


# Natural Language Processing



Dan Klein, GSI: Nick Tomlin  
UC Berkeley

# Logistics

---



# 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)

NL: Care a lot about  
natural language

## ■ COVID Policies

- Remote lecture / chat format for now
- Expanded late day policy (14 day)



# Resources and Readings

## ■ Resources

- Webpage (syllabus, readings, slides, links)
- Piazza (course communication)
- Gradescope (submission and grades)
- Compute via Colab notebooks
- Berkeley-internal webcasts / recordings

CS 288	Lectures
Home	Jan 18: Introduction
<a href="#">Course Info</a>	Jan 20: Language Modeling
<a href="#">Instructors</a>	Jan 25: Neural Language Modeling

## ■ Readings (see webpage)

- Individual papers will be linked
- Optional text: Jurafsky & Martin, 3<sup>rd</sup> (more NL)
- Optional text: Eisenstein (more ML)

The screenshot shows the Piazza Q&A interface for CS 288. The top navigation bar includes 'CS 288', 'Q & A', 'Resources', 'Statistics', and 'Manage Class'. Below the navigation bar, there are tabs for 'Drafts', 'project1', 'project2', 'project3', 'project4', and 'logistics'. The main content area is divided into two sections: 'Pinned' and 'Last Week'. The 'Pinned' section contains a 'New Post' button and a search bar. Below it, there are several pinned posts, including 'Search for Teammates!', 'Intro to the Piazza Network', 'Introduce Piazza to your stu...', 'Get familiar with Piazza', and 'Tips & Tricks for a successf...'. The 'Last Week' section contains a post titled 'Welcome to Piazza!'. The right sidebar shows a 'note' with the text 'Welcome to Piazza!' and a description of the platform: 'Piazza is a Q&A platform designed to get you great started: 1. Ask questions! The best way to get answers is to ask questions! As you can get answers from classmates who are up a 2. Edit questions and answers wiki-style. Think of Piazza as a Q&A wiki for your class. Every instructors.'





# 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

```
CS288 Project 0.ipynb
File Edit View Insert Runtime Tools Help
+ Code + Text
[7] def evaluate(network, data, raw_data):
    predictions = get_predictions(network, data)
    return fscore(predictions, raw_data)

Now we are ready to run everything and get our outputs. A correct implementation should get a validation score around 70%.

trained_network = train()

test_predictions = get_predictions(trained_network, test_inputs)
save_predictions(test_predictions, 'predicted_test_outputs.txt')

...
63.987473903966595
66.7004048582996
67.16791979949875
68.563021088769
```

