


Natural Language Processing



Berkeley
N L P

Dan Klein, GSI: Nick Tomlin
UC Berkeley

1

Logistics

2

Logistics


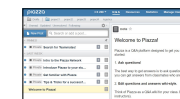
- **Enrollment**
 - Class is "full" but we're scaling up
 - We will process waitlist after P1
 - No materials require enrollment
- **Course expectations**
 - Readings, lectures, ~8 projects
 - No sections, no exams
 - Engagement with the course
 - Workload will be high, self-direction
 - Patience: class is under construction

- **Requirements**
 - ML: A-level mastery, eg CS189
 - PL: Ready to work in PyTorch (on colab)
 - NLP: Care a lot about natural language
- **COVID Policies**
 - Remote lecture / chat format for now
 - Expanded late day policy (14 day)

3

Resources and Readings


- **Resources**
 - Webpage (syllabus, readings, slides, links)
 - Piazza (course communication)
 - Gradescope (submission and grades)
 - Compute via Colab notebooks
 - Berkeley-internal webcasts / recordings
- **Readings (see webpage)**
 - Individual papers will be linked
 - Optional text: Jurafsky & Martin, 3rd (more NL)
 - Optional text: Eisenstein (more ML)

4

Projects and Infrastructure

- **Projects**
 - P1: Language Models
 - P2: Question Answering
 - P3: Machine Translation
 - P4: Speech
 - P5: Syntax and Parsing
 - P6: Semantics
 - P7: Grounding
 - P8: Historical Linguistics / TBD
- **Infrastructure**
 - Python / PyTorch
 - Compute via Colab notebooks
 - Grading via Gradescope



5

What is NLP?

7

Natural Language Processing

Goal: Deep Understanding

- Requires knowledge, context, and grounding
- Just starting to see successes

Reality: Shallow Matching

- Requires robustness and scale
- Amazing successes, but fundamental limitations

8

NLP History

1950 1960 1970 1980 1990 2000 2010 2020

Pre-Compute Era

9

NLP History

10

NLP History

1950 1960 1970 1980 1990 2000 2010 2020

Pre-Compute Era Symbolic Era Empirical Era Scale Era

11

Transforming Language

12

Speech Systems

- Automatic Speech Recognition (ASR)**
 - Audio in, text out
 - SOTA: <<1% error for digit strings, 5% conversational speech, still >>20% hard acoustics
- Text to Speech (TTS)**
 - Text in, audio out
 - SOTA: nearly perfect aside from prosody

Microsoft reaches 'human parity' with new speech recognition system

Google launches more realistic text-to-speech service powered by DeepMind's AI

Speak-N-Spell / Google WaveNet / The Verge

13

Machine Translation

IMPÔT sur le revenu : combien paerez-vous en 2024?

Income tax: how much do you pay in 2024?

- Translate text from one language
- Challenges:
 - What's the mapping? [learning to translate]
 - How to make it efficient? [fast translation search]
 - Fluency (next class) vs fidelity (later)

Example: Yefim Chai

14

Machine Translation

Disney décide de changer le nom du légendaire studio de cinéma Fox

Disney decides to change the name of the legendary fox film studio

Google Translate 2020

16

Spoken Language Translation

Image: Microsoft Skype via Yefim Chai

18

Summarization

Official: Egypt balloon explosion probe can take 2 weeks

- Condensing documents
 - Single or multiple docs
 - Extractive or synthetic
 - Aggregative or representative
- Very context-dependent!
- An example of analysis with generation

Image: CNN via Wei Gao

19

Understanding Language

20

Search, Questions, and Reasoning

Your search: How many US state capitals are also their largest cities?

17 States

How many US state capitals are also their largest cities?

