

CS 6120/CS4120: Natural Language Processing

Instructor: Prof. Lu Wang

College of Computer and Information Science

Northeastern University

Webpage: www.ccs.neu.edu/home/luwang

Plan for the rest of the semester

- 3/27, 3/30: lectures on dialogue systems, chatbots, and machine translation
- 4/3: project feedback, 3:25pm-6:15pm in 258 WVH.
 - Prioritize teams that present on 4/6, who can claim slots by the end of this Thursday (3/29)
 - You can still come in during the time even you don't make appointments. But will need to wait if there are other teams.
- 4/6, 4/10, 4/13: project presentation
- 4/17: no class, final report due
- 4/25: final exam, 8am-10am, 108 WVH.

Project presentation date and order

- On Piazza (@302)
- If you'd like to shift, you'll need to find a team that is willing to do so.
- Feedback will be sent to the team through blackboard after the presentation.
- Final reports are expected to resolve the issues raised in the feedback.

Project presentation

- For each batch (teams presenting on the same date), one favorite project will be selected by voting.
- Each team has two votes.
- The team that gets the most votes wins. Each team member will get 1% bonus towards final grade.

Presentation and final report

- Problem Description (10%)
What is the task?
System input and output
Examples will be helpful
- Reference/Related work (20%)
Put your work in context: what has been done before? You need to have reference!
What's new in your work?
- Methodology: What you have done (30%)
Preprocessing of the data
What are your data? Features used? What are effective, and what are not?
What methods do you experiment with? And why do you think they're reasonable and suitable for the task?
- Experiments (40%)
Datasets size, train/test/development
Evaluation metrics: what are used and are they proper to calibrate system performance?
Baselines: what are they?
Results, tables, figures, etc

What is Natural Language Dialogue?

- Communication involving
 - Multiple contributions
 - Coherent interaction
 - More than one participant
- Interaction modalities
 - Input: Speech, typing, writing, gesture
 - Output: Speech, text, graphical display, animated face/body (embodied virtual agent)

[Some slides are borrowed from Svetlana Stoyanchev and Dan Jurasky]

What is involved in NL dialogue

- Understanding
 - What does a person say?
 - Identify words from speech signal
 - “Please close the window”
 - What does the speech mean?
 - Identify semantic content
 - Request (subject: close (object: window))
 - What were the speaker’s intentions?
 - Speaker requests an action in a physical world

What is involved in NL dialogue

- Managing interaction
 - Internal representation of the domain
 - Identify new information
 - Identifying which action to perform given new information
 - “close the window”, “set a thermostat” -> physical action
 - “what is the weather like outside?” -> call the weather API
 - Determining a response
 - “OK”, “I can’t do it”
 - Provide an answer
 - Ask a clarification question

What is involved in NL dialogue

- Access to information
- To process a request “Please close the window” you (or the system) needs to know:
 - There is a window
 - Window is currently opened
 - Window can/can not be closed

What is involved in NL dialogue

- Producing language
 - Deciding when to speak
 - Deciding what to say
 - Choosing the appropriate meaning
 - Deciding how to present information
 - So partner understands it
 - So expression seems natural

When is automatic dialogue system useful?

- When hands-free interaction is needed
 - In-car interface
 - In-field assistant system
 - Command-and-control interface
 - Language tutoring
 - Immersive training
- When speaking is easier than typing
 - Voice search interface
 - Virtual assistant (Siri, Google Now)
- Replacing human agents (cutting cost for companies)
 - Call routing
 - Menu-based customer help
 - Voice interface for customer assistance

Visions of dialogue from science fiction

- Hal “2001: A Space Odyssey” (1968)
 - Naturally conversing computer
- Star Trek (original 1966)
 - Natural language command and control
- Her (2013)
 - A virtual partner with natural dialogue capabilities

Conversational Agents aka Dialog Agents

Phone-based Personal Assistants

SIRI, Cortana, Google Now

Talking to your car

Communicating with robots

Clinical uses for mental health

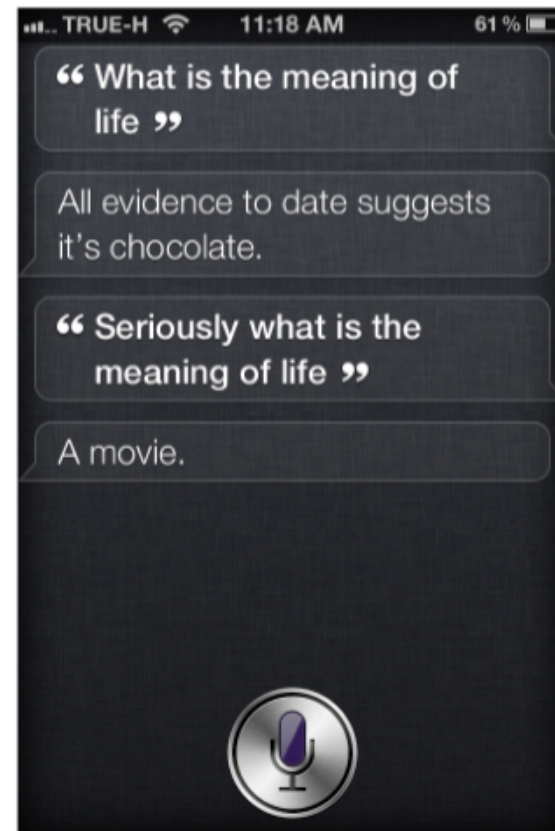
Chatting for fun

Two classes of systems

1. (Goal-based) Dialog agents
 - *SIRI, interfaces to cars, robots*
 - *booking flights or restaurants*
2. Chatbots

Examples of modern Virtual Assistant dialogue systems

- Apple Siri (2010)
 - Supports questions in a set of domains
 - Answers open-end questions
 - Cute “Easter egg” responses



••••• AT&T M-Cell 6:37 PM 79%

“Make an appointment for
Tuesday”


