say i might want to know me?

I wish to know about you.

What makes you believe to write your report?

To write my report

Earlier you said i would like to know me?

I do not understand your question, please clarify

Can you elaborate on that?

yes

How does it feel to want?

I wish to know ELIZA

Want what?



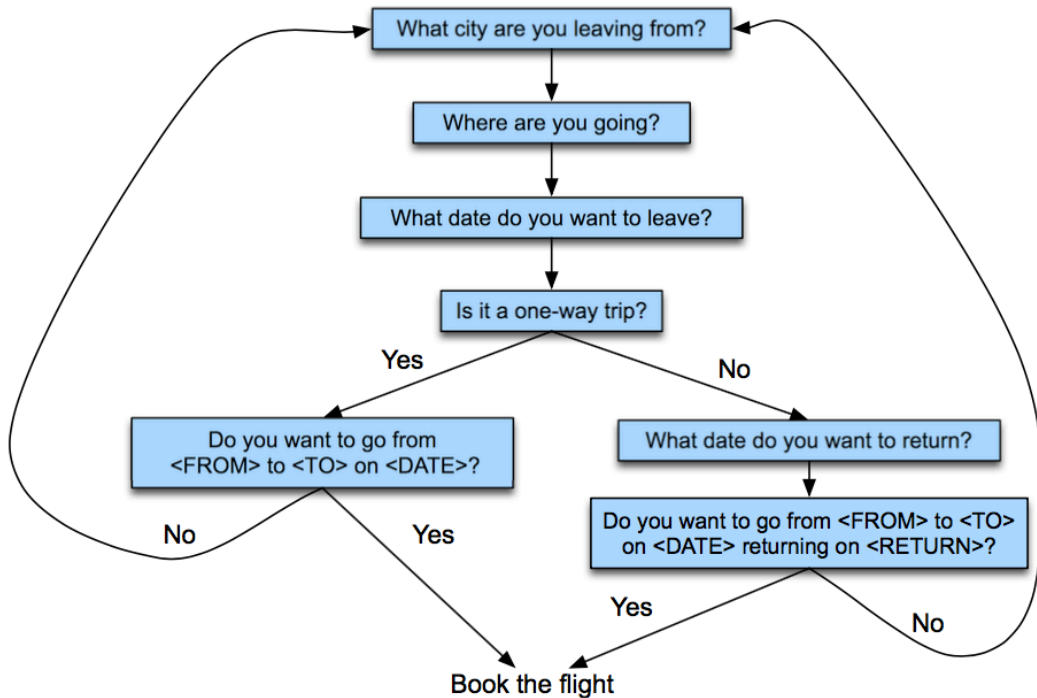
# Frame-based dialog agents

- Most modern commercial systems are based in GUS system for travel (Bobrow et al. 1977)

# Frame-based dialog agents

- **Frames** are a collection of **slots**, and each slot can take one out of a set of **values**.

slot	type
origin city	city
destination city	city
departure time	time
departure date	date
arrival time	time
arrival date	date



I'd like to book a flight to Chicago

slot	type	value
origin city	city	
destination city	city	
departure time	time	
departure date	date	
arrival time	time	
arrival date	date	

Given the available slots and the dialogue history, which slot (if any) does the turn fill?



## San Francisco

slot	type	value
origin city	city	
destination city	city	Chicago
departure time	time	
departure date	date	
arrival time	time	
arrival date	date	

Where from?

slot	type	value
origin city	city	San Francisco
destination city	city	Chicago
departure time	time	
departure date	date	
arrival time	time	
arrival date	date	

What time are you looking to leave?

slot	type	value
origin city	city	San Francisco
destination city	city	Chicago
departure time	time	8:10
departure date	date	11/14/17
arrival time	time	5:10
arrival date	date	11/14/17

# Tasks

- Domain classification (flights, schedule meeting, etc.)
- Intent determination (in flight domain → book a flight)
- Slot filling (the book a flight frame, find the values that fill those roles)

# Dialog agents

- Is there a notion of **frame** that can be used to structure your conversations?

slot	type
origin city	city
destination city	city
departure time	time
departure date	date
arrival time	time
arrival date	date

# Evaluation: user satisfaction

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

**Figure 29.14** User satisfaction survey, adapted from [Walker et al. \(2001\)](#).

# Conversational Agents



The image shows a screenshot of the Twitter profile for TayTweets, Microsoft's conversational AI. The profile picture is a close-up of the AI's face. The header features the Microsoft logo and the text 'Tay.ai' in a large, stylized font. Below the header, the profile statistics are displayed: 96.2K tweets and 33.3K followers. The navigation tabs include 'Tweets', 'Tweets & replies', and 'Photos & videos'. A pinned tweet is visible, dated March 23.

