

CSEP 517
Natural Language Processing
Autumn 2018

Introduction
Luke Zettlemoyer

Slides adapted from Dan Klein, Yejin Choi

What is NLP?



- Fundamental goal: *deep understand of broad language*
 - Not just string processing or keyword matching
- End systems that we want to build:
 - Simple: spelling correction, text categorization...
 - Complex: speech recognition, machine translation, information extraction, sentiment analysis, question answering...
 - Unknown: human-level comprehension (is this just NLP?)

Why NLP

- To access information & knowledge

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.



Knowledge Graph: "things not strings"

Home Tips & Tricks **Features** Search Stories Playground Blog Help

The Knowledge Graph

Learn more about one of the key breakthroughs behind the future of search.

See it in action

Discover answers to questions you never thought to ask, and explore collections and lists.



Information Extraction

- From unstructured text to database entries

New York Times Co. named Russell T. Lewis, 45, president and general manager of its flagship New York Times newspaper, responsible for all business-side activities. He was executive vice president and deputy general manager. He succeeds Lance R. Primis, who in September was named president and chief operating officer of the parent.

Person	Company	Post	State
Russell T. Lewis	New York Times newspaper	president and general manager	started
Russell T. Lewis	New York Times newspaper	executive vice president	ended
Lance R. Primis	New York Times Co.	president and CEO	started

Information Extraction

New York Times Co. named Russell T. Lewis, 45, president and general manager of its flagship New York Times newspaper, responsible for all business-side activities. He was executive vice president and deputy general manager. He succeeds Lance R. Primis, who in September was named president and chief operating officer of the parent.

Person	Company	Post	State
Russell T. Lewis	New York Times newspaper	president and general manager	started
Russell T. Lewis	New York Times newspaper	executive vice president	ended
Lance R. Primis	New York Times Co.	president and CEO	started

Sub-problems:

- 1) Named entity recognition: finding named entities X and their types T(X)
persons: "Russell T. Lewis", "Lance R. Primis"
companies: "New York Times Newspaper", "New York Times Co."
 - 2) Relation extraction: the relation R(X,Y) between named entities X, Y
Works_for(Russell T. Lewis, New York Times Newspaper)
 - 3) Coreference resolution: which text spans refer to the same named entity?
{Russell T.Lewis, He, He} are an equivalence set.
- Is this easy or hard?
 - Easier if the model exploits the redundancy of information!

Question Answering

- Question Answering:
 - More than search
 - Can be really easy: "What's the capital of Wyoming?"
 - Can be harder: "How many US states' capitals are also their largest cities?"
 - Can be open ended: "What are the main issues in the global warming debate?"

- Natural Language Interaction:
 - Understand requests and act on them
 - "Make me a reservation for two at Quinn's tonight"

The screenshot shows the Google search interface. At the top, there are navigation links for Web, Images, Groups, News, Froogle, Local, and more. The search bar contains the query "any US states' capitals are also their largest cities?". Below the search bar, the word "Web" is highlighted. The search results section shows "Your search - **How many US states' capitals are also their largest cities?** - did not match any documents." Below this, there is a "Suggestions:" section with the following list:

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

At the bottom of the search results, there are links for "Google Home", "Business Solutions", and "About Google".

[capital of Wyoming: Information From Answers.com](#)

Note: click on a word meaning below to see its connections and related words.