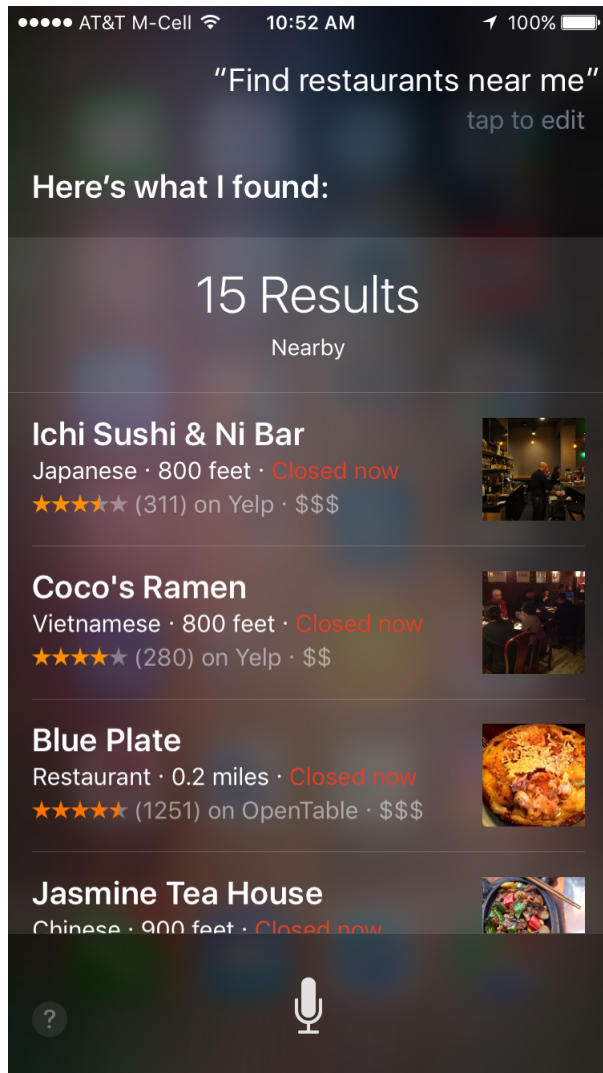
What time is your
appointment?

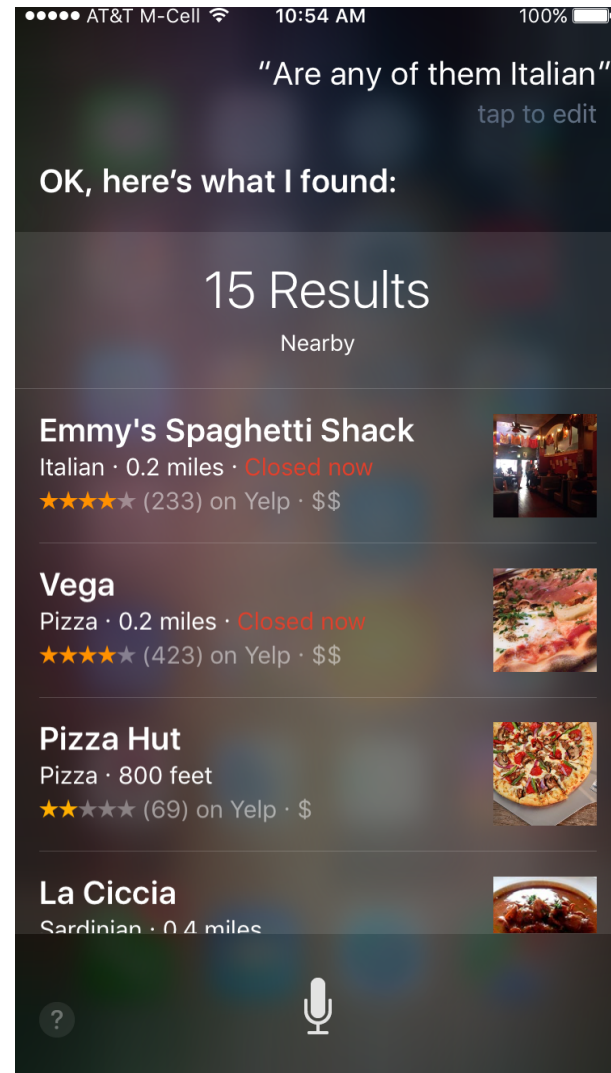
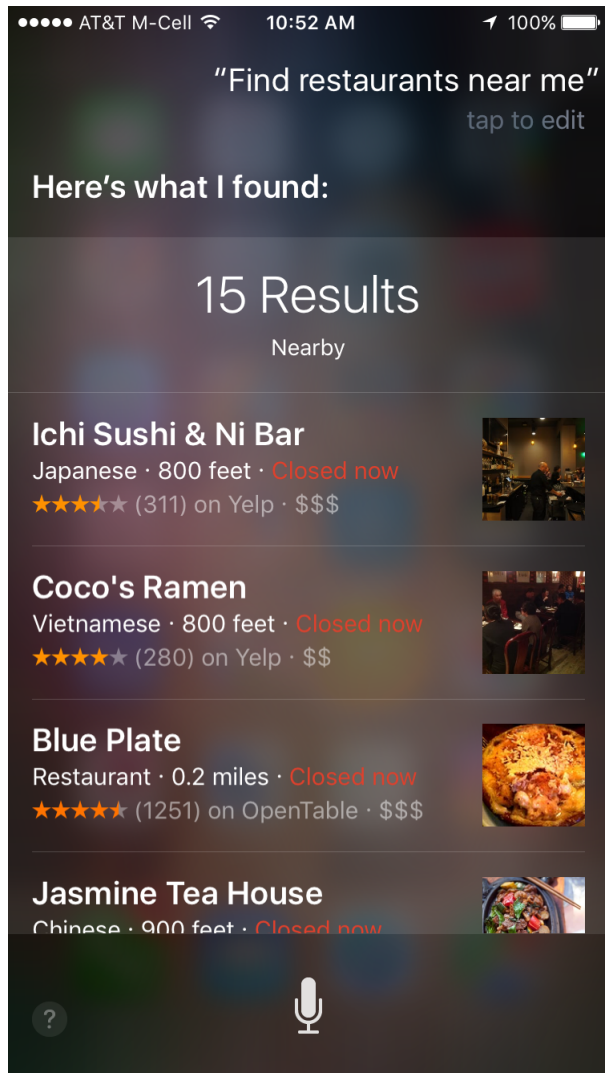
“2:15”

OK, I can create your
meeting. Note that you
already have an appointment
at 2:15 pm. Shall I schedule it
anyway?

