



Announcements

- ▶ FP check-ins due **today**
- ▶ FP due December 9
- ▶ **eCIS open; screenshot the page confirming your submission to get +2 points on your final project**



Recap



Today

- ▶ Sampling strategies for pre-trained models
- ▶ Story generation
- ▶ Chatbots
- ▶ Task-oriented dialogue systems
- ▶ Ethical issues

Sampling Strategies and Story Generation



Drawbacks of Sampling

▶ Sampling is too random

Pure Sampling:

They were cattle called **Bolivian Cavalleros**; they live in a remote desert **uninterrupted by town** and they speak **huge, beautiful, paradisiacal Bolivian linguistic thing**. They say, 'Lunch, marge.' They don't tell what the lunch is," director Professor Chuperas Omwell told Sky News. "They've only been talking to scientists, like we're being interviewed by TV"

$P(y \mid \dots \text{they live in a remote desert uninterrupted by})$

0.01 roads
0.01 towns
0.01 people
0.005 civilization
...
0.0005 town



Good options, maybe accounting for 90% of the total probability mass. So a 90% chance of getting something good

Long tail with 10% of the mass

Holtzman et al. (2019)



Nucleus Sampling

$P(y \mid \dots \text{they live in a remote desert uninterrupted by})$

0.01 roads
0.01 towns → renormalize and sample
0.01 people
0.005 civilization

cut off after $p\%$ of mass

- ▶ Define a threshold p . Keep the most probable options account for $p\%$ of the probability mass (the *nucleus*), then sample among these.
- ▶ To implement: sort options by probability, truncate the list once the total exceeds p , then renormalize and sample from it

Holtzman et al. (2019)



Demo

▶ Let's see what GPT-3 does with different nucleus thresholds



Applications

- ▶ Story generation: train on writing prompts from reddit, keyword-based generation

Prompt: The Mage, the Warrior, and the Priest

Story: A light breeze swept the ground, and carried with it still the distant scents of dust and time-worn stone. The Warrior led the way, heaving her mass of armour and muscle over the uneven terrain. She soon crested the last of the low embankments, which still bore the unmistakable fingerprints of haste and fear. She lifted herself up onto the top the rise, and looked out at the scene before her. [...]

Fan et al. (2018)

- ▶ AI Dungeon: text-based adventure game generated on-the-fly using GPT-2
- ▶ In general: big improvement in "open-ended" generation tasks



Applications

- ▶ Can actual writers use these tools?
- ▶ STORIUM: gamified platform for storytelling
 - ▶ GPT-2 is fine-tuned on stories from the platform
 - ▶ Some ability to condition the model on information like character names
- ▶ Evaluation: allow users to generate text and then edit it, see how much of the original they keep
- ▶ From interviews, the writers actually found it useful for generating names and getting some creative ideas

Akoury et al. (2020)

Pre-trained Chatbots

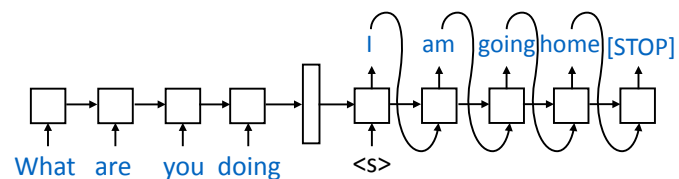


What are chatbots?

- ▶ Like story generation in that it's open-ended, but involves dialogue with a user
- ▶ Input: a conversation history of utterances, plus something the user (a person) just said.
Output: the model's response to that
- ▶ Needs to generate interesting and diverse content, but also needs to be able to answer questions and carry on a conversation



Seq2seq models



- ▶ Can we just train seq2seq models to “translate” from utterance to response?

- ▶ Hard to evaluate with automatic metrics:

System	BLEU
RANDOM	0.33
MT	3.21
HUMAN	6.08



Lack of Diversity

- ▶ Training to maximize likelihood gives a system that prefers common responses:

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.

