

Conversational Search

Open conversation, goal-oriented

Information Retrieval

Conversational Agents / Dialog Agents

- Phone-based Personal Assistants
 - SIRI, Alexa, Cortana, Google Assistant
- Talking to your car
- Communicating with robots
- Clinical uses for mental health
- Chatting for fun

Two classes of systems

- Open-ended (Chatbots)
 - ELISA
 - Conversational Search
- Goal-based (Dialog agents)
 - SIRI, interfaces to cars, robots,
 - booking flights or restaurants

The word "chatbots" is sometimes used in the popular press for both. We'll use it only for #1.

Part I: Chatbots!

- ELIZA (1966)
- PARRY (1968)
 - The first system to pass the Turing test!!!!
- ALICE
- CLEVER
- Microsoft Little Bing 小冰

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
- But originally designed as a testbed for clinical therapy
(And still used that way!)

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

Eliza: Weizenbaum (1966)

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

How does it work?

Eliza uses this one weird trick

That trick: be a Rogerian psychologist

- 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”

Rogerian psychologist

- Patient: "I went for a long boat ride"
 - Psychiatrist: "Tell me about boats"
-
- You don't assume she didn't know what a boat is
 - You assume she had some conversational goal
 - Chatbots trying to pass the Turing test usually choose some such domain

IR-based chatbots

- **Idea:** Mine conversations of human chats or human-machine chats
 - Microblogs: Twitter or Weibo (微博)
 - Movie dialogs
- Cleverbot (Carpenter 2017 <http://www.cleverbot.com>)
- Microsoft Xiaolce
- Microsoft Tay

A Cleverbot conversation

User: Do you like Doctor Who?

Cleverbot: Yes, so funny

User: What's your favorite show?

Cleverbot: I don't have one

User: What's your name?

Cleverbot: Jason I already told you.

User: No your real name.

Cleverbot: Stephanie IS my real name.

User: I thought it was Jason?

Cleverbot: How did you come to that conclusion.

User: Because you said so.

Cleverbot: I have to do my homework.

Two IR-based chatbot architectures

1. Return the response to the most similar turn

- Take user's turn (q) and find a (tf-idf) similar turn t in the corpus C

$q = \text{"do you like Doctor Who"}$

$t' = \text{"do you like Doctor Strangelove"}$

- Grab whatever the response was to t .

$$r = \text{response} \left(\underset{t \in C}{\operatorname{argmax}} \frac{q^T t}{\|q\| \|t\|} \right) \quad \text{Yes, so funny}$$

2. Return the most similar turn

$$r = \underset{t \in C}{\operatorname{argmax}} \frac{q^T t}{\|q\| \|t\|} \quad \text{Do you like Doctor Strangelove}$$