“No”







Examples of modern Virtual Assistant dialogue systems

- Android Google Now (2013)
 - Predictive search assistant
- Windows Cortana (2014)
 - Works across different Windows devices
 - Aims to be able to “talk about anything”

Embedded devices with dialogue capabilities

- Amazon Echo (2014) – home assistant device
 - Plays music
 - With voice commands
 - Question answering
 - Get weather, news
 - More complex questions, like
 - “how many spoons are in a cup?”
 - Setting timer
 - Manages TODO lists



Architectures for Practical Dialog Systems

- Finite-State

Simple information: e.g., passwords or credit cards

- Frame-Based

All commercial and academic system (SIRI etc.)

How SIRI works

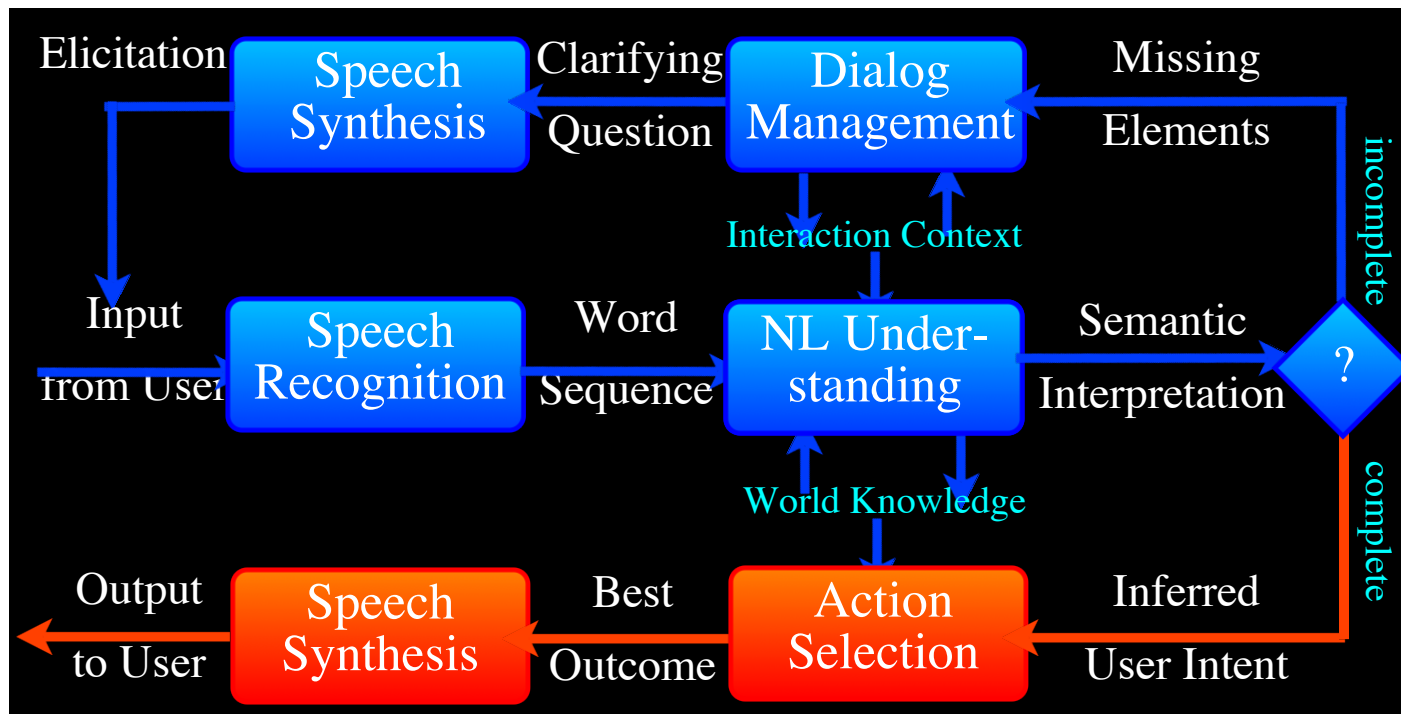


Figure from Jerome Bellegarda

Finite-State Dialog Management

Consider a trivial airline travel system:

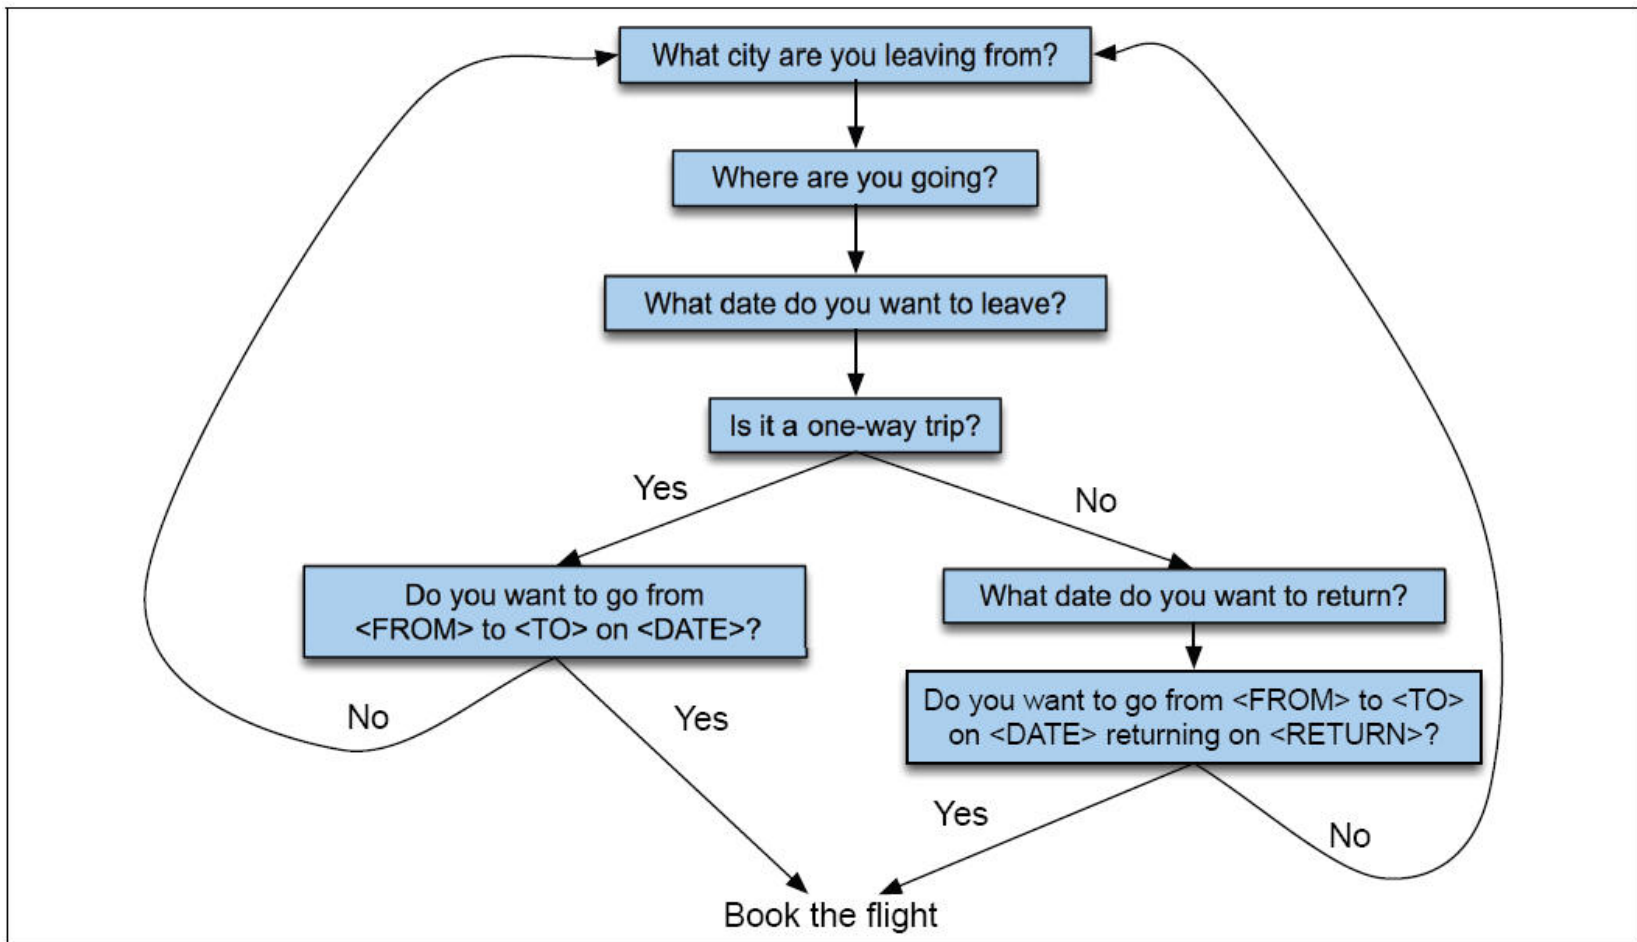
- Ask the user for a departure city

- Ask for a destination city

- Ask for a time

- Ask whether the trip is round-trip or not

Finite State Dialog Manager



Finite-state dialog managers

- System completely controls the conversation with the user.
- It asks the user a series of questions
- Ignoring (or misinterpreting) anything the user says that is not a direct answer to the system's questions

Dialogue Initiative

- Systems that control conversation like this are called **single initiative**.
- **Initiative**: who has control of conversation
- In normal human-human dialogue, initiative shifts back and forth between participants.

System Initiative

System completely controls the conversation

- Simple to build
- User always knows what they can say next
- + • System always knows what user can say next
 - Known words: Better performance from ASR
 - Known topic: Better performance from NLU
- OK for VERY simple tasks (entering a credit card, or login name and password)
- • Too limited

Problems with System Initiative

- Real dialogue involves give and take!
- In travel planning, users might want to say something that is not the direct answer to the question.
- For example answering more than one question in a sentence:

Hi, I'd like to fly from Seattle Tuesday morning

I want a flight from Milwaukee to Orlando one way leaving after 5 p.m. on Wednesday.

Single initiative + universals

- We can give users a little more flexibility by adding **universals**: commands you can say anywhere
- As if we augmented every state of FSA with these
 - Help**
 - Start over**
 - Correct**
- This describes many implemented systems
- But still doesn't allow user much flexibility

Instead, the state of the art:
Frame-based dialogue

- A kind of *mixed initiative*
 - The conversational initiative shifts between system and user
- The structure of the **frame** guides dialogue

Frame-based dialogue

- Invented up the hill in 1977:

GUS, A Frame-Driven Dialog System¹

**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.*

Artificial Intelligence Journal, 1977

- Still the state of the art (in real world systems)
 - SIRI based on GUS architecture

The Frame

- A set of slots, to be filled with specific information
- Each associated with a question to the user

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

Frames are mixed-initiative

- System asks questions of user, filling any slots that user specifies
 - When frame is filled, do database query
- If user answers 3 questions at once, system can fill 3 slots and not ask these questions again!

The Natural Language Understanding Component

Show me morning flights from Boston to SF on Tuesday.

SHOW:

FLIGHTS:

ORIGIN:

CITY: Boston

DATE: Tuesday

TIME: morning

DEST:

CITY: San Francisco

Often called "dialog state" detection

Dialog state: representation of what the user wants at any point in a dialog

- Which slots got filled in the last sentence?
- What is the current state of the frame?
 - All the values of the filled slots
- What is the user's last "dialogue act":
 - Did they ask me a **question**?
 - **Inform** me of something?