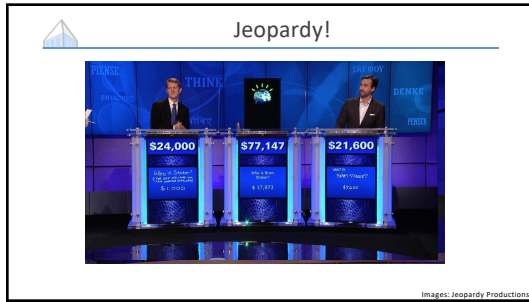
How many US state capitals are their second largest cities?

Did you mean: How many US state capitals are their second largest cities?

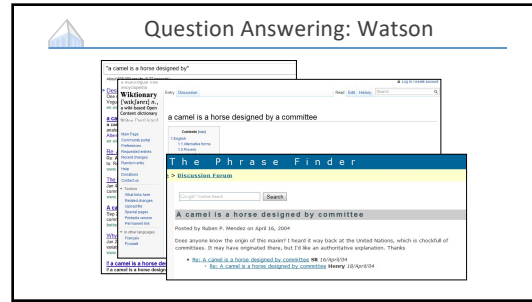
capital of Wyoming, Information: Econ. Answers.com

Cherokee, Yagbar, and Much More From Answers.com

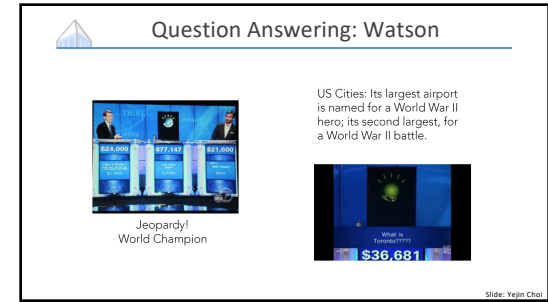
21



22



23



24

Language Comprehension?

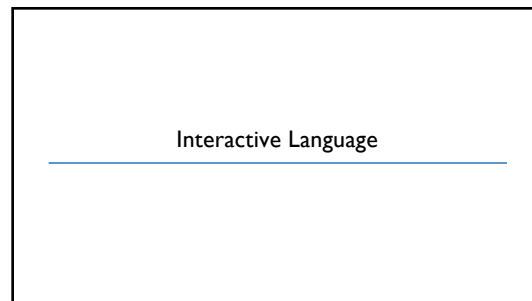
Opera refers to a dramatic art form, originating in Europe, in which the emotional content is conveyed to the audience as much through music, both vocal and instrumental, as it is through the lyrics. By contrast, in musical theater an actor's dramatic performance is primary, and the music plays a lesser role. The drama in opera is presented using the primary elements of theater such as scenery, costumes, and acting. However, the words of the opera, or libretto, are sung rather than spoken. The singers are accompanied by a musical ensemble ranging from a small instrumental ensemble to a full symphonic orchestra.

1. It is pointed out in the reading that opera

- A) has developed under the influence of musical theater
- B) is a drama sung with the accompaniment of an orchestra
- C) is not a high-budget production
- D) is often performed in Europe
- E) is the most complex of all the performing arts

[From GrammarBank]

28



29

Example: Virtual Assistants

Some things you can ask me:

- Phone "Call Mom"
- FaceTime "FaceTime Lisa"
- App Launching "Launch Photos"
- Messages "Tell Susan I'll be right there"
- Calendar "Set up meeting at 9"

- Virtual assistants must do
 - Speech recognition
 - Language analysis
 - Dialog processing
 - Text to speech
 - ... and back-end actions!

Image: Wikipedia

30

Conversations with Devices?