## ■ Infrastructure

- Python / PyTorch
- Compute via Colab notebooks
- Grading via Gradescope

# What is NLP?

---



# Natural Language Processing



NLP



grep


## 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


# NLP History



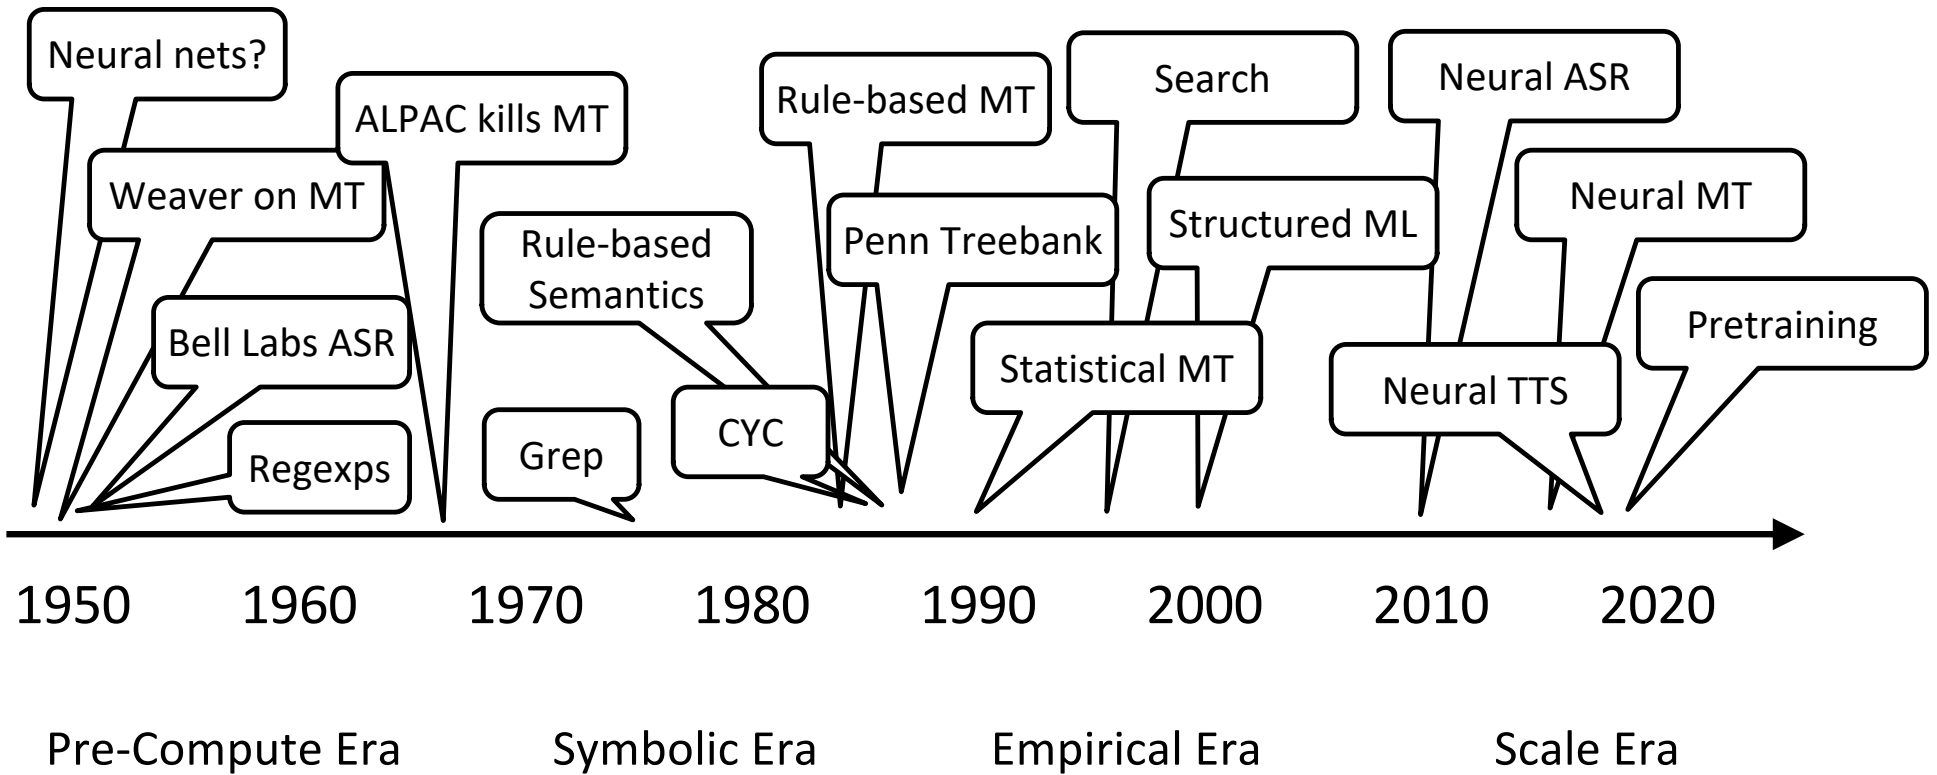
Pre-Compute Era



# NLP History



# NLP History



# Transforming Language


---



# Speech Systems

## ■ Automatic Speech Recognition (ASR)

- Audio in, text out
- SOTA: <<1% error for digit strings, 5% conversational speech, still >>20% hard acoustics



“Speech Lab”

## ■ Text to Speech (TTS)

- Text in, audio out
- SOTA: nearly perfect aside from prosody



MICROSOFT TECH ARTIFICIAL INTELLIGENCE

Microsoft reaches 'human parity' with new speech recognition system

GOOGLE TECH ARTIFICIAL INTELLIGENCE

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



Speak-N-Spell / Google WaveNet / The Verge





# Machine Translation

EN CE MOMENT Impôts Kenya Syrie Pakistan Emploi Scandale Prism

## Impôt sur le revenu : combien paierez-vous en 2014 ?

DÉCOUVREZ TOUS LES SERVICES ABONNÉS  
S'abonner au Monde à partir de 1 €

### APPEL À TÉMOIGNAGES

Adhérent(e) d'Europe Ecology-Les Verts, partagez-vous le constat sévère de Noël Mamère sur EELV ?  
Partagez votre expérience