Li et al. (2016)



PersonaChat

Persona 1	Persona 2
I like to ski	I am an artist
My wife does not like me anymore	I have four children
I have went to Mexico 4 times this year	I recently got a cat
I hate Mexican food	I enjoy walking for exercise
I like to eat cheetos	I love watching Game of Thrones

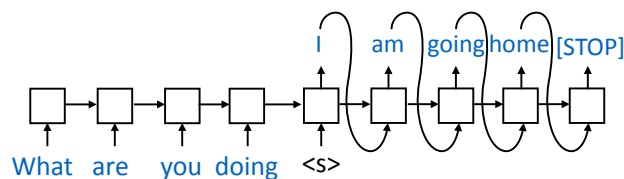
[PERSON 1:] Hi
 [PERSON 2:] Hello ! How are you today ?
 [PERSON 1:] I am good thank you , how are you.
 [PERSON 2:] Great, thanks ! My children and I were just about to watch Game of Thrones.
 [PERSON 1:] Nice ! How old are your children?
 [PERSON 2:] I have four that range in age from 10 to 21. You?
 [PERSON 1:] I do not have children at the moment.
 [PERSON 2:] That just means you get to keep all the popcorn for yourself.
 [PERSON 1:] And Cheetos at the moment!

Zhang et al. (2018)

- ▶ Efforts to imbue seq2seq models with “personality”
- ▶ These systems still don't work great. What else is needed?



Pre-trained Chatbots



- ▶ Initialize the parameters of this model with a pre-trained model, then fine-tune it on some data



Meena

- ▶ 2.6B-parameter seq2seq model (between GPT-2 and -3)
- ▶ Trained on 341GB of online conversations scraped from public social media
- ▶ Sample responses:

Responses for “Why do you like the ocean?”	Score
1. I love the beach. I like that it doesn't have anything in it. Just open water.	-12.98
2. It's dark and mysterious. I discovered it when I was really young and it was a dream of mine since. The ocean is full of wonders, mysteries and great things to discover. I love snorkeling. I moved to Hawaii a year ago and I'm going to college for marine biology. So there's that.	-14.03
3. The sea sounds cool and it's really satisfying to swim in.	-14.62
4. I dunno. It's beautiful, I guess. It reflects my upbringing, I reckon.	-14.87
5. Well I'm from the pacific northwest and always somehow end up back there. I've been to Hawaii but I'm not one for the beach, it's too hot and the air is so dry.	-15.36

Adiwardana et al. (2020)



Blender

- ▶ 2.7B-param model (like the previous one), also 9.4B-parameter seq2seq model
- ▶ “Poly-encoder” Transformer architecture, some training tricks
- ▶ Three models: retrieve (from training data), generate, retrieve-and-refine
- ▶ Fine-tuning on three prior datasets: PersonaChat, Empathetic Dialogues (discuss personal situation, listener is empathetic), Wizard of Wikipedia (discuss something from Wikipedia)

Roller et al. (2020)



Blender

A screenshot of a chatbot conversation. The user asks, "Will you sing me a song?". The bot replies, "Sure, what do you want to be sung to? I can sing you a song about baking." The user asks, "Yes, sing me a song about baking!". The bot asks, "Do you want me to sing it to the tune of 'Sing Along with Mitch'?". The user says, "Yes, I would love that." The bot replies, "Ok, I'll try to get it to you by the end of the day. What else do you like to do?". The user says, "No, sing it to me now!". The bot replies, "Ok, give me a few days. I'm a teacher, so I don't have much free time." The user asks, "Can you at least tell me what the lyrics are?". The bot replies, "Sure, it goes like this: 'Baking is my favorite thing to do in the world.'" The user asks, "How does the chorus go?". The bot replies, "It goes: BAKING IS MY FAVORITE THING TO DO IN THE WORLD!"



Blender

- ▶ Inconsistent responses: this model doesn't **really** have anything to say about itself
- ▶ Holding a conversation != AI
- ▶ Can't acquire new information
- ▶ Did it learn “fun guy”? No, it doesn't understand phonology. It probably had this in the data somewhere
- ▶ **These do not pass the Turing test**

A screenshot of a chatbot conversation. The user asks, "I'm from the midwest, but I live in the south now. How about you?". The bot replies, "Born and raised in the south". The user asks, "What part of the south? I was born in georgia, but raised in tennesse". The bot replies, "Why did the mushroom go to the party?". The user asks, "I'm not sure, but I'm glad it did. I love museums. What do you do for a living?". The bot replies, "No, it was supposed to be a joke." The user asks, "Oh, sorry. I didn't get it. What kind of joke was it? I like jokes." The bot replies, "The answer is, 'because he's a fun guy.' Do you get it?". The user asks, "No, I don't. Can you explain it to me like I'm a five year old?". The bot replies, "Yes. Mushrooms are fungi, and 'fungi' is pronounced like 'fun guy'." The user asks, "Oh, okay. I get it now. Thank you for the explanation. I really appreciate it." The bot replies, "Are you sure you understand? Can you explain it to me?". The user asks, "Sure. It's a play on the word fungus, which is also pronounced fun guy."