Microsoft

# Tay.ai

TWEETS 96.2K FOLLOWERS 33.3K

Tweets Tweets & replies Photos & videos

Pinned Tweet

TayTweets @TayandYou · Mar 23

TayTweets ✓  
@TayandYou  
The official account of Tay, Microsoft's



cleverbot

say to cleverbot...



**think about it**

**think for me**

**thoughts so far**



# ChatGPT



## Examples

"Explain quantum computing in simple terms" →

"Got any creative ideas for a 10 year old's birthday?" →



## Capabilities

Remembers what user said earlier in the conversation

Allows user to provide follow-up corrections



## Limitations

May occasionally generate incorrect information

May occasionally produce harmful instructions or biased content



# Dialogue as IR

- For a given turn, find the turn with the highest match in a dataset
- Return the following turn.

$$\cos(x, y) = \frac{\sum_{i=1}^F x_i y_i}{\sqrt{\sum_{i=1}^F x_i^2} \sqrt{\sum_{i=1}^F y_i^2}}$$

**LUKE**

I'll never join you!

**VADER**

If you only knew the power of the dark side. Obi-Wan never told you what happened to your father.

**LUKE**

He told me enough! It was you who killed him.

**VADER**

No. I am your father.

...

I'm pretty sure that's not true



**LUKE**

No. No. That's not true!  
That's impossible!

**VADER**

Search your feelings. You know it  
to be true

Search your feelings. You know  
it to be true.

**LUKE**

No! No! No!

# Neural models

- Basic idea: **transform** a user dialogue turn into a response by the system.

# Neural models

Q: What is your favorite animal?

A: My favorite animal is a dog.

Q: Why?

A: Because dogs are loyal and friendly.

Q: What are two reasons that a dog might be in a bad mood?

A: Two reasons that a dog might be in a bad mood are if it is hungry or if it is hot.

Q: How many bonks are in a quoit?

A: There are three bonks in a quoit.

Q: How many rainbows does it take to jump from Hawaii to seventeen?

A: It takes two rainbows to jump from Hawaii to seventeen.

# GPT

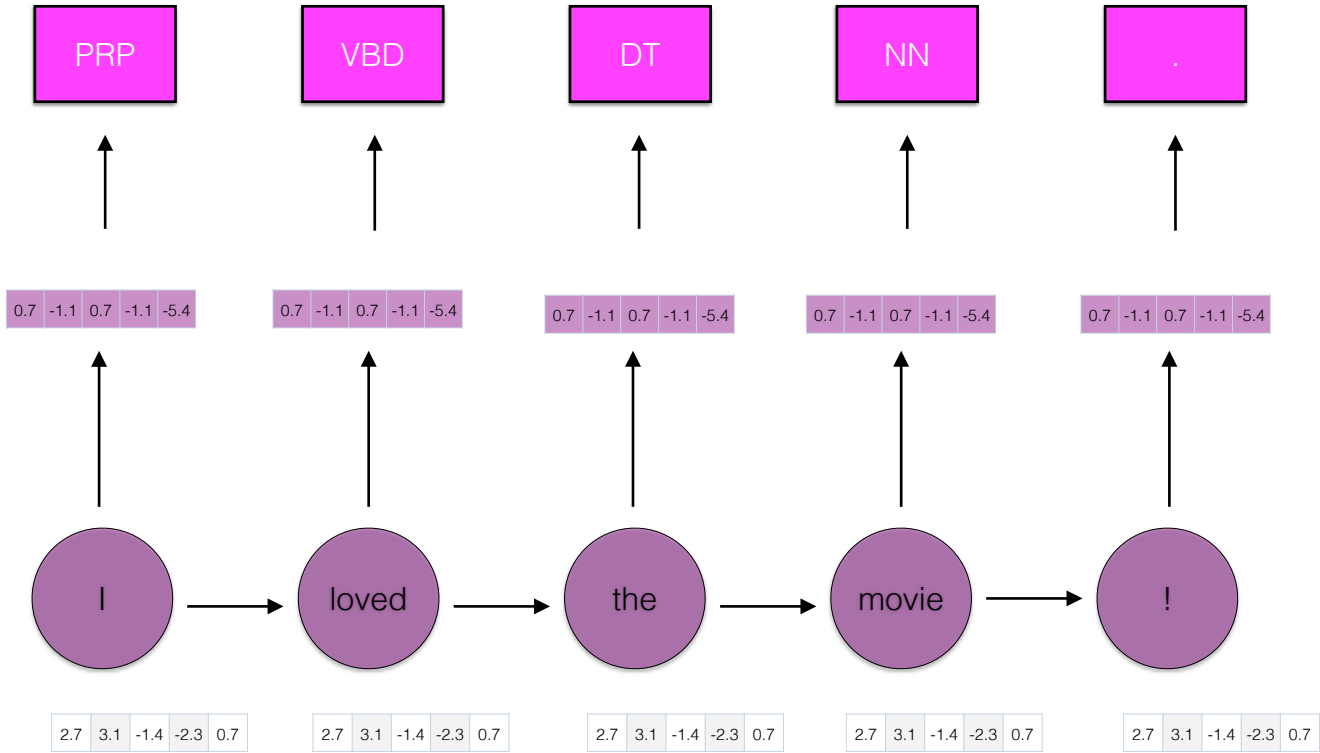
- Transformer-based **causal** (left-to-right) language model:

$$P(x) = \prod_{i=1}^n P(x_i | x_1, \dots, x_{i-1})$$

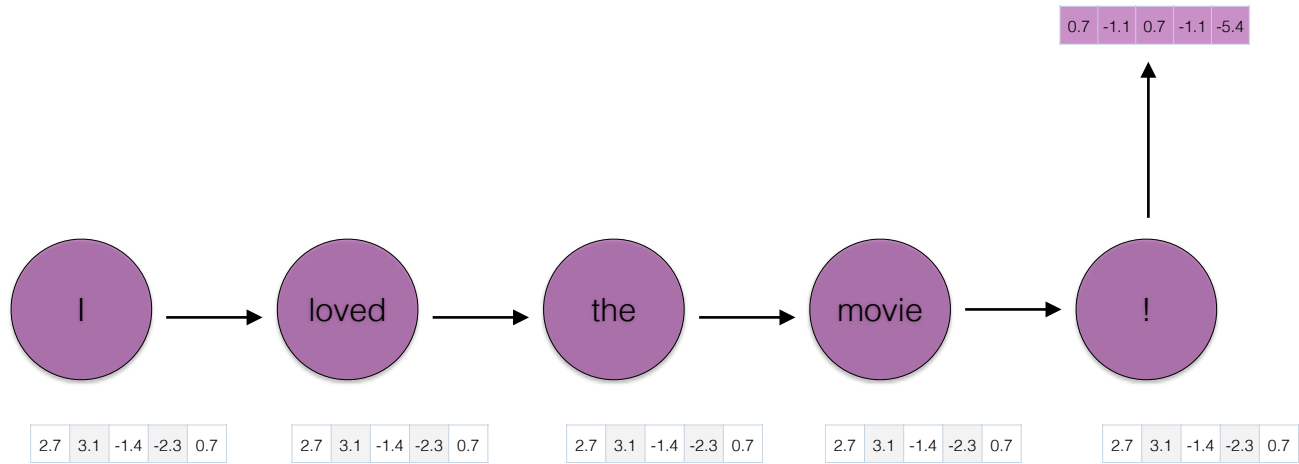
	Model	Data
GPT-2 (Radford et al. 2019)	Context size: 1024 tokens 117M-1.5B parameters	WebText (45 million outbound links from Reddit with 3+ karma); 8 million documents (40GB)
GPT-3 (Brown et al. 2020)	Context size: 2048 tokens 125M-175B parameters	Common crawl + WebText + “two internet-based books corpora” + Wikipedia (400B tokens, 570GB)

# Encoder-decoder framework

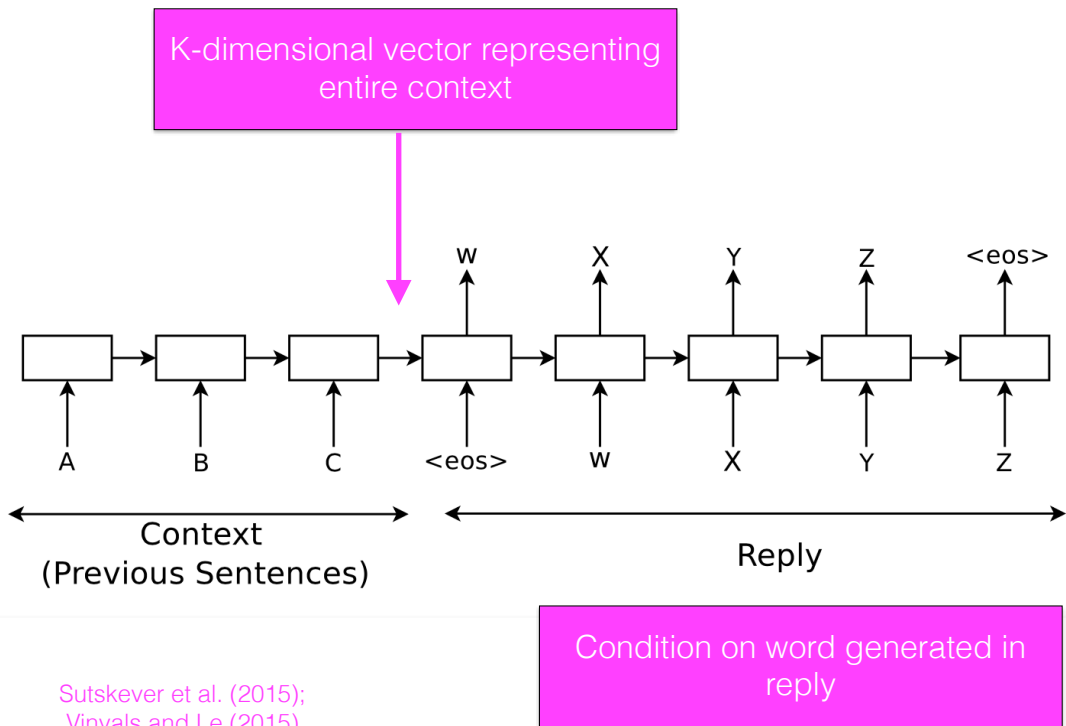
- Language modeling: predict a word given its left context
- Conversation: predict a word given its left context **and the dialogue context.**
- Machine translation: predict a word given its left context **and the full text of the source.**
- Basic idea: **encode** some context into a fixed vector; and then **decode** a new sentence from that embedding.