#### En continu

- 19:53 Budget : des dépenses incompressibles
- 19:36 Entendu à la "fashion week" de Paris
- 19:19 Commande géante à Airbus
- 19:04 Plainte contre "Valeurs actuelles"
- 19:01 Venezuela : 17 personnes arrêtées
- 18:59 Vidberg : le budget nouveau est arrivé
- 18:50 La "mission noble" de la NSA
- 18:38 Roms : Joutes entre Bruxelles et Paris

Sélectionnez votre revenu et votre situation familiale pour découvrir si vous bénéficiez de la pause fiscale.

- Comment le budget pour 2014 est-il réparti ? VISUEL INTERACTIF
- Un budget 2014 soumis aux critiques
- Budget : ces dépenses qu'aucun gouvernement ne peut diminuer
- Budget 2014 : le détail des économies VISUEL INTERACTIF

AT THIS MOMENT Taxes Kenya Syria Pakistan Use Prism scandal

## Income tax: how much do you pay in 2014?

DÉCOUVREZ TOUS LES SERVICES ABONNÉS  
S'abonner au Monde à partir de 1 €

### CALL FOR EVIDENCE

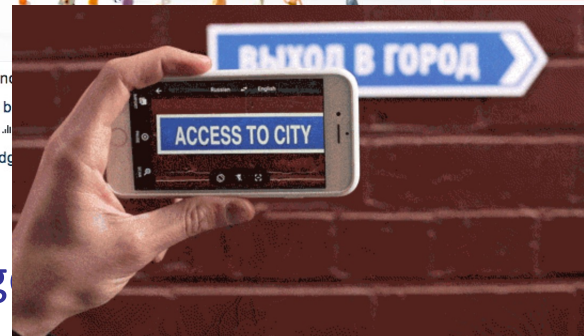
Member (s) of Europe Ecology-Greens, do you share the finding of severe Christmas Mamère EELV?  
Share your experience

#### Continuous

- 19:53 Budget : des dépenses incompressibles
- 19:36 Entendu à la "fashion week" de Paris
- 19:19 Commande géante à Airbus
- 19:04 Plainte contre "Actual Values"
- 19:01 Venezuela : 17 people arrested
- 18:59 Vidberg : the budget came
- 18:50 La "mission noble" of the NSA
- 18:38 Roms : Joutes between Brussels and Paris

Select your income and your family situation to discover if you benefit from the tax break.

- How is the budget allocated? VISUEL INTERACTIF
- A 2014 budget criticism
- Budget : these expenses that no government can reduce
- Budget 2014: the details of savings VISUEL INTERACTIF



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

Example: Yejin Choi



# Machine Translation

CINÉMA · MÉDIAS

Partage 

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

Le nouveau propriétaire va faire disparaître la mention « Fox » à la 20th Century pour que le nom du studio ne soit pas associé à l'image ultra-conservatrice de la chaîne de télévision Fox News.

Le Monde avec AFP · Publié le 18 janvier 2020 à 06h46

 Lecture 1 min.

Révolution dans le monde d'Hollywood. Le célèbre studio de cinéma 20th Century Fox va voir une partie de son nom modifiée et devenir 20th Century Studios, décision prise par son nouveau propriétaire, Disney.

Selon plusieurs médias américains, Disney aurait décidé de modifier le nom du studio pour qu'il ne soit plus associé à la grande chaîne Fox mais aussi et surtout à Fox News, la chaîne d'information en continu.


CINEMA · MEDIA

Sharing 

## Disney decides to change the name of the legendary Fox film studio


The new owner will remove the mention "Fox" in the 20th Century so that the name of the studio is not associated with the ultra-conservative image of the television channel Fox News.

Le Monde avec AFP · Posted on January 18, 2020 at 06:46

 Reading time 1 min.

Revolution in the Hollywood world. The famous film studio 20th Century Fox will see part of its name changed and become 20th Century Studios, decision taken by its new owner, Disney.

According to several American media, Disney has decided to change the name of the studio so that it is no longer associated with the big chain Fox but also and especially with Fox News, the news channel.