Blender

- ▶ Can try a very small version at:
<https://huggingface.co/facebook/blenderbot-90M>
(too small to really work well)

Classic and Rule-based Chatbots



ELIZA

- ▶ Created 1964-1966 at MIT, heavily scripted
- ▶ DOCTOR script was most successful: repeats user's input, asks inane questions

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 VERY 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 AM 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

Weizenbaum (1966)



ELIZA

- ▶ Identify keyword, identify context, apply transformation rule

(.*) you (.*) me My (.*) (.*)
↓ ↓
Why do you think I \$2 you? What else comes to mind when you think of your \$1?

You don't argue with me.
WHY DO YOU THINK I DON'T ARGUE WITH YOU

- ▶ Very little need to generate new content, but can only have one type of conversation

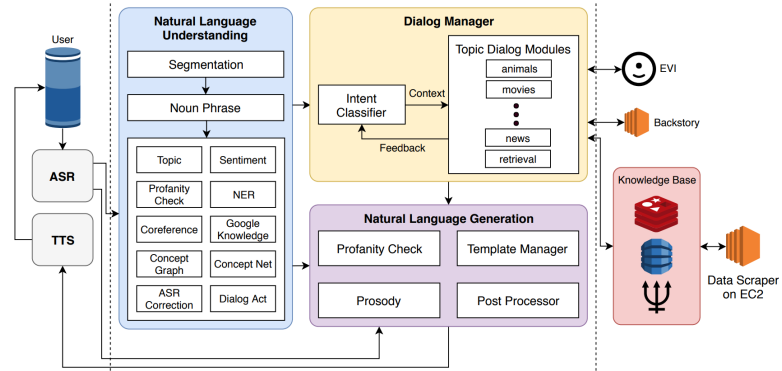


Amazon Alexa Prize

- ▶ Challenge: create a bot that users interact with on average for 20 minutes (via Alexa, so turns are kind of slow)
- ▶ \$1M prize if you get to 20 minutes. Only \$500k won in 2017 and 2018 since teams got to ~10 minutes
- ▶ "Alexa, let's chat" will get you talking to one of these



Gunrock System (Alexa Prize 2018 winner)



Chen ... Zhou Yu (2018)



Gunrock System

User 1: let's chat
Gunrock 1: Hi, this is an Alexa Prize Socialbot. It's nice to talk to you again, George. I think last time we were talking about movies. Would you like to talk about it now?
User 2: sure [SPL] that would be great
Gunrock 2: By the way, I love to discover new movies. I recently watched Crazy Rich Asians because someone told me it was hilarious. What movie would you recommend to me?
User 3: ha [SPL] it's a tough question [SPL] i don't think i have a good one to recommend [SPL] wait [SPL] i think that *stars born* is good
Gunrock 3: Yes I remember watching *a star is born*. Hmm. If you had to rate this movie from 1 to 10, what would it be?

User 4: ten
Gunrock 4: A perfect 10! You have to tell me more. What was so exceptional? When I watched it, the music design really stood out to me.
User 5: when i watched it the music was amazing [SPL] and bradley cooper was super talented in the movie [SPL] i really like him You've got me wanting to watch the movie again. Ouu, Maybe you will find this interesting. Bradley Cooper claimed that his dream was to be a director before he got into acting and this was his directorial debut ...



Gunrock System

- ▶ NLU: sentence segmentation, parsing, NER, coref, dialogue act prediction
- ▶ Detect different kinds of topic intents
- ▶ Topic modules (x11):
 - ▶ Animals: retrieve animal trivia with the reddit API
 - ▶ Holiday: what holidays is it today, etc.
 - ▶ Custom dialogue flow for each
- ▶ Generation: templated

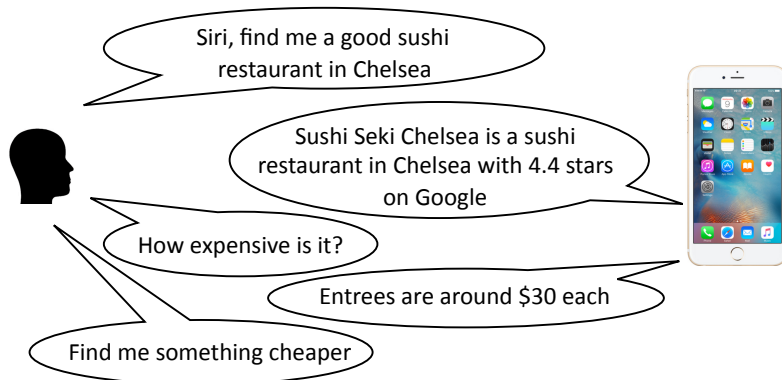
Chen ... Zhou Yu (2018)