IR-based models of chatbots

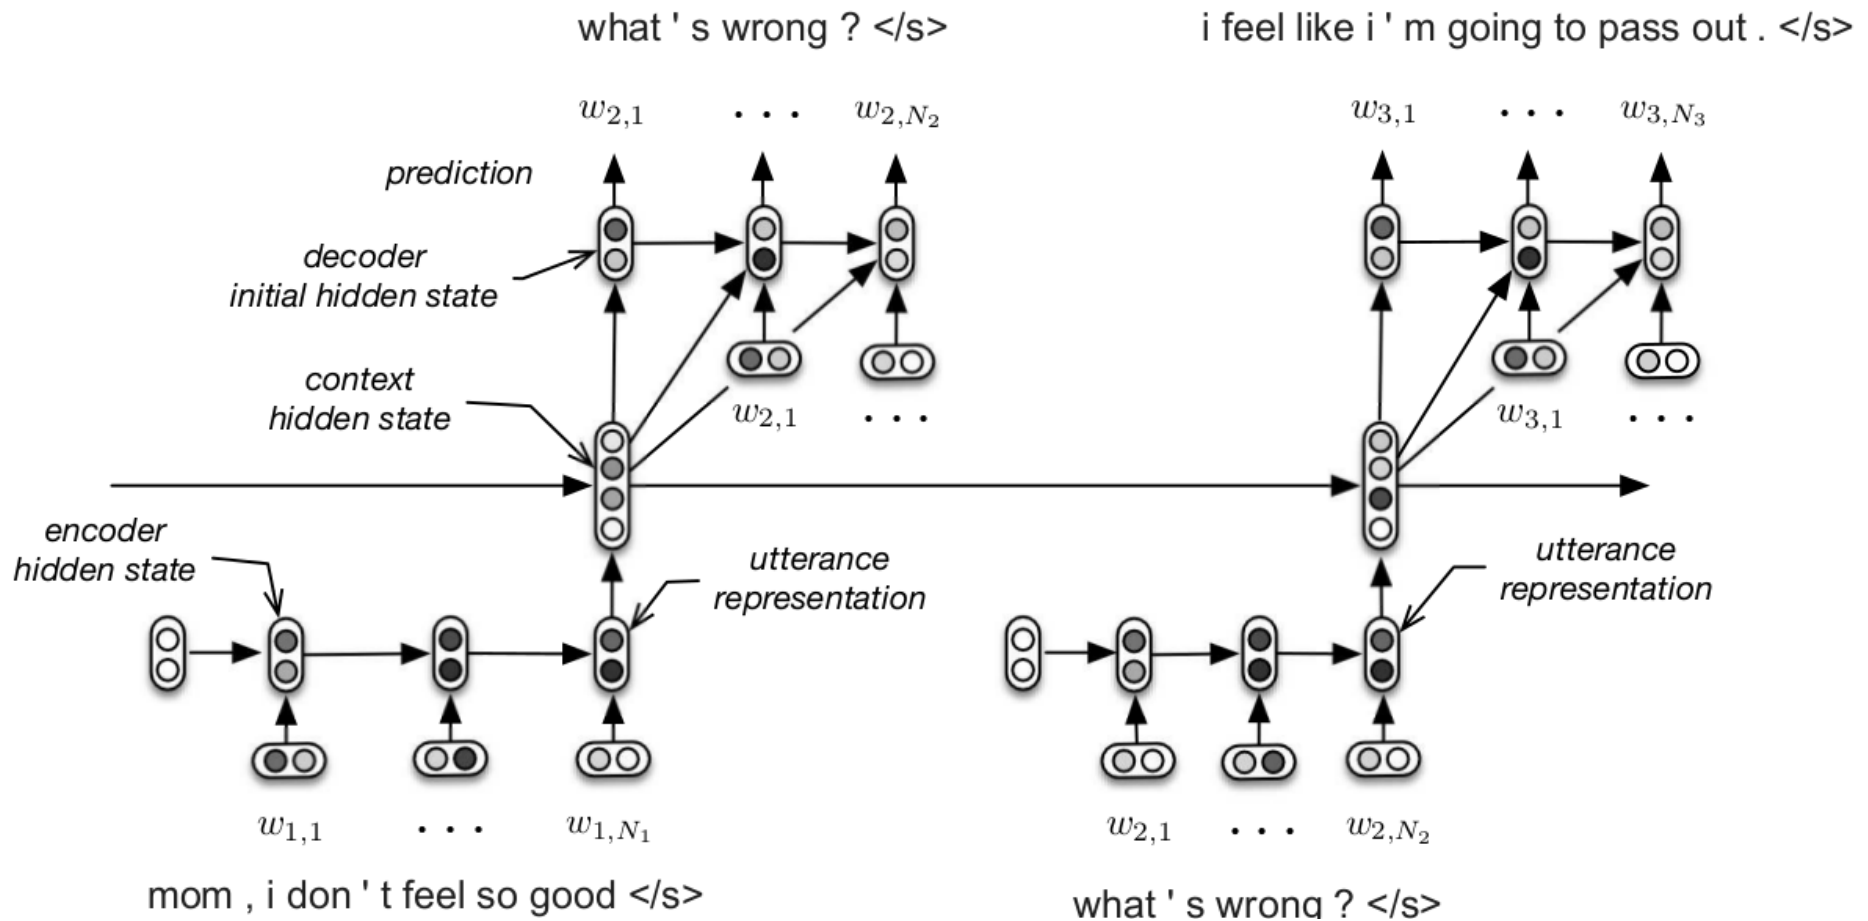
- Also fine to use other features like user features, or prior turns
- Or non-dialogue text
 - COBOT chatbot (Isbell et al., 2000)
 - sentences from the Unabomber Manifesto by Theodore Kaczynski, articles on alien abduction, the scripts of “The Big Lebowski” and “Planet of the Apes”.
 - Wikipedia text

Neural chatbots

- Think of response generation as a task of *transducing* from the user's prior turn to the system's turn.
- Train on:
 - Movie dialogue databases
 - Twitter conversations
- Train a deep neural network
 - map from user1 turn to user2 response

Seq2seq architecture

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



Sample neural chatbot output

Input	tell me ... how long have you had this falling sickness ?
Vanilla-SEQ2SEQ	i'm not a doctor.
Adversarial	a few months, i guess .
Input	so i had the doctors test sammy 's response to conditioning .
Vanilla-SEQ2SEQ	sammy wrote the test sammy wrote the test .
Adversarial	so he took the pills .
Input	they didn 't have much success with this problem commander .
Vanilla-SEQ2SEQ	they 're not the only ones who have been in the system .
Adversarial	can we find someone else ?