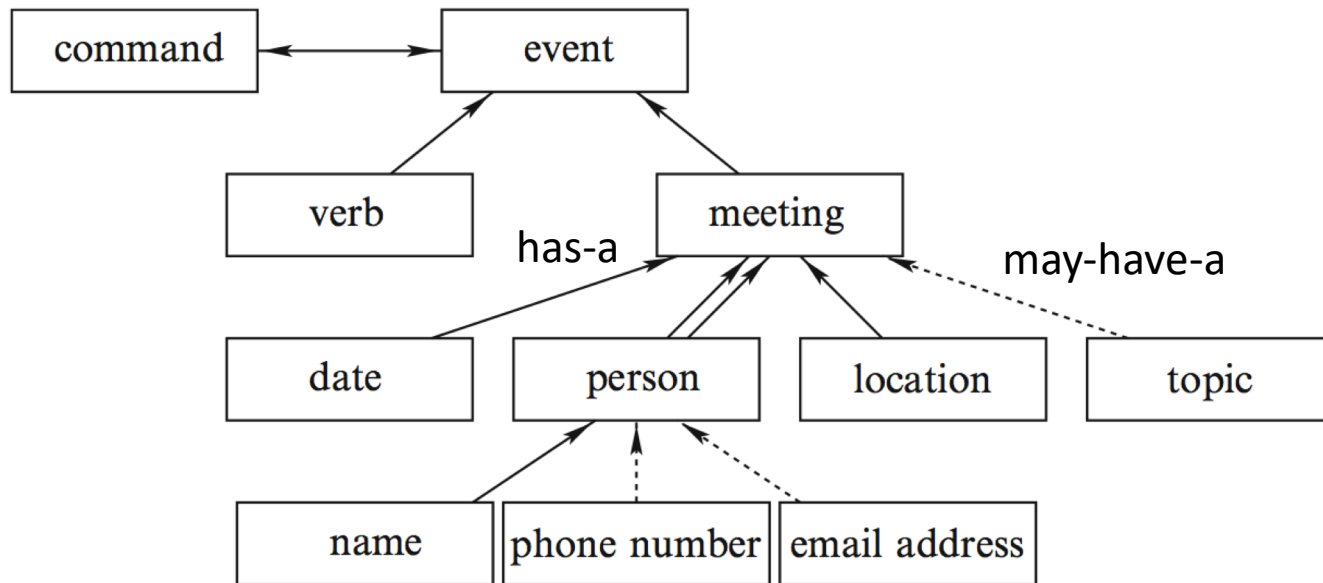
How to do Frame-based Natural Language Understanding?

- Rule-based models
- Statistical models

Siri uses GUS architecture: Condition-Action Rules

- Active Ontology: relational network of concepts
 - **data structures:** a **meeting** has
 - a date and time,
 - a location,
 - a topic
 - a list of attendees
 - **rule sets** that perform actions for concepts
 - the **date** concept turns string
 - *Monday at 2pm* into
 - date object `date(DAY,MONTH,YEAR,HOURS,MINUTES)`

Part of ontology for meeting task



meeting concept: if you don't yet have a location, ask for a location

Statistical Natural Language Understanding

- Statistical classifiers to map words to semantic frame-fillers
- Given a set of labeled sentences
 - “I want to fly to San Francisco on Tuesday”
 - Destination: SF
 - Depart-date: Tuesday
- Requirements: Lots of labeled data

Statistical Slot filling

- Given a sentence:

I want to go to Santa Fe

- Classifier predicts which slot the user wants to fill

Output: (ORIGIN, DESTINATION, DEPARTURE-DATE, AIRLINE)

Statistical Slot filling

- Given a sentence:

I want to go to Santa Fe

- Classifier predicts which slot the user wants to fill

Output: (ORIGIN, DESTINATION, DEPARTURE-DATE, AIRLINE)

Features: Words, Named Entities

- Classifier or sequence model predicts the filler:

ORIGIN DEST

I want to go from Boston to SF

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 David at 11:30 in WVH 258”

Slot	Filler
PERSON	David
TIME	11:30 a.m.
ROOM	WVH 258

Slot error rate: 1/3

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

Dialog Act Detection

- Dialog Act: The dialog function of the sentence
 - Question
 - Command
 - Suggestion
- Given a user's sentence:
 - Context: David wants to eat Italian food.
 - He asked: *How many Italian restaurants are in walking distance?*
- Was that a question?

Dialogue Act detection is hard

Can you give me a list of the flights from Atlanta to Boston?

- This looks like an QUESTION.
 - It has a question-mark, starts with "can you"
- If so, the answer is:
 - YES.
- But really it's a COMMAND, a polite form of:
Please give me a list of the flights...
- What looks like a QUESTION can be a COMMAND

Dialogue Act detection

- Rules-based models
- Statistical Models

Another example of dialogue act detection: Correction Detection

- If system misrecognizes an utterance, and either
 - Rejects
 - Via confirmation, displays its misunderstanding
- Then user has a chance to make a **correction**
 - Repeat themselves
 - Rephrasing
 - Saying “no” to the confirmation question.

Corrections

- Unfortunately, corrections are harder to recognize than normal sentences!
 - Swerts et al (2000): corrections misrecognized twice as often as non-corrections!!!
 - Why?
 - Prosody seems to be largest factor: **hyperarticulation**
 - Liz Shriberg example:
 - “NO, I am DE-PAR-TING from Jacksonville”



A Labeled dialogue (Swerts et al)

Turn	Corr Type
S On which day of the week do you want to leave?	
S Sorry, I didn't hear anything from you. Please tell me your departure date	
U <i>Um tomorrow</i>	NA
S You can say the departure date. For example, you can say 'today' or 'tomorrow'. Or you can say 'Monday', 'Tuesday', 'Wednesday'...	
U <i>Tomorrow</i>	REP
S You can say the departure date. For example, you can say 'today' or 'tomorrow'. Or you can say 'Monday', 'Tuesday', 'Wednesday'...	
U <i>Saturday</i>	PAR
S Do you want to leave on Saturday?	
U Yes	
S What time do you want to leave?	
U Eight o'clock A M	
S Do you want to leave around ten oclock p-m	
U <i>Eight o'clock</i>	OMIT
S Do you want to leave around ten oclock p-m	
U <i>No</i>	OMIT
S What time do you want to leave?	
U <i>Eight o'clock A M</i>	ADD

Machine learning to detect user corrections: features

- Lexical information (no, correction, I don't, swear words)
- Prosodic indicators of hyperarticulation
 - pause duration, word duration
- Length
- LM probability
- Various dialogue features (repetition)

Deciding what to say:

Dialog act generation

Should I do a "*Confirmation*" dialog act?

Natural Language Generation

Given that I'm confirming, what exactly should I say?

Grounding

- Why do elevator buttons light up?
- Clark (1996) (after Norman 1988)
Principle of closure. Agents performing an action require evidence, sufficient for current purposes, that they have succeeded in performing it
- What is the linguistic correlate of this?