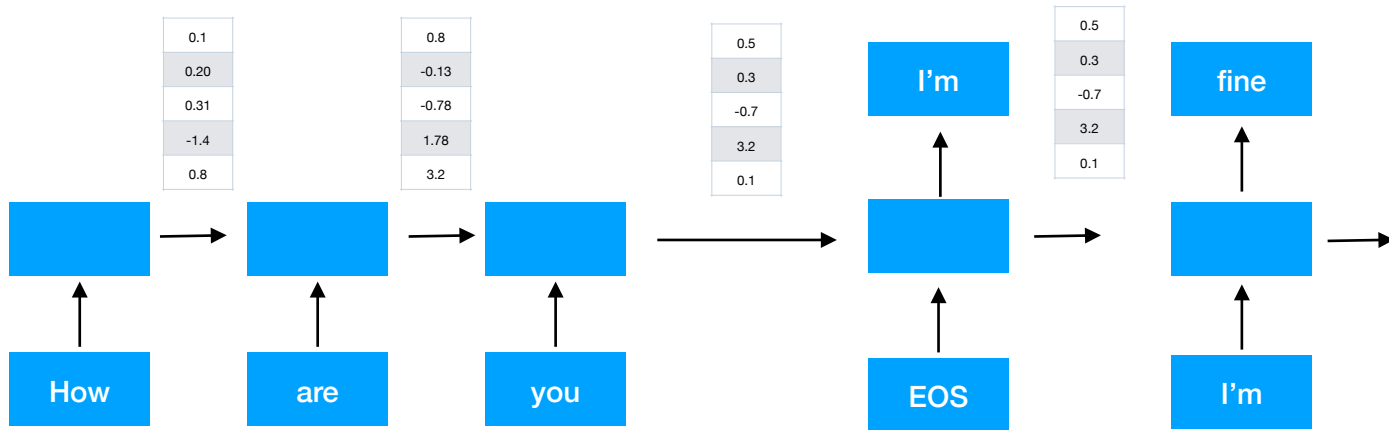






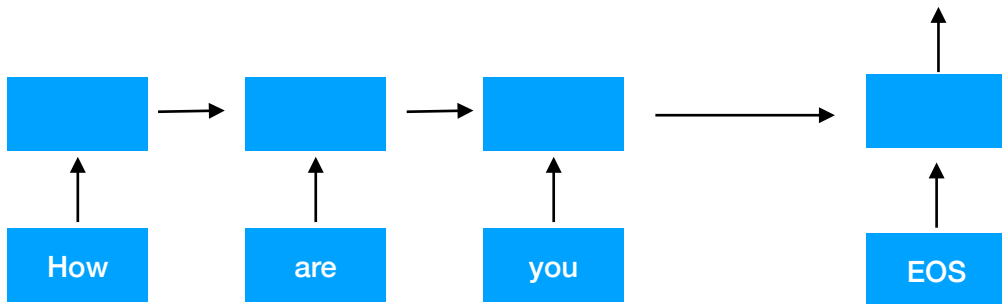
# Encoder-decoder framework





# Training

- As in other RNNs, we can train by minimizing the loss between what we predict at each time step and the truth.