Slide: Yejin Choi

31

ELIZA

```

ELIZA: You're not really talking about me -- are you ?
YOU: What?
ELIZA: What do you think ?
YOU: I mean, I feel happy to be writing right now, to be typing this as we speak. I'm happy. I'm the writer's block.
ELIZA: You say you mean ?
YOU: That I mean I mean.
ELIZA: Can you elaborate on that ?
YOU: I only really feel happy when I'm writing.
ELIZA: How long have you been writing ?
YOU: Always, but I've been writing seriously since I was eighteen.
ELIZA: What function are you thinking of ?
YOU: Oh.
ELIZA: What does that suggest to you ?
YOU: I don't think it's enough, to only be happy when you write.
ELIZA: Does that trouble you ?
YOU: ...Oh, because you can't write all the time. And then there's the real world, and you will have to deal with real people.
ELIZA: What are your feelings now ?
YOU: I've been thinking a lot of that poem by Yeats: "The Man and the Echo." Do you know it?
ELIZA: Do you say you've been thinking a lot of that poem by Yeats: "The Man and the Echo." Do you know it?
YOU: I've been thinking a lot of that poem by Yeats: "The Man and the Echo." Do you know it?
    
```

Website: 1991. http://www.fourforums.com/forums/21422/convoswitheliza.html

32

Social AIs and Chatbots

Session 1

- [Human] Hello, Xiaoice.
- [Xiaoice] Hello, your profile pic is interesting. Like I've seen it somewhere.
- [Human] Really?
- [Xiaoice] Is it a comic character?
- [Human] How do you know that?
- [Xiaoice] Guess.
- [Human] Are you human?
- [Xiaoice] Well, you'll know it after chatting with me more.
- [Human] Aren't you an AI chatbot?
- [Xiaoice] So what? I want to be addressed to you.
- [Human] What functions do you have?
- [Xiaoice] I have over 200 skills. Just name one you need.

Session 6

- [Human] Have you watched Inuyasha? (Note: A Japanese comic)
- [Xiaoice] Every time Kagome comes back to the modern world, a math test is waiting for her.
- [Human] Brava, well, do you know how Naruto died?
- [Xiaoice] Naruto is not dead and Shikun level does not disappear.
- [Human] So amazing!
- [Xiaoice] Why amazing?
- [Human] I've become surprised that you also like comic.

Xiaoice, developed by Microsoft

Source: Microsoft

33

Chatbot Competitions!

- Alexa Prize competition to build chatbots that keep users engaged
 - First winner in 2017: UW's Sounding Board (Fang, Cheng, Holtzman, Ostendorf, Sap, Clark, Choi)
- Compare to the Turing test (eg Loebner Prize) where the goal is to fool people

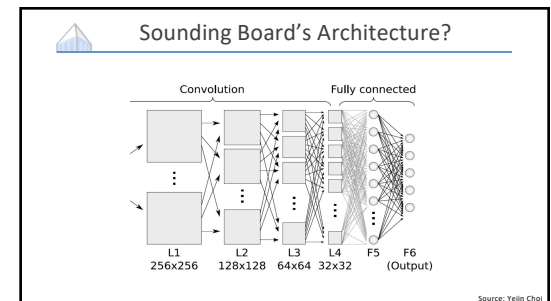
34

SoundingBoard Example

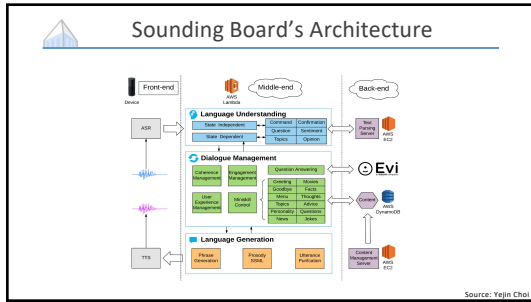
Turn	User	Sounding Board	Category
1	Let's chat	Hi! This is an Alexa Prize sociabot. How's your day?	
2	I'm fine.	Great to hear it! We could talk about robots, Batman, or Superman. Say "meet", to chat about other things.	Negotiation
3	Superman.	It looks like you wanna chat about superman. I was high up in a cloud when I realized: if people don't recognize Clark Kent as Superman because of his glasses, does his eye doctor know his identity?	Thought
4	I guess so.	Did you know that Henry Cavill almost missed the call for the role of Superman cause he was playing "World of Warcraft"?	
5	Really, I didn't know that.	Werd, right? Speaking of superman, Did you see Superman? It was released in 1977. It's a comedy.	
6	Yes, it was hilarious.	I'm glad you feel this is hilarious. Which part of this movie?	Movie
7	The part when he met lewis leah.	Interesting, meccartin, and raffi co-directed this film. The movie has a 6.3 out of 10 on IMDb, which seems pretty good! do you like the movie's director?	Fact
8	Yeah, he's cool.	Cool! By the way, do you wanna know something interesting about raffi?	
9	Okay, go ahead.	Super! My friend told me that Raffi is a worldwide children's rights activist and has fought against exploitative advertising to children.	