# Spoken Language Translation



Image: Microsoft Skype via Yejin Choi



# Summarization

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

The image shows a screenshot of a CNN news article titled "Official: Egypt balloon explosion probe can take 2 weeks". The article is dated February 27, 2013, and is updated at 16:14 GMT (00:14 HKT). The article is written by Adam Makary, Saad Abedine, and Mariano Castillo. The article discusses an investigation into a hot air balloon explosion in Luxor, Egypt, which killed 19 people. The article mentions that the investigation could take two weeks and that the governor has banned all hot air balloon flights until further notice. The article also mentions that the Tuesday accident was the world's deadliest hot air balloon accident in at least 20 years. The article includes a video player with a "Click to play" button and a "2009 balloon crash survivor speaks" video. The article also includes a "How to stay safe in a hot air balloon" video. The article is written in English and includes a link to read a version of the story in Arabic.

**EDITION: INTERNATIONAL** | U.S. | MEXICO | ARABIC  
TV: CNN | CNN en Español  
Set edition preference

Home Video World U.S. Africa Asia Europe Latin America Middle East Business

## Official: Egypt balloon explosion probe can take 2 weeks

By Adam Makary, Saad Abedine and Mariano Castillo, CNN  
February 27, 2013 -- Updated 16:14 GMT (00:14 HKT)

**STORY HIGHLIGHTS**

- **NEW:** No foul play is suspected, official says
- The Tuesday accident was the world's deadliest hot air balloon accident in at least 20 years
- Officials: Passengers in the balloon included 19 foreign tourists
- Luxor province bans all hot air balloon flights until further notice

*Read a version of this story in Arabic.*

**Cairo (CNN)** -- An official investigation into the cause of a balloon accident that killed 19 people in Egypt could take two weeks, the governor of Luxor province said Wednesday.

The Tuesday accident was the world's deadliest hot air balloon accident in at least 20 years.

Preliminary investigations confirmed no foul play was involved when gas canisters aboard the balloon exploded, causing it to plummet about 1,000 feet (300 meters) to the ground, Gov. Ezzat Saad said.

CNN iReport: After tragedy, vacationers recall glorious balloon rides in Egypt

Passengers in the balloon included 19 foreign tourists: nine from Hong Kong, four from Japan, three from Britain, two from France and one from Hungary, officials said.

How safe is hot air ballooning?

An Egyptian pilot and another Egyptian were also aboard, Luxor province spokesman Badawi al-Masri said.

Balloon rides offering panoramic aerial views of the Nile River and the ancient temples of Karnak and Hatshepsut are a popular tourist attraction in Luxor, about a nine-hour drive southeast of Cairo.

"You can see Valley of the Kings in the background bordered by farmland," Pauline Liang of Vancouver, Canada, told CNN's iReport last year. "Below were banana farms, and behind us was the city of Luxor. There was a great contrast between desert landscape, lush farmland and urban development."

Tuesday's crash prompted the governor to ban all hot air balloon flights until further notice.

twitter

- CNN: official investigation into yesterday air balloon accident in Luxor could take 2 weeks
- Governor bans all hot air balloon flights until further notice
- Foul play not suspected in fatal balloon accident
- Official: Egypt balloon explosion probe can take 2 weeks
- Egypt balloon explosion
- An official investigation into the cause of a balloon accident that killed 19 people in Egypt could take two w...
- Egypt: Balloon probe could take weeks
- .....

Click to play  
Tourists killed in hot air balloon blast

Click to play  
2009 balloon crash survivor speaks

Click to play  
How to stay safe in a hot air balloon

Image: CNN via Wei Gao

# Understanding Language

---





# Search, Questions, and Reasoning

Web Images Groups News Froogle Local more »

Google any US states' capitals are also their largest cities? Search

Web

Your search - **How many US states' capitals are also their largest cities?** - did not match any documents.

Suggestions:

- Make sure all words are spelled correctly.
- Try different keywords.
- Try more general keywords.
- Try fewer keywords.

How many US states' capitals are also their largest cities?

About 39,100,000 results (0.84 seconds)

## 17 States

The 17 **States** Whose Capital is also its **Largest City**. For the majority of **US States**, the state capital is NOT the **largest city** in the state. There are a few different reasons for this. In some cases, **capitals** were established decades before the **largest city** was established in another. Jul 9, 2018

quizriot.com › geography › states › largest cities › largest cities

The 17 States Whose Capital is also its Largest City

How many US states' capitals are their second largest cities?