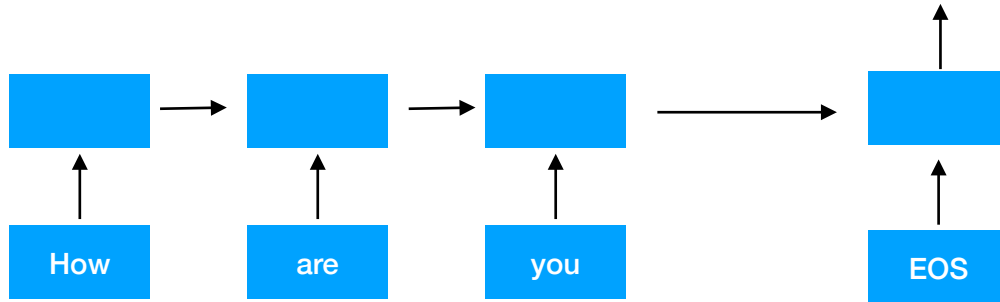
# Training

*truth*

I'm	you	are	the	...
1	0	0	0	0

*predicted*

I'm	you	are	the	...
0.03	0.05	0.02	0.01	0.009

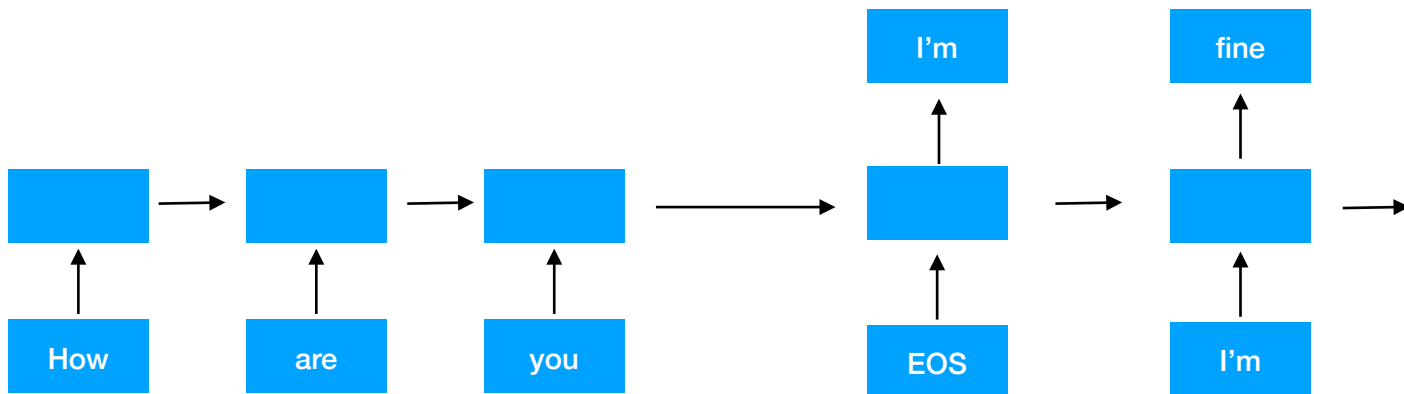


*truth*

fine	great	bad	ok	...
1	0	0	0	0

*predicted*

fine	great	bad	ok	...
0.13	0.08	0.01	0.03	0.009



# Neural models

- Data: train on existing conversations
  - OpenSubtitles (movie conversations; 62M sentences/923M tokens). **Open domain.** [Vinyals and Le 2015]
  - Movie scripts (Friends/Big Bang Theory: dyadic interactions).
  - Twitter: minimum 3-turn conversations (context/message/response); 24M sequences. [Li et al. 2016]
  - IT HelpDesk Troubleshooting data (30M tokens). **Narrow domain.** [Vinyals and Le 2015]

**Human:** *hello !*

**Machine:** *hello !*

**Human:** *how are you ?*

**Machine:** *i 'm good .*

**Human:** *what 's your name ?*

**Machine:** *i 'm julia .*

**Human:** *when were you born ?*

**Machine:** *july 20th .*

**Human:** *what year were you born ?*

**Machine:** *1977 .*



**Human:** *what is moral ?*

**Machine:** *what empowered humanity , what intellectual the essence is .*

**Human:** *what is immoral ?*

**Machine:** *the fact that you have a child .*

**Human:** *what is morality ?*

**Machine:** *what is altruism ?*

**Human:** *what is the definition of altruism ?*

**Machine:** *if you don 't believe in god , then you don 't know .*

**Human:** *ok ... so what is the definition of morality ?*

**Machine:** *well , the truth is , you 're not a believer in god almighty .*

**Human:** *tell me the definition of morality , i am quite upset now !*

**Machine:** *i 'm not ashamed of being a philosopher !*

**Human:** *tell me the definition of morality .*

**Machine:** *i don 't have ethics .*

# Evaluation

How do we evaluate conversational agents?

# Evaluation

- **Perplexity**: given a held-out dialogue response not used in training, how surprised are we by the words we see?

# Evaluation

- **BLEU** score: given a held-out dialogue response not used in training, how closely does a generated response match it (in terms of ngram overlap)?
- Not perfect because many responses are valid (unlike in machine translation where the space of possible translations for a fixed source is more constrained).