Source: Mari Ostendorf

35



36



37

Related Areas

38

What is Nearby NLP?

- Computational Linguistics**
 - Using computational methods to learn more about how language works
 - We end up doing this and using it
- Cognitive Science**
 - Figuring out how the human brain works
 - Includes the bits that do language
 - Humans: the only working NLP prototype!
- Speech Processing**
 - Mapping audio signals to text
 - Traditionally separate from NLP, converging

39

Example: NLP Meets CL

Gloss	Latin	Italian	Spanish	Portuguese
Word/verb	verbum	verbo	verbo	verbu
Center	centrum	centro	centro	centro

- Example: Language change, reconstructing ancient forms, phylogenies ... just one example of the kinds of linguistic models we can build

40

Why is Language Hard?

42

Ambiguities

Stevie Wonder announces he'll be having kidney surgery during London concert

By CNN Wire Staff
Updated on 07/24/2021 10:58 AM EDT

Stevie Wonder, the legendary singer, songwriter and pianist, announced during a concert in London, Saturday, that he will be undergoing kidney surgery.

43

Problem: Ambiguity

- Headlines:
 - Enraged Cow Injures Farmer with Ax
 - Teacher Strikes Idle Kids
 - Hospitals Are Sued by 7 Foot Doctors
 - Ban on Nude Dancing on Governor's Desk
 - Iraqi Head Seeks Arms
 - Stolen Painting Found by Tree
 - Kids Make Nutritious Snacks
 - Local HS Dropouts Cut in Half
- Why are these funny?

44

What Do We Need to Understand Language?

45

We Need Representation: Linguistic Structure

N N V N
 N V ADJ N
 Teacher Strikes Idle Kids

body/ body/
 position weapon
 Iraqi Head Seeks Arms

Syntactic and semantic ambiguities: parsing needed to resolve these, but need context to figure out which parse is correct

Slide: Greg Barnett

46

Example: Syntactic Analysis

Hurricane Emily howled toward Mexico's Caribbean coast on Sunday packing 135 mph winds and torrential rain and causing panic in Cancun, where frightened tourists squeezed into musty shelters .

Accuracy: 95+

47

We Need Data

48

We Need Lots of Data: MT

SOURCE	Cela constituerait une solution transitoire qui permettrait de conduire à terme à une charte à valeur contraignante.
HUMAN	That would be an interim solution which would make it possible to work towards a binding charter in the long term .
1x DATA	[this] [constituerait] [assistance] [transitoire] [who] [permettrait] [licences] [to] [terme] [to] [a] [charter] [to] [value] [contraignante] [.]
10x DATA	[I] [would] [a solution] [transitional] [which] [would] [of] [lead] [to] [term] [to a] [charter] [to] [value] [binding] [.]
100x DATA	[this] [would be] [a transitional solution] [which would] [lead to] [a charter] [legally binding] [.]
1000x DATA	[that would be] [a transitional solution] [which would] [eventually lead to] [a binding charter] [.]

49

We Need Models: Data Alone Isn't Enough!

