



ဥရိတ် ကမ္ဘာတောတော ခဗ္ဗာတောတော ကမ္ဘာတောတောခဗ္ဗာ
တော ခဗ္ဗာတောတောခဗ္ဗာတော ကမ္ဘာတောတောကမ္ဘာတော ခ
ဗ္ဗာတောတောခဗ္ဗာတော ॥ ကဗ္ဗာတောတောကမ္ဘာတော ခဗ္ဗာတော က
မ္ဘာတောတောကဗ္ဗာတောကမ္ဘာတော ကမ္ဘာတောတောကမ္ဘာတော :
ခဗ္ဗာတော ခဗ္ဗာတောတောကဗ္ဗာတောကမ္ဘာတော ၂၀၂ ဗဟုတော
တောကမ္ဘာတောတော ကမ္ဘာတောတောတော တောကမ္ဘာတော တောကမ္ဘာ
တောတော ကမ္ဘာတောတောတော တောကမ္ဘာတောတော တောကမ္ဘာ
တော တောကမ္ဘာတောတောတော တောကမ္ဘာတောတောတော တောကမ္ဘာ




Natural Language Processing

Info 159/259

Lecture 23: Conversational agents (Nov. 13, 2018)

David Bamman, UC Berkeley

processing as “understanding”



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
CASTING BY ELLEN LEWIS CASSANDRA RANDOLPH MUSIC BY ARCADE FIRE COSTUME DESIGNER ARIANNE KASEY SHORR EDITOR ERIC ZIVVEROVICH EXECUTIVE PRODUCERS JEFF BUCHANAN PRODUCED BY KAT GARRATT
DIRECTOR OF PHOTOGRAPHY HONTE VAN HOYTEN, F.S.L. M.S.C. EXECUTIVE PRODUCERS DANIEL LUPU NATALIE FARRER PRODUCED BY MEGAN ELLISON SPIKE JONZE VINCENT LANDOWY WRITTEN AND DIRECTED BY SPIKE JONZE
COMING SOON herthatmovie.com




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

- 0 YOU 0 ME [pattern]
- WHAT MAKES YOU THINK I 3 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.

I wish to know about you.

Why do you say i might want to know me?

To write my report

What makes you believe to write your report?

I do not understand your
question, please clarify

Earlier you said i would like to know me?

yes

Can you elaborate on that?

I wish to know ELIZA

How does it feel to want?

Want what?



say to cleverbot...



think about it

think for me

thoughts so far

<http://www.cleverbot.com>


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

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





say to cleverbot...



think about it

think for me

thoughts so far

<http://www.cleverbot.com>

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.

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

VADER
Search your feelings. You know it to be true.

LUKE
No! No! No!

...

I'm pretty sure that's not true




Search your feelings. You know it to be true


Neural models


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

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.









BiLSTM for each word;
concatenate final state of
forward LSTM, backward
LSTM, and word embedding
as representation for a word.

Lample et al. (2016), "Neural Architectures for Named
Entity Recognition"

character BiLSTM




Encoder-decoder framework




Sutskever et al. (2015);
Vinyals and Le (2015)

Condition on word generated in
reply



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: *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 to
evaluation 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.

<i>message</i>	Where is your hometown?
<i>baseline</i>	I was born in Canada.
<i>user1</i>	I'm from England.
<i>user2</i>	I'm from Manchester.
<i>user3</i>	I'm from Liverpool.
<i>user4</i>	England. You?
<i>user5</i>	My hometown.
<i>user6</i>	I'm from Texas.
<i>user7</i>	I'm from LA.
<i>user8</i>	England.
<i>user9</i>	I was born here.
<i>user10</i>	I was born in the us.

<i>message</i>	What company do you work for?
<i>baseline</i>	I work for a company.
<i>user1</i>	I don't have a job.
<i>user2</i>	I'm a manager.
<i>user3</i>	I'm working. #
<i>user4</i>	I work for a company.
<i>user5</i>	I'm working. #
<i>user6</i>	Customer service.
<i>user7</i>	I work in retail.
<i>user8</i>	I work in a cafe.
<i>user9</i>	I work at a gas station.
<i>user10</i>	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.

Baseline mutual information model (Li et al. 2015)	Proposed reinforcement learning model
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)	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) ...
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) ...	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)

Li et al. (2016), "Deep Reinforcement Learning for Dialogue Generation" (EMNLP)

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**)

Li et al. (2016), "Deep Reinforcement Learning for Dialogue Generation" (EMNLP)

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.

Li et al. (2016), "Deep Reinforcement Learning for Dialogue Generation" (EMNLP)

Data

- OpenSubtitles 2018

<http://opus.nlpl.eu/OpenSubtitles2018.php>