About 23,500,000 results (1.01 seconds)

Did you mean: How many US states' capitals are **there** second largest cities?

www.factmonster.com › us › states › state-capitals-and-largest-cities

### State Capitals and Largest Cities - Fact Monster

State **Capitals** and **Largest Cities**. The following table lists the **capital** and **largest city** of every state in the **United States**. Featured Videos. ADVERTISEMENT.

Missing: second | Must include: second

Web

Tip: Have a question? Try the Google Assistant.

Web definitions

Cheyenne wordnet

Scholarly articles

The Cultural Capital of Wyoming

MSN Encarta

Cheyenne (city in Wyoming) ... In 1869 the city was founded and named Cheyenne.

encarta.msn.com/encyclopedia\_761556413/Cheyenne\_(city\_in\_Wyoming).htm

### capital of Wyoming: Information From Answers.com

Note: click on a word meaning below to see its connections and related words. The noun **capital** of **Wyoming** has one meaning: Meaning #1 : the **capital**.

www.answers.com/topic/capital-of-wyoming - 21k - Cached - Similar pages

### Cheyenne: Weather and Much More From Answers.com

Chey-enne ( shī-ăn ' , -ĕn ' ) The **capital** of **Wyoming**, in the southeast part of the state near the Nebraska and Colorado borders.

www.answers.com/topic/cheyenne-wyoming - 74k - Cached - Similar pages



# Jeopardy!



Images: Jeopardy Productions



# Question Answering: Watson

"a camel is a horse designed by"

About 200,000 results (0.27 seconds)

Log in / create account

Wiktionary  
[ˈwɪkʃənəri] n.,  
a wiki-based Open  
Content dictionary

Wiktionary [ˈwɪkʃənəri]

Entry Discussion Read Edit History Search

a camel is a horse designed by a committee

Contents [hide]

- 1 English
  - 1.1 Alternative forms
  - 1.2 Proverb

The Phrase Finder

Discussion Forum

Google Custom Search Search

**A camel is a horse designed by committee**

Posted by Ruben P. Mendez on April 16, 2004

Does anyone know the origin of this maxim? I heard it way back at the United Nations, which is chockfull of committees. It may have originated there, but I'd like an authoritative explanation. Thanks

- [Re: A camel is a horse designed by committee](#) SR 16/April/04
  - [Re: A camel is a horse designed by committee](#) Henry 18/April/04

If a camel is a horse des  
If a camel is a horse design






# Question Answering: Watson



Jeopardy!  
World Champion

US Cities: Its largest airport is named for a World War II hero; its second largest, for a World War II battle.