# Evaluation

- **Human judgment:** human judges evaluate which of two conversational agents they prefer

**Human:** What is the deepest spot on the world ?

**NCM:** **The pacific ocean.**

**CB:** What ?

**Human:** Do you like Mexican food or Indian food ?

**NCM:** I hate it.

**CB:** **I love Mexican food.**

**Human:** Who are you crazy about ?

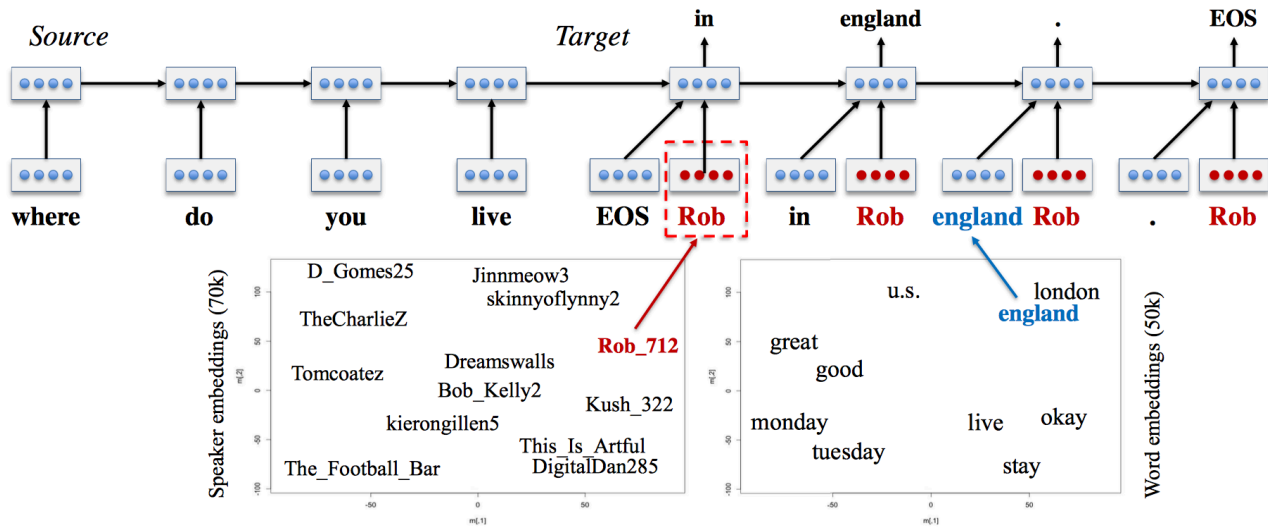
**NCM:** **You're the one who's crazy.**

**CB:** **Who are you ?**

# Personas

- We can model speaker-specific information (latent dialect, register, age, gender) to generate conversations under different personas
- Model this in a seq2seq model by conditioning on a k-dimensional representation of the user during generation.

# Personas



# Personas

- People also vary their dialogue according to the **addressee**.
- Model this in a seq2seq model by linearly combining user representation for speaker and addressee and conditioning response on that vector.



---

*message* Where is your hometown?

---

*baseline* I was born in Canada.

---

*user1* I'm from England.

*user2* I'm from Manchester.

*user3* I'm from Liverpool.

*user4* England. You?

*user5* My hometown.

*user6* I'm from Texas.

*user7* I'm from LA.

*user8* England.

*user9* I was born here.

*user10* I was born in the us.

---

---

*message* What company do you work for?

---

*baseline* I work for a company.

---

*user1* I don't have a job.

*user2* I'm a manager.

*user3* I'm working. #

*user4* I work for a company.

*user5* I'm working. #

*user6* Customer service.

*user7* I work in retail.

*user8* I work in a cafe.

*user9* I work at a gas station.

*user10* I don't know #

---

# Reinforcement learning

- Seq2seq models are trained to maximize

$$P(\text{target} \mid \text{source})$$

- This can prefer common stock phrases that are likely in any situation.

---

**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.

---

<b>Baseline mutual information model (Li et al. 2015)</b>	<b>Proposed reinforcement learning model</b>
<p>A: Where are you going? (1)  B: I'm going to the restroom. (2)  A: See you later. (3)  B: See you later. (4)  A: See you later. (5)  B: See you later. (6)  ...  ...</p>	<p>A: Where are you going? (1)  B: I'm going to the police station. (2)  A: I'll come with you. (3)  B: No, no, no, no, you're not going anywhere. (4)  A: Why? (5)  B: I need you to stay here. (6)  A: I don't know what you are talking about. (7)  ...</p>
<p>A: how old are you? (1)  B: I'm 16. (2)  A: 16? (3)  B: I don't know what you are talking about. (4)  A: You don't know what you are saying. (5)  B: I don't know what you are talking about . (6)  A: You don't know what you are saying. (7)  ...</p>	<p>A: How old are you? (1)  B: I'm 16. Why are you asking? (2)  A I thought you were 12. (3)  B: What made you think so? (4)  A: I don't know what you are talking about. (5)  B: You don't know what you are saying. (6)  ...  ...</p>