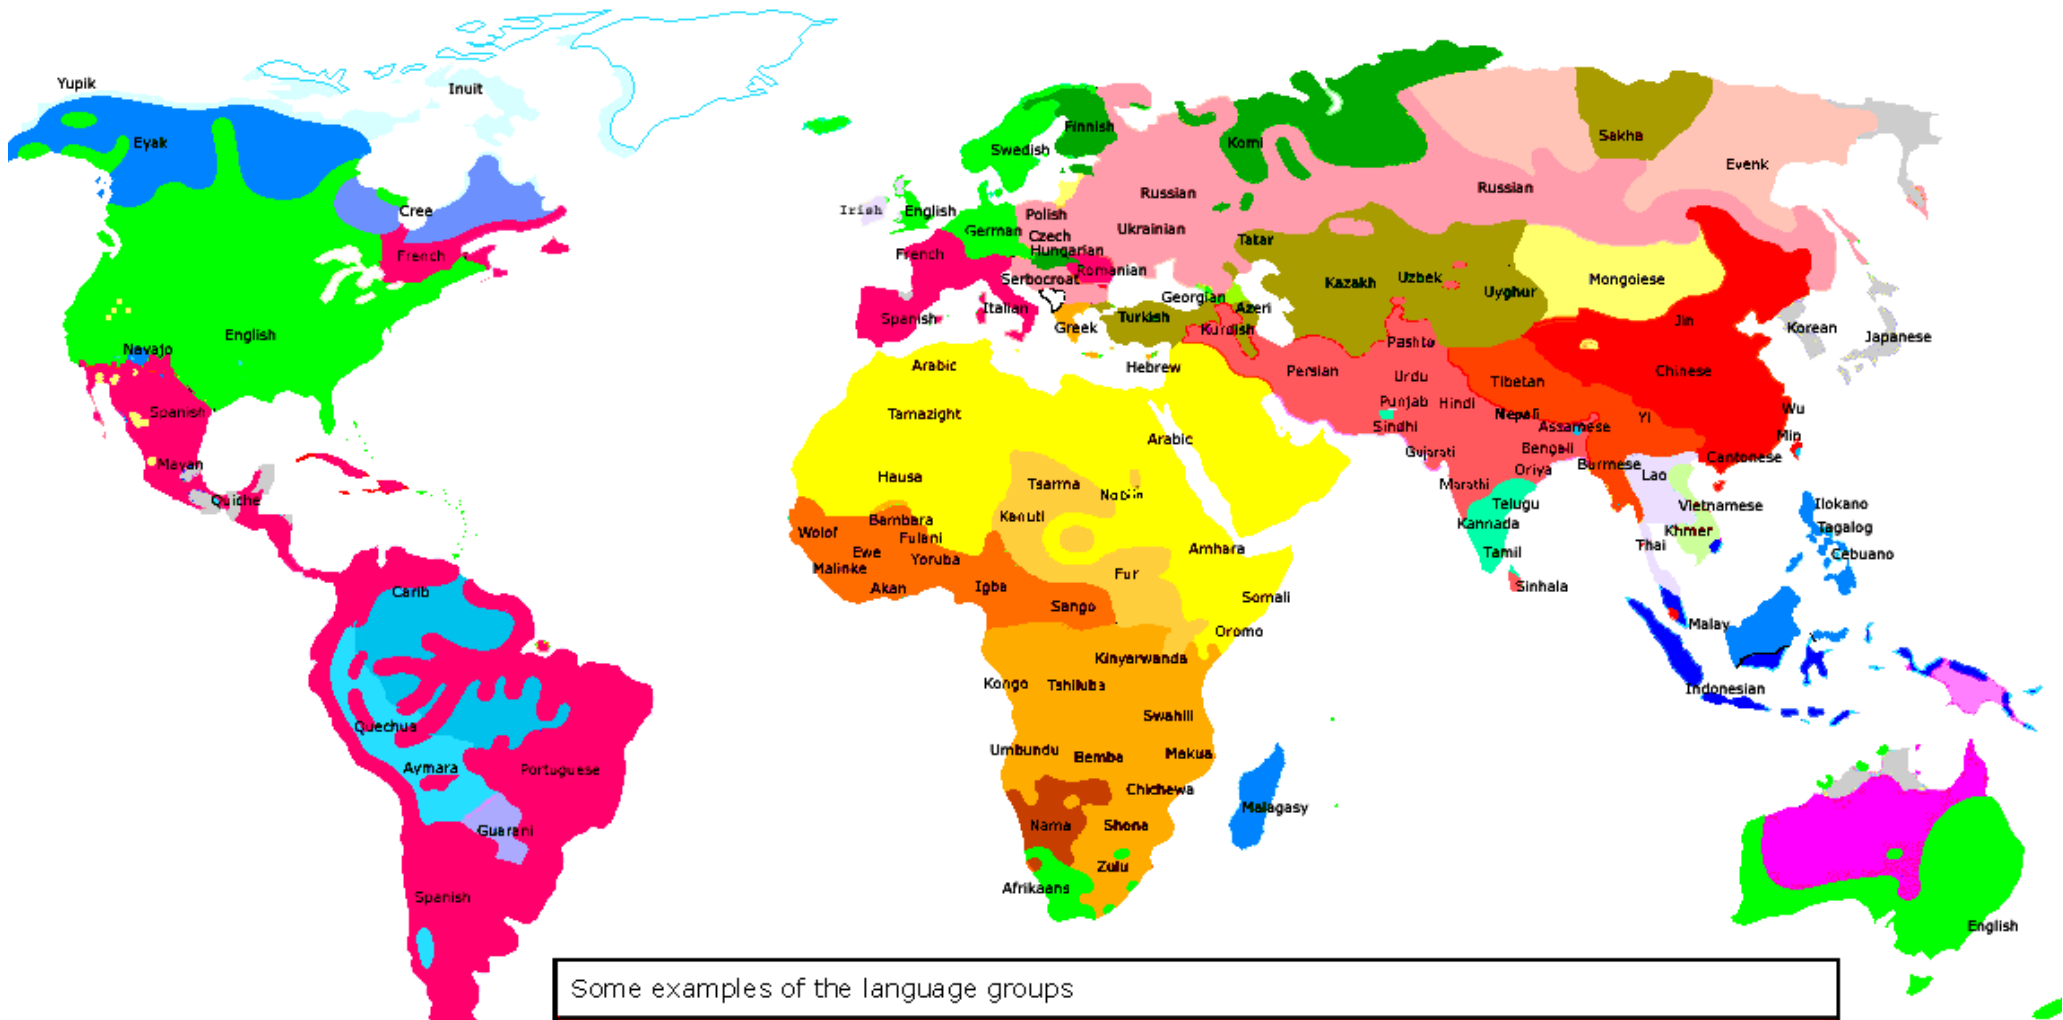
The noun **capital** of **Wyoming** has one meaning: Meaning #1 : the **capital**.