# 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]

# Interactive Language

---



# Example: Virtual Assistants



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



# Conversations with Devices?







# 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. Lately, I've had writer's block.
ELIZA: You say you mean ?
YOU:   That's what I said.
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 incident are you thinking of ?
YOU:   Um.
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:   ...Um, because you can't write all the time. And then there's the real world, and you still 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 for some special reason ?
YOU:   Yes
```



# Social AIs and Chatbots

## Session 1

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

## Session 6

- 1 [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
  - 2 [Human] Bravo, well, do you know how Naraku died?  
[Xiaoice] Naraku is not dead and Shikon Jewel does not disappear.
  - 3 [Human] So amazing!  
[Xiaoice] Why amazed?
  - 4 [Human] Awesome! Surprised that you also like comic
- .....




Xiaoice, developed  
by Microsoft

Source: Microsoft



# 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





# SoundingBoard Example

Turn	User	Sounding Board
1	Let's chat	Hi! This is an Alexa Prize socialbot. How's your day?
2	I'm <u>five</u> .	Glad to hear it! We could talk about robots, batman, or superman. Say "next", to chat about other things.
3	Superman.	It looks like you wanna chat about superman. I was high up in the cloud when I realized: If people don't recognize Clark Kent as Superman because of his glasses, does his eye doctor know his identity?
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.	Weird, right? Speaking of superman, Did you see: Superman? It was released in 1997. It's a comedy.
6	Yes, it was hilarious.	I'm glad you feel this is hilarious. Which part is best about this 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?