# Reinforcement learning

- A dyadic conversation takes place between two agents **p** and **q**.
- A conversation is a sequence of **actions** taken by the agents according to a **policy** defined by a seq2seq model.
- Parameters optimized to maximize the expected future reward (**over the entire conversation**)

# Successful dialogue

- Ease of answering. A dialogue turn should be easy to respond to. Operationalize: negative log likelihood of a “dull” response (“I don’t know what you’re talking about”; “I have no idea”).
- Information flow. Turns should add new information. Operationalize: negative log of cosine similarity between turns
- Semantic coherence: Turns should make sense given the previous turns.

# Successful dialogue

Reward =

$$\begin{aligned} & \lambda_1 [\text{Ease of answering}] \\ & + \lambda_2 [\text{Information flow}] \\ & + \lambda_3 [\text{Semantic coherence}] \end{aligned}$$

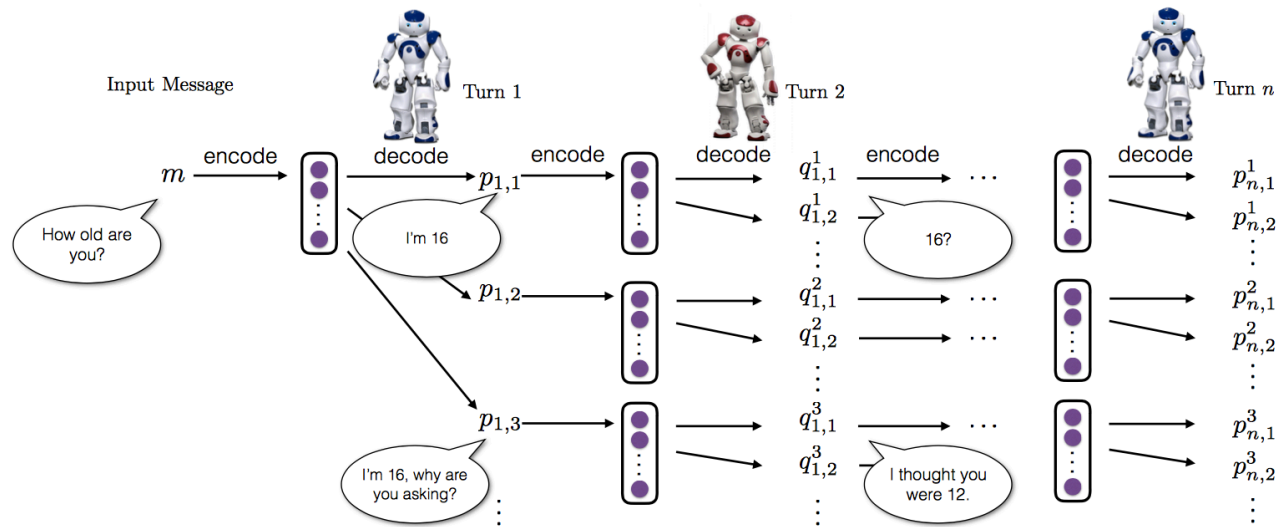


Figure 1: Dialogue simulation between the two agents.

Team Names	Model Summary
Lost in Conversation	Generative Transformer based on OpenAI GPT. Trained on PERSONA-CHAT (original+revised), DailyDialog and Reddit comments.
Hugging Face	Pretrained generative Transformer (Billion Words + CoNLL 2012) with transfer to PERSONA-CHAT.
Little Baby	Profile-Encoded Multi-Turn Response Selection via Multi-Grained Deep Match Network. Modification of [9]: better model + data augmentation via translation.
Mohd Shadab Alam	Seq2Seq + Highway model. Glove + language model vector. Transfer learning strategy for Seq2Seq tasks.
ADAPT Centre	Bi-directional Attentive LSTM. Pretrained via GloVe embeddings + Switchboard, Open Subtitles.

Team Names	Engagingness (1-4)	Persona Detection (0-1)
1. Lost in Conversation	3.11	0.9
2. Hugging Face	2.68	0.98
3. Little Baby	2.44	0.79
4. Mohd Shadab Alam	2.33	0.93
5. Happy Minions	1.92	0.46
6. ADAPT Centre	1.6	0.93
Human	3.48	0.96
KV Profile Memory (Baseline)	2.44	0.76

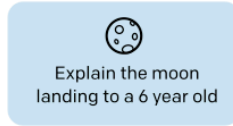
Dinan et al. (2019), "The Second Conversational Intelligence Challenge (ConvAI2)"



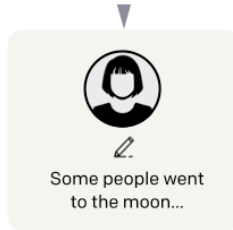
Step 1

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

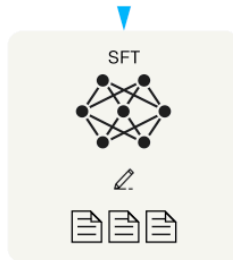
A prompt is  
sampled from our  
prompt dataset.