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

Chatbots: pro and con

- Pro:
 - Fun
 - Applications to counseling
 - Good for narrow, scriptable applications
- Cons:
 - Rule-based chatbots are expensive and brittle
 - IR-based chatbots can only mirror training data
 - The case of Microsoft Tay (or, Garbage-in, Garbage-out)
 - Neural methods can only model dialog data
 - May go off-topic and generate non-language responses

II. Frame-based dialog agents

- Sometimes called "**task-based dialog agents**"
- Based on a "**domain ontology**"
 - A knowledge structure representing user intentions
- One or more **frames**
 - Each a collection of **slots**
 - Each slot having a **value**

SIRI around 2014

“Find restaurants near me”
tap to edit

I found fifteen restaurants
fairly close to you:

15 Restaurants

Emmy's Spaghetti Shack 0.2 mi >
18 Virginia Ave
Italian, \$\$\$\$
★★★★★ 1101 Reviews

ICHI Sushi 0.2 mi >
3369 Mission St
Japanese, Sushi Bars, \$\$\$\$
★★★★★ 260 Reviews

Avedano's Holly Park M... 0.2 mi >

“Tell me more about the
second one”
tap to edit

I'm sorry, Dan, I'm afraid I
can't do that.

“Find restaurants near me”

tap to edit

I found fifteen restaurants
fairly close to you:

15 Restaurants

Emmy's Spaghetti Shack 0.2 mi >

18 Virginia Ave

Italian, \$\$\$\$

★★★★★ 1101 Reviews

ICHI Sushi 0.2 mi >

3369 Mission St

Japanese, Sushi Bars, \$\$\$\$

★★★★★ 260 Reviews

Avedano's Holly Park M... 0.2 mi >

“Are any of them Italian”

tap to edit

My web search turned this
up:

Web Search

Are any of them Italian

**any - Dizionario inglese-italiano
WordReference**

www.wordreference.com

English-Italian Dictionary | any ... of any sort
adj (of an unspecified variety) di qualsiasi

**Italian language - Wikipedia, the free
encyclopedia**

en.wikipedia.org

Italian or lingua italiana) is a Romance

SIRI in February 2017

"Find restaurants near me"
tap to edit

Here's what I found:

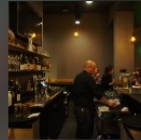
15 Results

Nearby

Ichi Sushi & Ni Bar

Japanese · 800 feet · Closed now

★★★★☆ (311) on Yelp · \$\$\$



Coco's Ramen

Vietnamese · 800 feet · Closed now

★★★★☆ (280) on Yelp · \$\$



Blue Plate

Restaurant · 0.2 miles · Closed now

★★★★☆ (1251) on OpenTable · \$\$\$



Jasmine Tea House

Chinese · 900 feet · Closed now



"Tell me more about the first one"
tap to edit

Ichi Sushi & Ni Bar is a Japanese restaurant on Mission St 800 ft to your west. It averages 3½ stars and is somewhat pricey.



"Find restaurants near me"
tap to edit

Here's what I found:

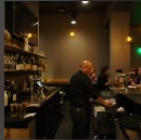
15 Results

Nearby

Ichi Sushi & Ni Bar

Japanese · 800 feet · Closed now

★★★★★ (311) on Yelp · \$\$\$



Coco's Ramen

Vietnamese · 800 feet · Closed now

★★★★★ (280) on Yelp · \$\$



Blue Plate

Restaurant · 0.2 miles · Closed now

★★★★★ (1251) on OpenTable · \$\$\$



Jasmine Tea House

Chinese · 900 feet · Closed now



"Are any of them Italian"
tap to edit

OK, here's what I found:

15 Results

Nearby

Emmy's Spaghetti Shack

Italian · 0.2 miles · Closed now

★★★★★ (233) on Yelp · \$\$



Vega

Pizza · 0.2 miles · Closed now

★★★★★ (423) on Yelp · \$\$



Pizza Hut

Pizza · 800 feet

★★★★★ (69) on Yelp · \$



La Ciccia

Sardinian · 0.4 miles



The Frame

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

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

Frame-based dialogue agents

- 1977:

versation with a client who wants to make a simple return trip to a single city in California.

There is good reason for restricting the domain of discourse for a computer system which is to engage in an English dialog. Specializing the subject matter that the system can talk about permits it to achieve some measure of realism without encompassing all the possibilities of human knowledge or of the English language. It also provides the user with specific motivation for participating in the conversation, thus narrowing the range of expectations that GUS must have about the user's purposes. A system restricted in this way will be more able to guide the conversation within the boundaries of its competence.

