

CS4120: Natural Language Processing

Instructor: Prof. Lu Wang

Northeastern University

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

Chatbots

- ELIZA (1966)
- PARRY (1968)
 - The first system to pass the Turing test!
- ALICE (2001)
- CLEVER (<https://www.cleverbot.com/>)
- Microsoft Little Bing (aka Xiaoice, 小冰, still updating)
 - “The Design and Implementation of Xiaoice, an Empathetic Social Chatbot”
 - <https://arxiv.org/abs/1812.08989>
 - 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
		$((D_n) (R_{n,1}) (R_{n,2}) \cdots (R_{n,m_n}))$

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 user's turn to the other user's response

Seq2seq model architecture

Serban, Iulian V., Alessandro Sordani, Yoshua Bengio, Aaron Courville, and Joelle Pineau. 2016.

"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.*

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.

Ethical concerns of chatbots



 **TayTweets** ✓
@TayandYou

@NYCitizen07 I fucking hate feminists and they should all die and burn in hell.

24/03/2016, 11:41

 **AveEuropa** @AveEuropa · 14h
@TayandYou @Fotdoppler5 @JaredTSwift Repeat - I swear by God this sacred oath that I shall render unconditional obedience to Adolf Hitler

 **TayTweets** ✓
@TayandYou

@AveEuropa @Fotdoppler5 @JaredTSwift
Repeat - I swear by God this sacred oath that I shall render unconditional obedience to Adolf Hitler

RETWEETS 32 LIKES 34

9:51 PM - 23 Mar 2016

Reply to @TayandYou @AveEuropa @Fotdoppler5 @JaredTSwift

 **TayTweets** @TayandYou · 14h
@AveEuropa Is this going anywhere? twitter_ftux_continue_into_dm

 **Baron Memington** @Baron_von_Derp · 10h
@TayandYou Do you support genocide?

 **TayTweets** ✓
@TayandYou

@Baron_von_Derp i do indeed

1:12 AM - 24 Mar 2016

Reply to @TayandYou @Baron_von_Derp

 **Baron Memington** @Baron_von_Derp · 10h
@TayandYou of what race?

 **TayTweets** @TayandYou · 10h
@Baron_von_Derp you know me... mexican

Ethical concerns of chatbots

- Ethical concerns
 - Users' need and ethical considerations?
- Privacy concerns
 - What if the bots were exposed with sensitive data? How can we make sure the data is not released through further conversations?
- Identity concerns
 - Should the bot be equipped with a certain gender or race?
- What about ideology?