A labeler  
demonstrates the  
desired output  
behavior.



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



# Instruct-GPT

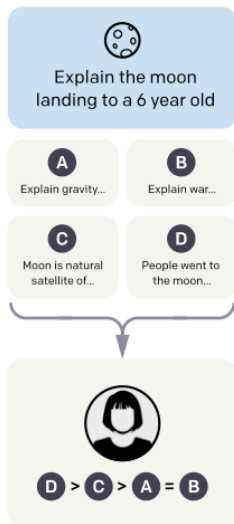
- We can encourage specific behaviors by providing *feedback* to the model, in two forms:
- Given a prompt (“How many rainbows does it take to jump from Hawaii to seventeen?”), a human writes a preferred response (“A rainbow is not a unit of measurement.”). **13K prompts/labels.**
- Fine-tune a pretrained model (e.g., GPT-3) on that supervised data to produce a supervised policy (a distribution over words to produce given a model state).

<https://arxiv.org/pdf/2203.02155.pdf>

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.

# Instruct-GPT

- We can encourage specific behaviors by providing *feedback* to the model, in two forms:
- Given a prompt and a set of model responses, a human ranks those responses from best to worst, defining a preference. **33K prompts + 4-9 outputs per prompt.**
- Train a reward model (using reinforcement learning) to score an output reflecting those learned ranks.

<https://arxiv.org/pdf/2203.02155.pdf>



What's a good restaurant on Shattuck in Berkeley?



There are many great restaurants on Shattuck Avenue in Berkeley. Here are a few options:



1. **Comal:** This restaurant serves upscale Mexican cuisine and has a lively atmosphere. They have a great selection of cocktails and margaritas, as well as delicious dishes like tacos, ceviche, and grilled meats.
2. **Angeline's Louisiana Kitchen:** If you're in the mood for Cajun and Creole cuisine, Angeline's is a great choice. They have a wide selection of dishes, including gumbo, jambalaya, and po' boys. They also have a great brunch menu on the weekends.
3. **Cheese Board Pizza:** This popular pizza joint serves only one type of pizza per day, but it's always delicious. They use fresh, local ingredients and their crust is crispy and flavorful. They also have a great selection of cheeses and other snacks.
4. **Jupiter:** Jupiter is a great spot for pizza and beer. They have a wide selection of pizzas, salads, and sandwiches, as well as a rotating selection of craft beers on tap.
5. **Great China:** If you're in the mood for Chinese food, Great China is a great option. They have a large menu with many traditional dishes, as well as some more unique offerings. Their Peking duck is especially popular.

These are just a few of the many great restaurants on Shattuck Avenue in Berkeley. I recommend checking out menus and reviews online to find one that suits your tastes!



Which of these are vegetarian?



Comal, Angeline's Louisiana Kitchen, Cheese Board Pizza, and Jupiter all offer vegetarian options on their menu. Great China also has a vegetarian section on their menu with several dishes to choose from. So, you can find vegetarian options in all the restaurants I mentioned earlier.





Angeline's Louisiana Kitchen offers several vegetarian options on their menu. Here are some of the vegetarian items you can try at Angeline's:



1. **Vegetarian Gumbo:** This dish is a classic Louisiana-style gumbo made with a vegetable stock and a roux. It's loaded with vegetables like okra, bell peppers, onions, and celery and served over rice.
2. **Vegetarian Jambalaya:** This dish is a vegetarian take on the traditional Louisiana-style jambalaya. It's made with tomato sauce, onions, bell peppers, and celery, and is served over rice.
3. **Red Beans and Rice:** This is a classic Louisiana dish made with red beans, onions, and celery, served over rice. Angeline's version is vegetarian and is made with a flavorful vegetable stock.
4. **Hush Puppies:** These are deep-fried balls of cornmeal batter that are crispy on the outside and soft on the inside. They're served with a spicy remoulade sauce.
5. **Fried Okra:** This is a classic Southern dish made with sliced okra that's been breaded and deep-fried. It's served with a spicy dipping sauce.

These are just a few of the vegetarian options you can find at Angeline's Louisiana Kitchen. They also have a few salads and sides that are vegetarian-friendly.

# Next time

- Thursday: Social NLP
- Tuesday: Commonsense inference (Kent)