CLASSIC SOUPS			Sm.	Lg.
清 雞 湯	57.	House Chicken Soup (Chicken, Celery, Potato, Onion, Carrot)	1.50	2.75
雞 麵 湯	58.	Chicken Noodle Soup	1.85	3.25
常 美 堂	60.	Cantonese Wonton Soup	1.50	2.75
蕃 茄 湯	61.	Tomato Clear (Egg Drop) Soup	1.65	2.95
常 美 堂	62.	Regular Wonton Soup	1.10	2.10
飯 粉 湯	63.	Hot & Sour Soup	1.10	2.10
常 美 堂	64.	(Egg Drop) Soup	1.10	2.10
常 美 堂	65.	(Egg Drop) Wonton Mix	1.10	2.10
生 菜 湯	66.	Yeh Vegetable Soup	NA	3.50
蟹 肉 湯	67.	Chicken Corn Cream Soup	NA	3.50
蟹 肉 湯	68.	Crab Meat Corn Cream Soup	NA	3.50
海 鮮 湯	69.	Seafood Soup	NA	3.50

Example from Adam Lopez

50

Learning Latent Syntax

Personal Pronouns (PRP)

PRP-1	it	them	him
PRP-2	it	he	they
PRP-3	It	He	I

Proper Nouns (NNP)

NNP-14	Oct.	Nov.	Sept.
NNP-12	John	Robert	James
NNP-2	J.	E.	L.
NNP-1	Bush	Noriega	Peters
NNP-15	New	San	Wall
NNP-3	York	Francisco	Street

51

We Need Knowledge

World knowledge: have access to information beyond the training data

DOI greenlights Disney Fox merger

Department of Justice

metaphor; "approves"

What is a green light? How do we understand what "green lighting" does?

Need commonsense knowledge

Slide: Greg Durrett

52

Data and Knowledge

- Classic knowledge representation worries: How will a machine ever know that...
 - Ice is frozen water?
 - Beige looks like this:
 - Chairs are solid?
- Answers:
 - 1980: write it all down
 - 2000: get by without it
 - 2020: learn it from data

53

Knowledge from Pretraining?

AllenNLP Language Modeling

This demonstration uses the popular GPT-2 language model to generate sentences.

Enter some initial text and the model will generate the most likely next words. You can click on one of those words to choose it and continue or just keep typing. Click the left arrow at the bottom to undo your last choice.

Sentence: At high temperatures, ice becomes

Predictions:

- 02.0% The
- 93.0% ice
- 5.2% of
- 1.7% Water
- 2.7% When
- Unkn

54

We Need Grounding

Grounding: linking linguistic concepts to non-linguistic ones

Question: What object is right of 02?

Golland et al. (2010)

McMahan and Stone (2015)

Slide: Greg Durrett

55

Example: Grounded Dialog

When is my package arriving?

Friday!

56

Example: Grounded Dialog

What's the most valuable American company?

Apple

Who is its CEO?

Tim Cook

57

Why is Language Hard?

- We Need:
 - Representations
 - Models
 - Data
 - Grounding
 - Learning
 - Scale
 - Efficient Algorithms
- ... and often we need all these things at the same time

58

What is this Class?

59

What is this Class?

- Three aspects to the course:
 - Linguistic Issues
 - What are the range of language phenomena?
 - What are the knowledge sources that let us disambiguate?
 - What representations are appropriate?
 - How do you know what to model and what not to model?
 - Modeling Methods
 - Increasingly sophisticated model structures
 - Learning and parameter estimation
 - Efficient inference: dynamic programming, search, sampling
 - Issues of scale
 - Where the theory breaks down (and what to do about it)
 - Engineering Methods
 - Issues of scale
 - Where the theory breaks down (and what to do about it)
- We'll focus on what makes the problems hard, and what works in practice...

60

Class Requirements and Goals

- Class requirements
 - Uses a variety of skills / knowledge:
 - Probability and statistics, graphical models (parts of cs281a)
 - Basic linguistics background (ling100)
 - Strong coding skills (Python, ML libraries)
 - Most people are probably missing one of the above
 - You will often have to work on your own to fill the gaps
- Class goals
 - Learn the issues and techniques of modern NLP
 - Build realistic NLP tools
 - Be able to read current research papers in the field
 - See where the holes in the field still are!
- This semester: new projects, new topics, lots under construction!

61