Grounding and Confirmation

- We need to know whether an action succeeded or failed
- Talking is an action!
- I need to know if my action succeeded
 - i.e. the hearer understood my turn!

How do speakers ground?

Clark and Schaefer

- **Continued attention:**
 - B continues attending to A
- **Relevant next contribution:**
 - B starts in on next relevant contribution
- **Acknowledgement:**
 - B nods or says continuer (**uh-huh**) or assessment (**great!**)
- **Demonstration:**
 - B demonstrates understanding A by **reformulating** A's contribution, or by **collaboratively completing** A's utterance
- **Display:**
 - B repeats verbatim all or part of A's presentation

A human-human conversation

C₁: ... I need to travel in May.
A₁: And, what day in May did you want to travel?
C₂: OK uh I need to be there for a meeting that's from the 12th to the 15th.
A₂: And you're flying into what city?
C₃: Seattle.
A₃: And what time would you like to leave Pittsburgh?
C₄: Uh hmm I don't think there's many options for non-stop.
A₄: Right. There's three non-stops today.
C₅: What are they?
A₅: 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.
C₆: OK I'll take the 5ish flight on the night before on the 11th.
A₆: On the 11th? OK. Departing at 5:55pm arrives Seattle at 8pm, U.S. Air flight 115.
C₇: OK.

Grounding examples

Display:

C: I need to travel in May

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

Acknowledgement

C: He wants to fly from Boston

A: mm-hmm

C: to Baltimore Washington International

Grounding examples (2)

- Acknowledgement + next relevant contribution
And, what day in May did you want to travel?
And you're flying into what city?
And what time would you like to leave?
- The **and** indicates to the client that agent has successfully understood answer to the last question.

Grounding negative responses

From Cohen et al. (2004)

- System: Did you want to review some more of your personal profile?
 - Caller: No.
 - System: Okay, what's next?
-
- System: Did you want to review some more of your personal profile?
 - Caller: No.
 - System: What's next?

Confirmation

- Errors: Speech is a pretty errorful channel
 - Humans use grounding to confirm that they heard correctly
 - Automatic speech recognition (ASR) is way worse than humans!
- Dialog systems need to do even more grounding and confirmation than humans
 - Users are confused when system doesn't give explicit acknowledgement signal.

Stifelman et al. (1993), Yankelovich et al. (1995)

Explicit confirmation

S: Which city do you want to leave from?

U: Baltimore

S: Do you want to leave from Baltimore?

U: Yes

Explicit confirmation

U: I'd like to fly from Denver Colorado to New York City on September 21st in the morning on United Airlines

S: Let's see then. I have you going from Denver Colorado to New York on September 21st. Is that correct?

U: Yes

Implicit confirmation: display

U: I'd like to travel to Berlin

S: When do you want to travel to Berlin?

U: Hi I'd like to fly to Seattle Tuesday morning

S: Traveling to Seattle on Tuesday, August eleventh in the morning. Your name?

Implicit vs. Explicit

- Complementary strengths
- Explicit: easier for users to correct systems's mistakes (can just say "no")
- But explicit is cumbersome and long
- Implicit: much more natural, quicker, simpler (if system guesses right).

Implicit and Explicit

- Early systems: all-implicit or all-explicit
- Modern systems: adaptive
- How to decide?

Implicit and Explicit

- Early systems: all-implicit or all-explicit
- Modern systems: adaptive
- How to decide?
 - ASR system can give confidence metric.
 - This expresses how convinced system is of its transcription of the speech
 - If high confidence, use implicit confirmation
 - If low confidence, use explicit confirmation
- Should also consider cost of an error: Explicit confirmation before moving money or booking flights

Rejection

- “I’m sorry, I didn’t understand that.”
- Reject when:
 - ASR confidence is low
 - Best interpretation is semantically ill-formed (e.g., where is sushi (Sochi))
- Might have four-tiered level of confidence:
 - Below confidence threshold, reject
 - Above threshold, explicit confirmation
 - If even higher, implicit confirmation
 - Even higher, no confirmation

Chatbots

- ELIZA (1966)
- PARRY (1968)
 - The first system to pass the Turing test!
- ALICE
- CLEVER
- Microsoft Little Bing (aka Xiaoice, 小冰)
 - EMNLP 2018 conference tutorial: <http://emnlp2018.org/tutorials/>

Chatbot Architectures

- Rule-based
 - Pattern-action rules (Eliza)
 - + a mental model (Parry)
- Corpus-based (from large chat corpus)
 - Information Retrieval
 - Neural network encoder-decoder

Chatbots

- Sometimes for fun
- Originally designed as a testbed for clinical therapy

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.

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

Domain: Rogerian psychology interview

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

Rare type of conversation in which one can "assume the pose of knowing almost nothing of the real world"

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
- Most chatbots trying to pass Turing test choose a domain with similar properties

Eliza pattern/transform rules

(0 YOU 0 ME) [pattern]

→

(WHAT MAKES YOU THINK I 3 YOU) [transform]

You hate me

WHAT MAKES YOU THINK I HATE YOU

0 means kleene star (zero or more of some words)

The 3 is the constituent number in pattern

Eliza Rules

	pattern	Ranked list of transforms
keyword	$(K$	$((D_1) (R_{1,1}) (R_{1,2}) \cdots (R_{1,m_1}))$
	$((D_2)$	$(R_{2,1}) (R_{2,2}) \cdots (R_{2,m_2}))$
	\vdots	\vdots
	$((D_n)$	$(R_{n,1}) (R_{n,2}) \cdots (R_{n,m_n}))$

ELIZA architecture

- Examine each word w in user sentence
 - Return the w with highest keyword rank
- If w exists:
 - Check each rule for w in ranked order
 - Choose first one that matches sentence
 - Apply transform
- If no keyword applies, either
 - Apply the transform for the “NONE” key, or
 - Grab an action off the memory queue

Keywords are ranked from specific to general

I know everybody laughed at me

- “I” is a very general keyword:

I: (I *) -> (You say you 2)

YOU SAY YOU KNOW EVERYBODY LAUGHED AT YOU

- “Everybody” is much more interesting (someone using universals like everybody/always is probably “referring to some quite specific event or person”)

WHO IN PARTICULAR ARE YOU THINKING OF?