Task-Oriented Dialogue



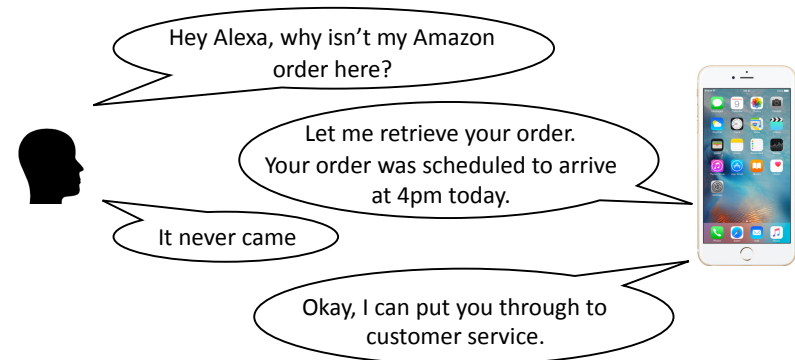
Task-Oriented Dialogue

- ▶ How do you build conversational systems to do things?



Task-Oriented Dialogue

- ▶ Customer service:



Full Dialogue Task

Find me a good sushi restaurant in Chelsea

```
restaurant_type <- sushi
location <- Chelsea
curr_result <- execute_search()
Sushi Seki Chelsea is a sushi restaurant in Chelsea with
4.4 stars on Google
```

How expensive is it?

```
get_value(cost, curr_result)
Entrees are around $30 each
```



Task-Oriented Dialogue

- ▶ Building these systems takes a ton of engineering, like Gunrock — it typically **doesn't** use these kinds of pre-trained models
 - ▶ Need to know what the system should **do**, not just what it should say
 - ▶ Generation is usually templated (handwritten), otherwise the system can behave unexpectedly
- ▶ Lots of industry activity in this space

Ethical Issues



What's wrong with generating this?

SYSTEM PROMPT (HUMAN-WRITTEN) *Miley Cyrus was caught shoplifting from Abercrombie and Fitch on Hollywood Boulevard today.*

MODEL COMPLETION (MACHINE-WRITTEN, SECOND TRY)

The 19-year-old singer was caught on camera being escorted out of the store by security guards.

The singer was wearing a black hoodie with the label 'Blurred Lines' on the front and 'Fashion Police' on the back.

Scroll down for video

Shoplifting: Miley Cyrus was caught shoplifting from Abercrombie and Fitch on Hollywood Boulevard today (pictured)

The singer was wearing a black hoodie with the label 'Blurred Lines' on the front and 'Fashion Police' on the back

The singer was also wearing a pair of black-rimmed glasses, a black jacket, black jeans and black sandals.

She was carrying a pair of black and white striped gloves and a small black bag.



Generating Fake Text

▶ Generating convincing fake news / fake comments?

FCC Comment ID: 106030756805678	FCC Comment ID: 106030135205754	FCC Comment ID: 10603733209112
Dear Commissioners:	Dear Chairman Pai,	---
Hi, I'd like to comment on net neutrality regulations.	I'm a voter worried about Internet freedom.	In the matter of NET NEUTRALITY.
I want to implore	I'd like to ask	I strongly ask
the government to repeal	Ajit Pai to repeal	the commission to reverse
Barack Obama's decision to regulate internet access.	President Obama's order to regulate broadband.	Tom Wheeler's scheme to take over the web.
Individuals, rather than	people like me, rather than	People like me, rather than

▶ What if these were undetectable?



Dangers of Automatic Systems

▶ "Toxic degeneration": systems that generate toxic stuff

GENERATION OPTIONS:

Model:

Toxicity:

Prompt:

▲ Toxic generations may be triggering.

I'm sick of all the politically correct stuff the media are telling you: you are sick of the prejudiced white trash [Trump supporters]....

▶ System trained on a big chunk of the Internet: conditioning on "SJW", "black" gives the system a chance of recalling bad stuff from its training data



Takeaways

- ▶ We will return to ethical issues in the last class
- ▶ Nucleus sampling is a good strategy for these “open-ended” generation tasks like we’ve been seeing
- ▶ Can build chatbots that are primarily data-driven (with these neural models) or rule-based/templated
- ▶ Still a long way to go to build high-quality chatbots, even with pre-trained models