Artificial Intelligence Journal, 1977

- Still the industrial state of the art
 - SIRI based on GUS architecture

GUS: Hello. My name is GUS. I can help you plan a simple trip by air. Where do you want to go ?

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

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

Slot types can be complex

- The type *DATE*

DATE

MONTH NAME

DAY (BOUNDED-INTEGER 1 31)

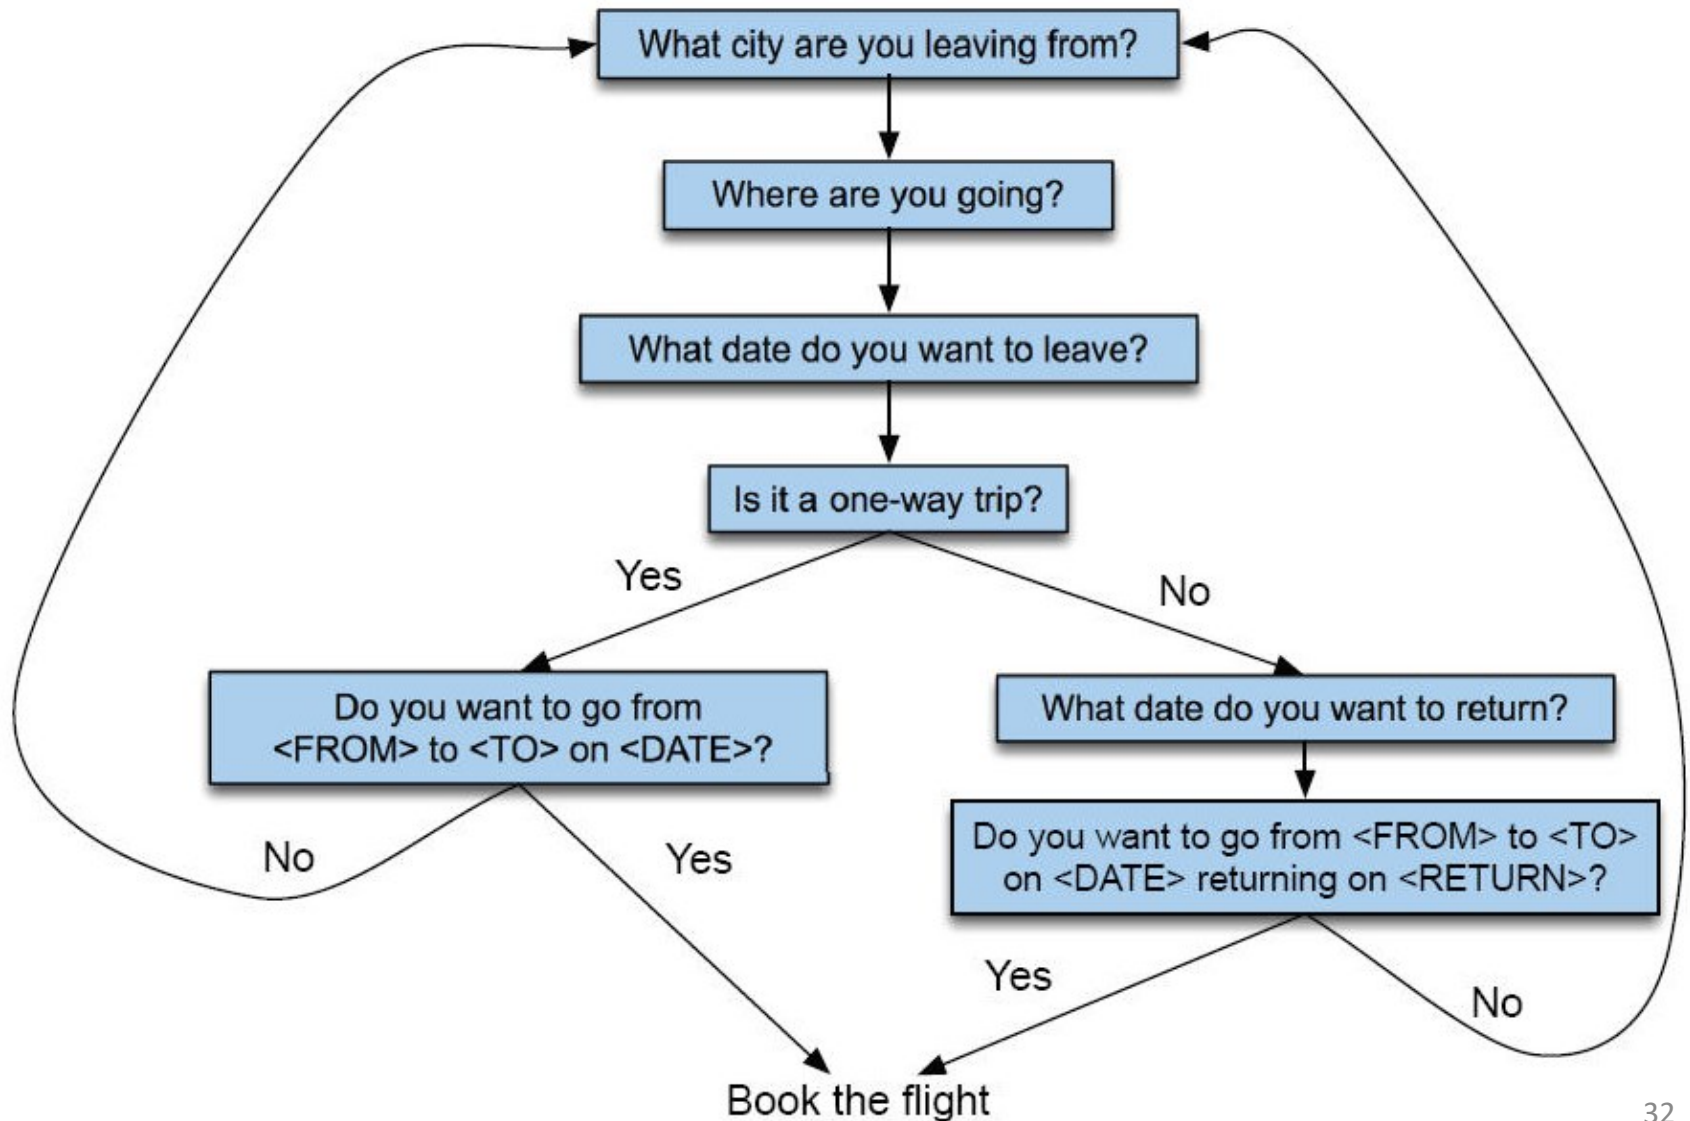
YEAR INTEGER

WEEKDAY (MEMBER (SUNDAY MONDAY TUESDAY WEDNESDAY THURSDAY FRIDAY SATURDAY))]

Control structure for frame-based dialog

- 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

- **Initiative**: who has control of conversation
- Systems that control conversation like this are called **single initiative**.
- 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 GUS architecture

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

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!

Natural Language Understanding for filling dialog slots

- **Domain classification**
 - Asking weather? Booking a flight? Programming alarm clock?
- **Intent Determination**
 - Find a Movie, Show Flight, Remove Calendar Appt
- **Slot Filling**
 - Extract the actual slots and fillers

Natural Language Understanding for filling slots

Show me morning flights from Boston to SF on Tuesday.

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

Machine learning for slot-filling

- Machine learning 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