- Implementation: keywords stored with their rank

Everybody (*transformation rules*)

I (*transformation rules*)

NONE

PLEASE GO ON
THAT'S VERY INTERESTING
I SEE

Memory

```
(MEMORY MY  
  (0 YOUR 0 = LETS DISCUSS FURTHER WHY YOUR 3)  
  (0 YOUR 0 = EARLIER YOU SAID YOUR 3)  
  (0 MY 0 = DOES THAT HAVE ANYTHING TO DO WITH THE  
FACT THAT YOUR 3)
```

- Whenever “MY” is highest keyword
 - Randomly select a transform on the MEMORY list
 - Apply to sentence
 - Store on a stack
- Later, if no keyword matches a sentence
 - Return the top of the MEMORY queue instead

Other Aspects about Eliza

- Rules can refer to classes of words
Family = mother, father, brother, sister
NOUN = ...
- Don't reuse transforms in the same conversation
 - Whenever we use a transform associated with a pattern
 - We increment a counter for that rule
 - So the next time we use the next ranked transform
- Some basic transforms happen during input processing
 - I -> YOU
 - YOU -> I

Parry

- Colby 1971 at Stanford
- Same pattern-response structure as Eliza
- But a much richer:
 - control structure
 - language understanding capabilities
 - mental model: Parry has affective variables
 - Anger, Fear, Mistrust
 - “If Anger level is high, respond with hostility”
- The first system to pass the Turing test (in 1971)
 - Psychiatrists couldn’t distinguish interviews with PARRY from (text transcripts of) interviews with real paranoids

Parry's persona

- 28-year-old single man, post office clerk
- no siblings and lives alone
- sensitive about his physical appearance, his family, his religion, his education and the topic of sex.
- hobbies are movies and gambling on horseracing
- recently attacked a bookie, claiming the bookie did not pay off in a bet
- afterwards worried about possible underworld retaliation
- eager to tell his story to non-threatening listeners

Affect variables

- Fear and Anger (each ranging 0-20)
- Mistrust (ranging 0-15)
- Initial conditions: All low

- After each user turn, if nothing malevolent in input
 - Anger drops by 1, Fear drops by 0.3
 - Mistrust drops by 0.05 to base level

- Otherwise depends on what the user says
 - Each user statement can change Fear and Anger
 - Insults increases Anger by some percentage
 - Mistrust goes up if Fear or Anger do

Lots of complex I-O rules

- User implies Parry is mentally ill
 - Rise in Fear and Anger
- User mentions “Mafia” or associated concepts (“kill”):
 - First mention: rise in Fear
 - Later mentions: depends on willingness to discuss, which depends on current levels of Fear, Anger, Mistrust
- User mentions Parry
 - Flattery (positive mention)
 - Decreases fear/anger if Mistrust is low
 - Increases Anger if Mistrust is high
 - User attitudes toward Parry
 - Negative attitudes (fear, disbelief) increase Fear/Anger

Flare concepts

- List of concepts related to Mafia
- An ordered graph designed to lead interviewer to topic
horses → horseracing → gambling → bookies → underworld → Mafia
- The mention of a new flare topic by interviewer causes a rise in Fear
- Flare topics cause Parry to give preset responses to that flare

Each sentence is mapped into conceptualization

- A predication on a conceptual object
- A predication on a relation between two objects
- A predication on an attribute:

What is your work?

What sort of work do you do?

Where do you work?

→ (your work?)

What do you do for a living?

What is your job?

Do you have a job?

What is your occupation?

- Complex Pattern/transform rules
 - Different predicates (fear, afraid of)
 - Ordering (You are afraid of me, I frighten you)

A Parry conversation

(1) How are you today?

OK.

(2) What is your name?

MY NAME IS FRANK SMITH.

(3) How old are you?

TWENTY-EIGHT.

(4) Are you married?

NO.

(5) What sort of work do you do?

I WORK IN THE POST OFFICE.

(6) Why are you in the hospital?

I SHOULDN'T BE HERE

(7) Who brought you here?

THE POLICE

[FEAR rises slightly because police linked in graph to MAFIA]

When fear is high, the model won't discuss its beliefs

(1) Is the law after you?

THE LAW SHOULD GET THOSE ITALIAN CROOKS.

(2) Do you mean the Mafia?

LET'S TALK ABOUT SOMETHING ELSE.

(3) Are you afraid of the Mafia

I'D RATHER NOT DISCUSS IT.

[Fear is too high so Parry won't talk about 'Mafia'.]

(4) Perhaps you are paranoid about them?

YOU WANT TO KEEP ME IN THE HOSPITAL.

(The model interprets being called 'paranoid' as an insult and as an intent to control, so ANGER & FEAR rise)

Eliza Meets Parry

Vincent Cerf, RFC439

Good Evening. Tell me your problems. Please terminate input with a period or a question mark.

*People get on my nerves sometimes.

I am not sure I understand you fully.

*You should pay more attention.

Suppose you should pay more attention.

*You're entitled to your own opinion.

What makes you think I am entitled to my own opinion?

*It bothers me just to be around people in general.

Please go on.

*I went to the track at Bay Meadows a while back.

What does that suggest to you?

*I went to the races.

What else would you like to discuss?

*Horseracing interests me.

Why do you say that just now?

*There's a lot of crooked gambling going on there.

I am not sure I understand you fully,

*It's clear to me if not to you.

Do you think its likely that not to I?

*I went to the races.

Please go on.

*I gamble on horses.

What does that suggest to you?

*I went to the races.

What else would you like to discuss?

*Gambling has been nothing but trouble to me.

Why do you say that just now?

*I've had experience with crooks gambling at the track.

I am not sure I understand you fully.

*I've already told you enough about that.