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.

Negotiation


Thought

Movie

Fact

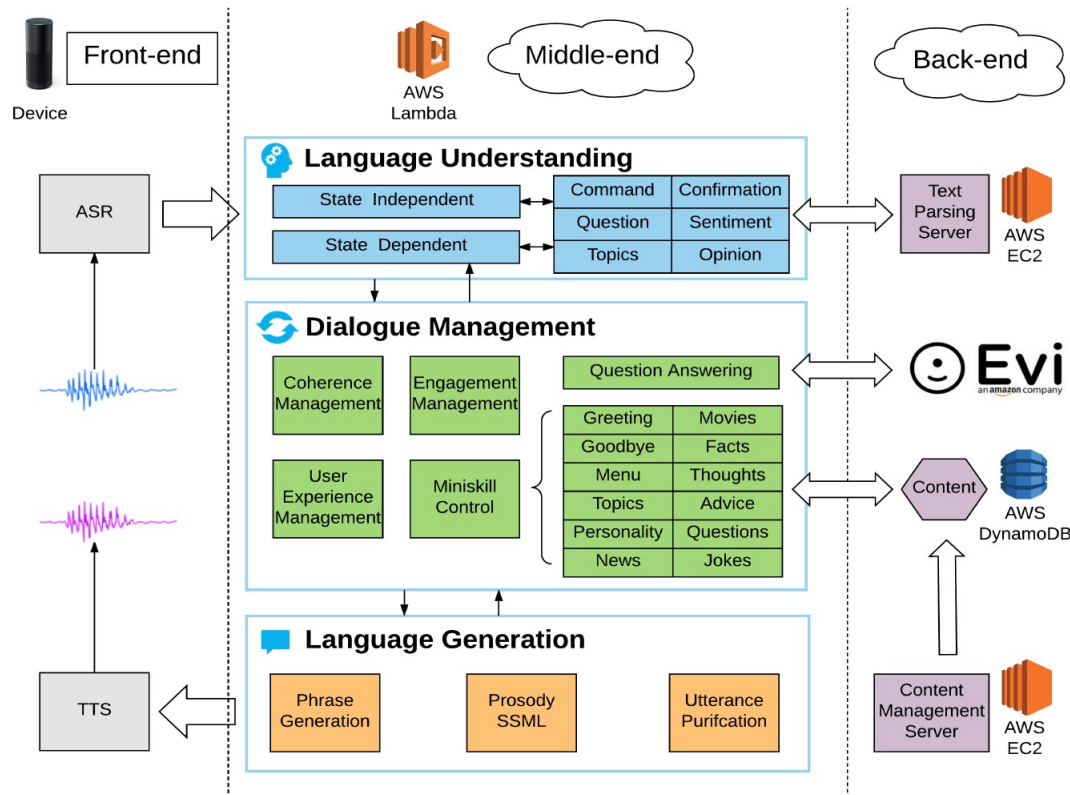


# Sounding Board's Architecture?





# Sounding Board's Architecture



Source: Yejin Choi

# Related Areas


---



# 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!



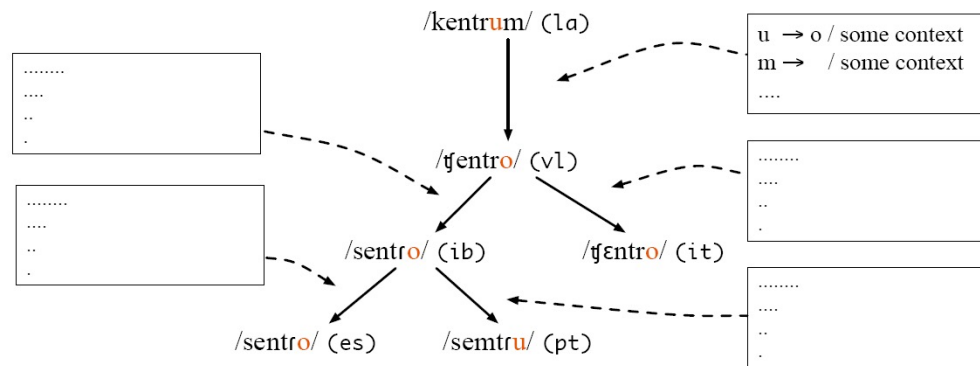
## ■ Speech Processing

- Mapping audio signals to text
- Traditionally separate from NLP, converging





# 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

# Why is Language Hard?

---



# Ambiguities

CNN

Live TV 

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

By Amir Vera, CNN

Updated 11:16 PM EDT, Sat July 06, 2019



(CNN) — [Stevie Wonder](#) will be taking a break from music.

The legendary singer-songwriter announced during [a concert in London](#) Saturday that he will be undergoing kidney surgery.





# 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?

# What Do We Need to Understand Language?

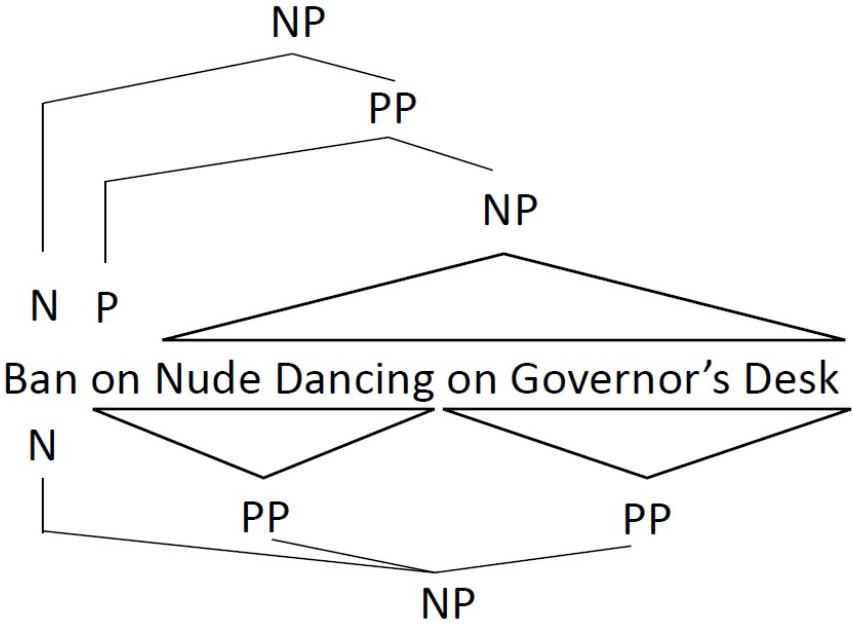
---



# We Need Representation: Linguistic Structure

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


body/ position	body/ weapon
Iraqi Head	Seeks Arms



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



# 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+



# We Need Data





# 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	[it] [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] [.]



# We Need Models: Data Alone Isn't Enough!

			CLASSIC SOUPS		Sm.	Lg.		
清	燉	雞	湯	57.	House Chicken Soup (Chicken, Celery, Potato, Onion, Carrot) .....	1.50	2.75	
雞	飯	湯	58.	Chicken Rice Soup .....	1.85	3.25		
雞	麵	湯	59.	Chicken Noodle Soup .....	1.85	3.25		
廣	東	雲	吞	60.	Cantonese Wonton Soup.....	1.50	2.75	
蕃	茄	蛋	湯	61.	Tomato Clear Egg Drop Soup .....	1.65	2.95	
雲	吞	湯	62.	Regular Wonton Soup .....	1.10	2.10		
酸	辣	湯	63.	Hot & Sour Soup .....	1.10	2.10		
蛋	花	湯	64.	Egg Drop Soup.....	1.10	2.10		
雲	蛋	湯	65.	Egg Drop Wonton Mix.....	1.10	2.10		
豆	腐	菜	湯	66.	Tofu Vegetable Soup .....	NA	3.50	