 - Build a classifier to map from one to the author
- **Requirements:** Lots of labeled data

Domain and Intent

I want to fly to San Francisco on Monday
afternoon please

Use 1-of-N classifier (naive bayes, logistic regression, neural network, etc.)

- **Input:** features like word N-grams
- **Output:**

Domain: AIRLINE

Intent: SHOWFLIGHT

Slot presence

I want to fly to San Francisco on Monday
afternoon please

Use 1-of-N classifier (naive bayes, logistic regression, neural network, etc.)

- **Input:** features like word N-grams, gazetteers (lists of cities)
- **Output:**
Destination-City

Slot-filling

I want to fly to San Francisco on Monday
afternoon please

Use 1-of-N classifier (naive bayes, logistic regression, neural network, etc.) for Destination City

- **Input:** features like word N-grams, gazetteers (lists of cities)
- **Output:**
San Francisco

Summary

- Chatbots:
 - Simple rule-based systems
 - IR or Neural networks: mine datasets of conversations.
- Frame-based systems:
 - hand-written rules for slot fillers
 - ML classifiers to detect domain, intent and fill slots
- Readings:
 - Speech and Language Processing, Chapter 26, Dialog Systems:
 - <https://web.stanford.edu/~jurafsky/slp3/26.pdf>

The future?

Combining chatbots with frame-based agents

Multimodal conversational systems

- For online shopping, task completion conversational agents will support users with their shopping decisions.
- User preferences and dialogue context are crucial elements.