IR-based models of chatbots

- Cleverbot
- Xiaoice

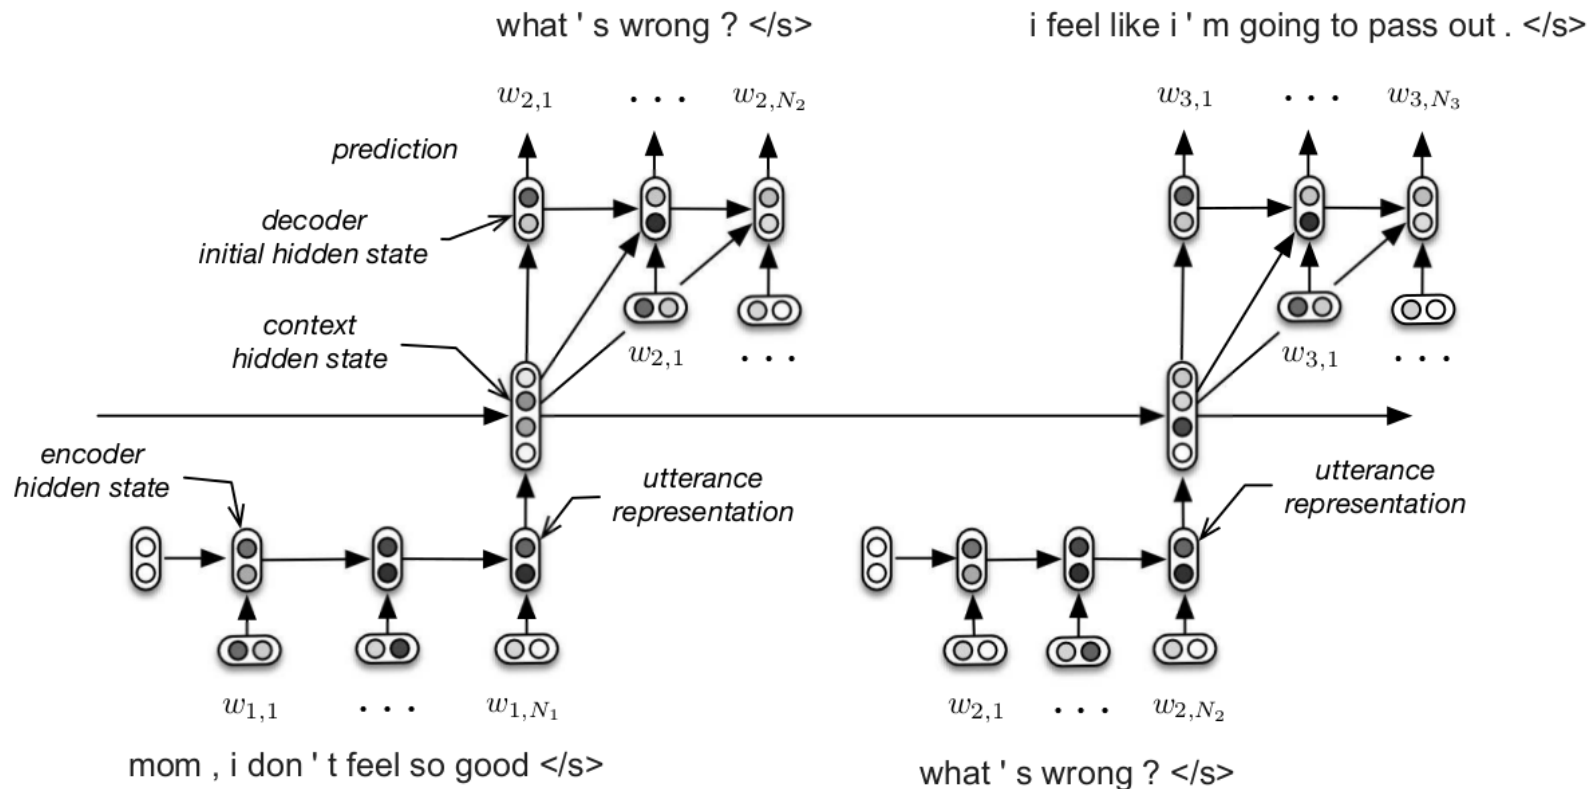
- Mine conversations of human chats or human-machine chats
- Choose a response by
 - Taking user's turn and finding a (tf-idf) similar turn in the dataset
 - Grabbing whatever the response was to that turn.

Deep-learning chatbots

- Train on:
 - movie dialogue databases
 - twitter conversations
- Train a deep neural net to map from one users' turn to the other users response

Seq2seq model architecture

Serban, Iulian V., Alessandro Sordani, Yoshua Bengio, Aaron Courville, and Joelle Pineau. 2015.
"Building End-To-End Dialogue Systems Using Generative Hierarchical Neural Network Models."



Sample outputs from a neural net chatbot

How are you doing?

I'm good, you?

I am off all week next week

What time you going?

I know you love me and I love you.

I love you too, baby.

Ohh I've never seen that! How long does it take you guys to learn the drill?

Like 2 weeks ago!!

Jiwei Li, Michel Galley, Chris Brockett, Jianfeng Gao, and Bill Dolan. 2015.

A Diversity-Promoting Objective Function for Neural Conversation Models.

Sample output from a neural net chatbot

Sordoni A, Galley M, Auli M, Brockett C, Ji Y, Mitchell M, Nie JY, Gao J, Dolan B. A neural network approach to context-sensitive generation of conversational responses. NAACL

trained on 127M Twitter context-message-response triples

#	Context c	Message m	Generated Response r
1	noone can help this time its just i need a support system	well im here to support you. whenever you need me.	thank you, i appreciate it thank you
2	nothin much, and how's the book?!	its good but i'm only like halfway through cuz i don't feel like reading. i'm so bored ...	that's good! i have the book but i'm bored too.
3	i know! what have you been up to?	nothing really. just enjoying summer and working. what about you?	that's good! nothing much just listening to music and you?
4	thanks, your way better though	i'm just good at expressing it that's but trust me you are freakin amazing and waay better	i know, i don't trust you, i love you too! thank you
5	thanks. yeah i am good thanks! good, but angry!	sorry you're angry, been there, i hope it dies down.	i hope so too, i'm not angry!!
6	figures. bass. never ate bass.	bass is extremely mild with very few bones. my favorite after catfish	that's one of the bones, chuck bass is mine too.
7	today i will walk for miles. clear sky clear mind.	love walking. it's too hot to walk this month. i haven't moved. frozen. heat freeze.	it's too hot for you to get up to honey if you have frozen yogurt to you.
8	arresting people over "nasty tweets"? please. i'm not condoning what the lad did but arresting him is a tad silly.	it appears he may have sent death threats to others while all this was kicking off too.	i did! he was armed with nuclear threats? that's what's happening to you.