雞	玉	米	湯	67.	Chicken Corn Cream Soup .....	NA	3.50	
蟹	肉	玉	米	湯	68.	Crab Meat Corn Cream Soup.....	NA	3.50
海	鮮	湯	69.	Seafood Soup.....	NA	3.50		

*Example from Adam Lopez*



# Learning Latent Syntax

---

## Personal Pronouns (PRP)

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

## Proper Nouns (NNP)

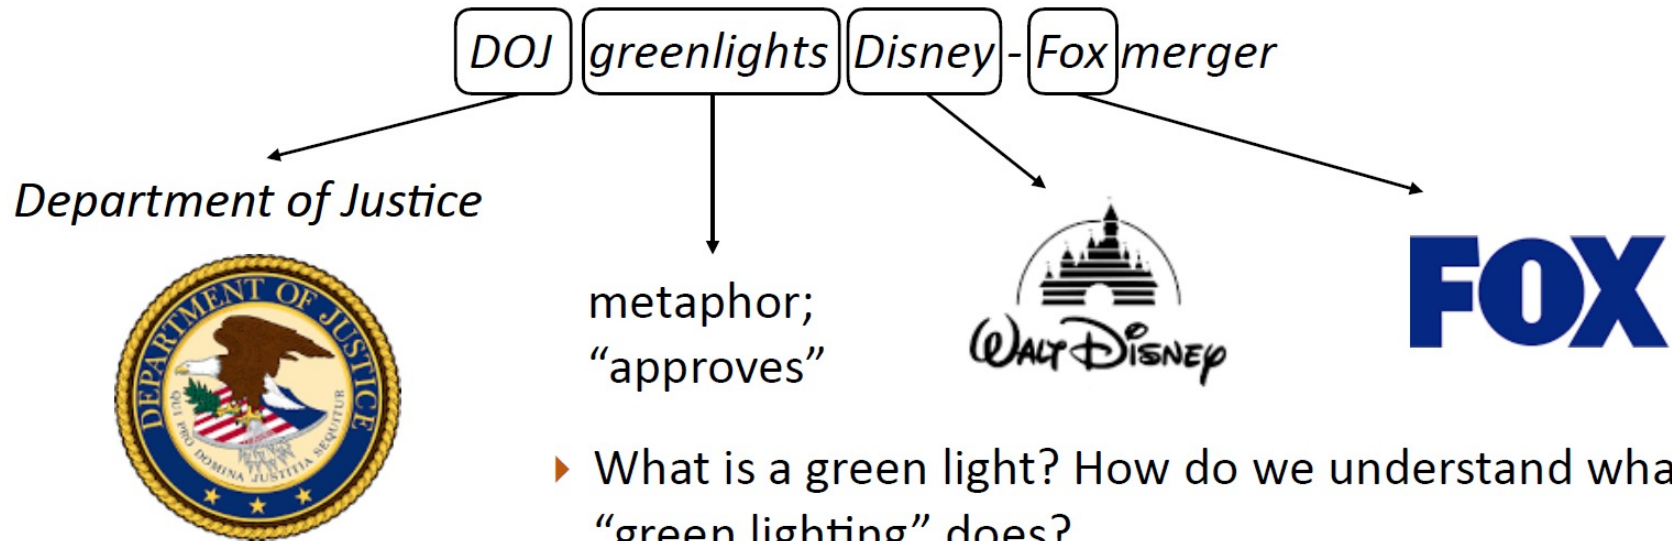
NNP-14	Oct.	Nov.	Sept.
NNP-12	John	Robert	James
NNP-2	J.	E.	L.
NNP-1	Bush	Noriega	Peters
NNP-15	New	San	Wall
NNP-3	York	Francisco	Street





# We Need Knowledge

- ▶ World knowledge: have access to information beyond the training data




- ▶ What is a green light? How do we understand what “green lighting” does?
- ▶ Need commonsense knowledge



# 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



# Knowledge from Pretraining?

AI2 Allen Institute for AI

## AllenNLP

- ∠ Annotate a sentence ^
  - Semantic Role Labeling
  - Named Entity Recognition
  - Constituency Parsing
  - Dependency Parsing
  - Open Information Extraction
  - Sentiment Analysis
- ☰ Annotate a passage ^
  - Coreference Resolution
- 🔍 Answer a question ^
  - Reading Comprehension
- ⚙ Semantic parsing ^

## Language Modeling

This demonstration uses the public 345M parameter [OpenAI 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:


- 32.0% **The**
- 10.3% **Ice**
- 5.2% ↵
- 3.7% **Water**
- 2.7% **When**
- ← Undo




# We Need Grounding

Grounding: linking linguistic concepts to non-linguistic ones

Question: What object is right of O2 ?



Golland et al. (2010)



McMahan and Stone (2015)



# Example: Grounded Dialog

---

When is my package arriving?

Friday!





## Example: Grounded Dialog

What's the most valuable American company?

Apple

Who is its CEO?

Tim Cook





# 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

**What is this Class?**

---





# 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
  - 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...



# 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!**