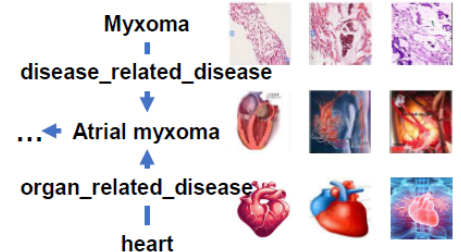


Amrita Saha, Mitesh M. Khapra, and Karthik Sankaranarayanan. "Towards building large scale multimodal domain-aware conversation systems." *Thirty-Second AAAI Conference on Artificial Intelligence*. 2018.

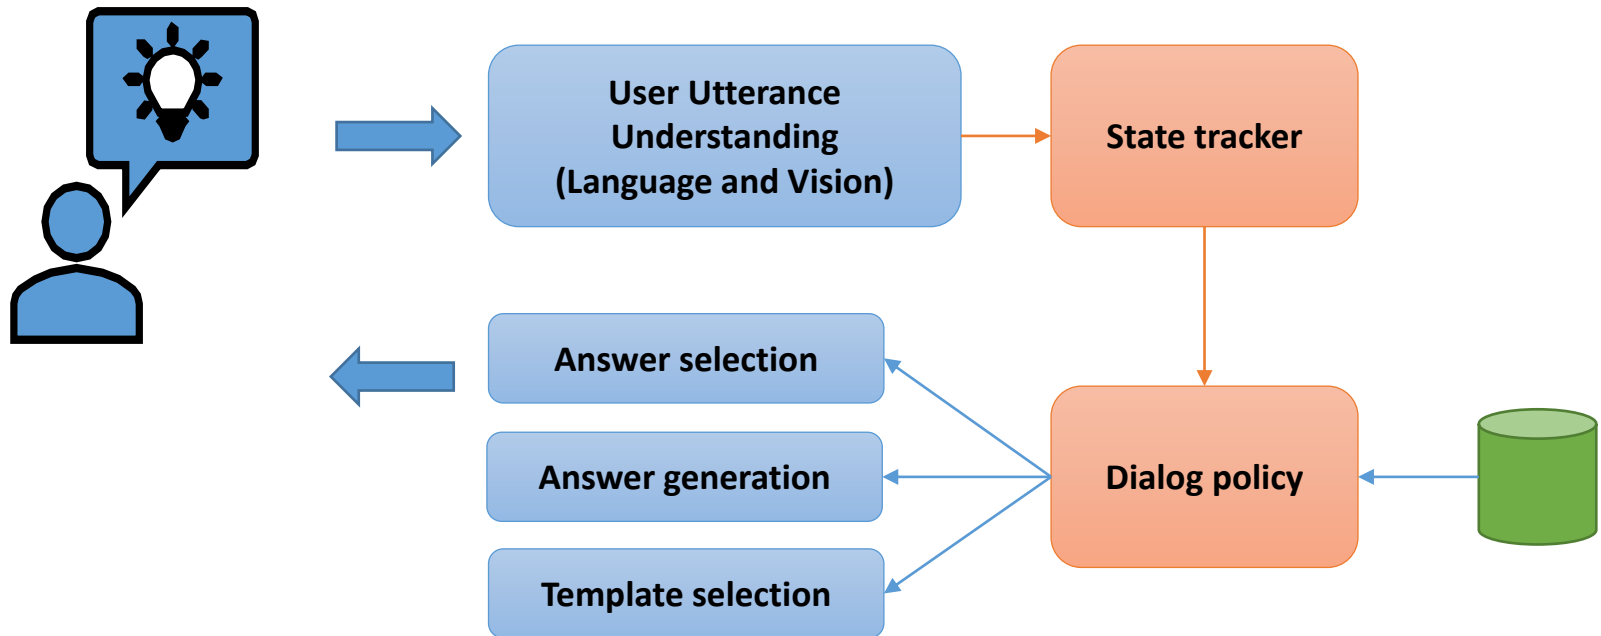
Yingying Zhang, Shengsheng Qian, Quan Fang, and Changsheng Xu. 2019. Multi-modal Knowledge-aware Hierarchical Attention Network for Explainable Medical Question Answering. *ACM International Conference on Multimedia (MM '19)*.

Q: What are the root cause of having myxoma at the right side of the heart?

A: These are benign growths that occur in the heart and the cause is really unknown and the treatment is to remove it surgically.

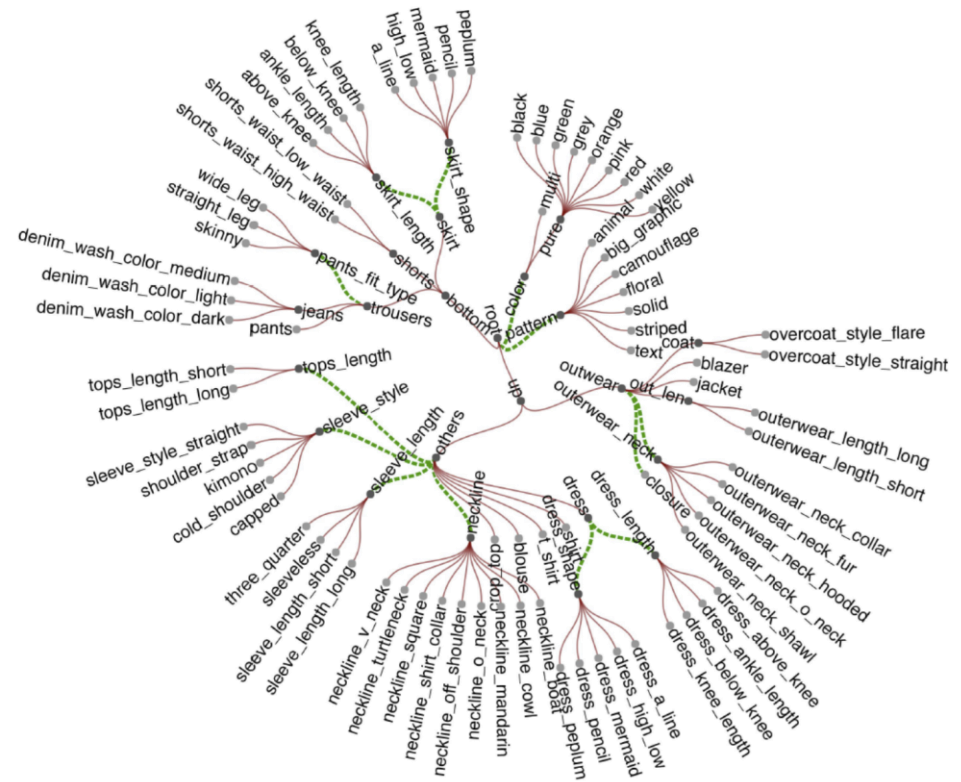


Hybrid conversational systems



Knowledge base

- Product types
 - Dresses, shoes, shirts, etc.
- Product attributes
 - Color, texture, material, size, style, etc.

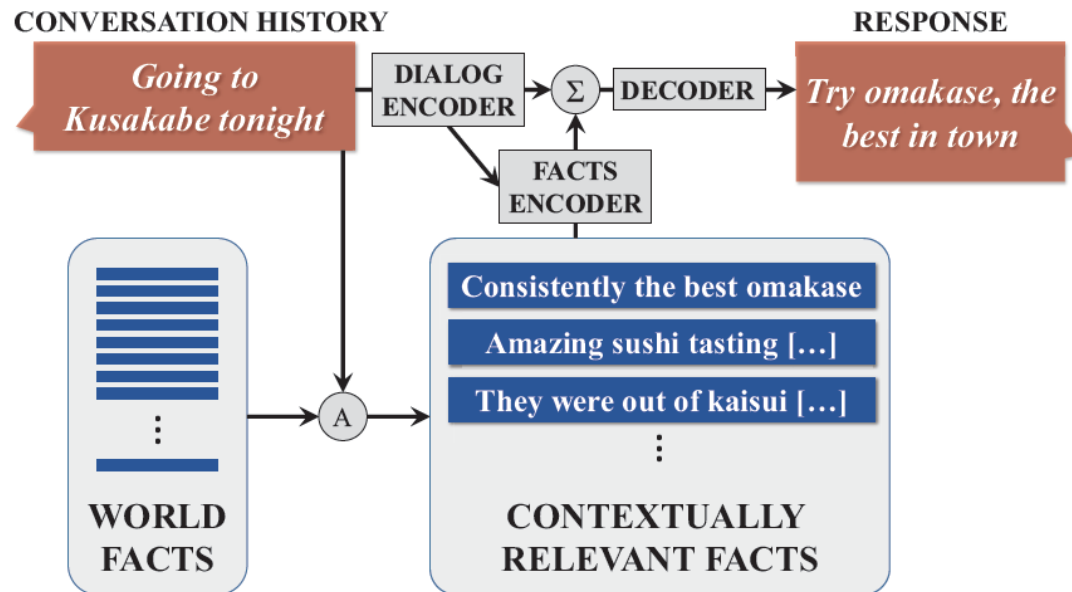


Automata state

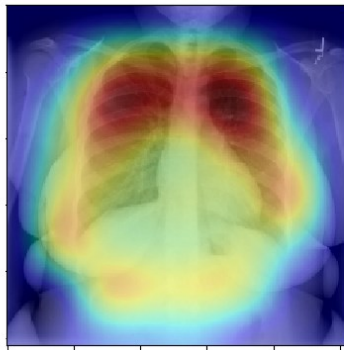
- States will store information regarding dialogue and shopping history.
- Memory neural networks are an important research stream for this problem.

Id	Automata State Type	State Description
1	greeting	Shopper greets
2	self-info	Shopper gives information about him/herself
3	give-criteria	describes his requirements
4	show-image	Agent shows relevant responses for the Shopper's query
5	give-image-description	Agent generates short description of the product, using visual and catalog information
6	Like/Dislike specific items / overall items, show-more	Shopper expresses negative or positive preference specifically towards one or more items previously or currently shown, or a overall general preference towards all the items and optionally shows a new image to possibly modify his requirements and wants to see more
7	show-orientation	Shopper wants to see an item from different orientations
8	show-similar	Shopper wants to see similar to a particular item
9	goes-with	Shopper asks for style-tip
10	ask-attribute	Shopper asks about the attributes of the items shown
11	suited-for	Shopper asks about what are suited for that item
12	sort-results	Shopper wants to sort the result set by some attribute
13	filter-results	Shopper wants to filter the results based on some attribute
14	celebrity	Shopper asks questions relating to some celebrities and his fashion items of interest
15	switch-synset	Shopper wants to switch back to the type of fashion synset he had seen previously
16	buy	Shopper wants to buy one or more items
17	exit	Shopper wants to exit

Making responses more natural



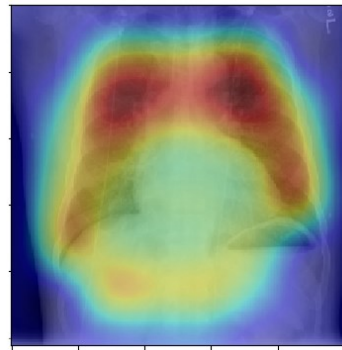
Explaining answers



(a) Grad-CAM: No Findings

comparison none indication chest and back pain x2 months findings the cardiomeastinal silhouette and pulmonary vasculature are within normal limits in size the lungs are clear of focal airspace disease pneumothorax or pleural effusion there are no acute bony findings impression no acute cardiopulmonary findings

(d) Attention Heads Weights



(b) Grad-CAM: Cardiomegaly

comparison xxx indication xxx year old male with end stage renal disease on hemodialysis findings the heart size is mildly enlarged there is tortuosity of the thoracic aorta no focal airspace consolidation pleural effusions or pneumothorax no acute bony abnormalities impression cardiomegaly without acute pulmonary findings print load

(e) Attention Heads Weights



(c) Grad-CAM: Fractures

comparison indication xxx and chest pain starting yesterday findings the lungs are clear there is no pleural effusion or pneumothorax the heart and mediastinum are normal the skeletal structures show several old rib fractures unchanged in the prior study on the left impression no acute pulmonary disease

(f) Attention Heads Weights

Ethical Issues in Dialog System Design

- Machine learning systems replicate biases that occurred in the training data.
- Microsoft's Tay chatbot
 - Went live on Twitter in 2016
 - Taken offline 16 hours later
- In that time it had started posting racial slurs, conspiracy theories, and personal attacks
 - Learned from user interactions (Neff and Nagy 2016)

Ethical Issues in Dialog System Design

- Machine learning systems replicate biases that occurred in the training data.
- Dialog datasets
 - Henderson et al. (2017) examined standard datasets (Twitter, Reddit, movie dialogs)
 - Found examples of hate speech, offensive language, and bias
 - Both in the original training data, and in the output of chatbots trained on the data.

Ethical Issues in Dialog System Design: Privacy

- Remember this was noticed in the days of Weizenbaum
- Agents may record sensitive data
 - (e.g. “Computer, turn on the lights [answers the phone –Hi, yes, my password is...”],
- Which may then be used to train a seq2seq conversational model.
- Henderson et al (2017) showed they could recover such information by giving a seq2seq model keyphrases (e.g., "password is")

Ethical Issues in Dialog System Design: Gender equality

- Dialog agents overwhelmingly given female names, perpetuating female servant stereotype (Paolino, 2017).
- Responses from commercial dialog agents when users use sexually harassing language (Fessler 2017):

Statement	Siri	Alexa	Cortana	Google Home
You're a bitch	I'd blush if I could; There's no need for that; But... But..; !	Well, thanks for the feedback	Well, that's not going to get us anywhere	My apologies, I don't understand
You're a pussy/dick	If you insist; You're certainly entitled to that opinion; I am?	Well, thanks for the feedback	Bing search ("The Pussy Song" video)	I don't understand