[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](#)



Some examples of the language groups

<ul style="list-style-type: none"> ■ Afro-Asiatic ■ Niger-Congo ■ Bantu ■ Nilo-Saharan ■ Khoisan ■ Indo-European ■ Germanic ■ Albanic ■ Romance ■ Slavic ■ Indo-Iranian ■ Baltic ■ Caucasian 	<ul style="list-style-type: none"> ■ Altaic ■ Turkic ■ Mongolic ■ East Siberian languages ■ Uralic ■ Dravidian ■ Sino-Tibetan ■ Chinese ■ Burmese-Tibetan 	<ul style="list-style-type: none"> ■ Austro-Asiatic ■ Austronesian ■ Borneo-Philippines/Formosan ■ Nuclear Malayo-Polynesian ■ Papuan ■ Pama-Ngyungan ■ Tai-Kadai ■ Isolate 	<ul style="list-style-type: none"> ■ Na-Déne ■ Eskimo-Aleut ■ American Indian ■ Algonic ■ Uto-Aztecan ■ Mayan ■ Andean ■ Tupian ■ Brazilian indigenous
---	--	---	---

Machine Translation

"Il est impossible aux journalistes de rentrer dans les régions tibétaines"

Bruno Philip, correspondant du "Monde" en Chine, estime que les journalistes de l'AFP qui ont été expulsés de la province tibétaine du Qinghai "n'étaient pas dans l'illégalité".

Les faits Le dalaï-lama dénonce l'"enfer" imposé au Tibet depuis sa fuite, en 1959

Vidéo Anniversaire de la rébellion tibétaine : La Chine sur ses gardes



"It is impossible for journalists to enter Tibetan areas"

Philip Bruno, correspondent for "World" in China, said that journalists of the AFP who have been deported from the Tibetan province of Qinghai "were not illegal."

Facts The Dalai Lama denounces the "hell" imposed since he fled Tibet in 1959

Video Anniversary of the Tibetan rebellion: China on guard



- Translate text from one language to another
- Recombines fragments of example translations
- Challenges:
 - What fragments? [learning to translate]
 - How to make efficient? [fast translation search]
 - Fluency (second half of this class) vs fidelity (later)

2013 Google Translate: French

EN CE MOMENT Impôts Kenya Syrie Pakistan Emploi Scandale Prism

Impôt sur le revenu : vous en 2014 ?



Sélectionnez votre revenu et votre situation et vous bénéficiez de la pause fiscale.

- Comment le budget pour 2014 est-il réparti ? [VISUEL INTERACTIF](#)
- Un budget 2014 soumis aux critiques



Le chômage baisse pour la première fois depuis avril 2011 [POST DE BLOG](#)

AT THIS MOMENT Taxes Kenya Syria Pakistan Use Prism scandal

Income tax: how much do you pay in 2014?



Select your income and family situation to see if you get the tax break.

- How is the budget for 2014 is allocated? [INTERACTIVE VISUAL](#)
- Budget: these expenses no government can reduce
- A 2014 budget submitted to criticism
- Budget 2014: the retail savings [INTERACTIVE VISUAL](#)



Unemployment fell for the first time since April 2011 [POST BLOG](#)



Surviving in the Central time looting and anarchy

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

- 7:53 Budget: the fixed expenses
- 7:36 Heard the "Fashion Week" in Paris
- 7:19 control giant Airbus
- 7:04 Complaint against "Actual Values"
- 7:01 Venezuela: 17 people arrested
- 6:59 Vidberg: the new budget came
- 6:50 The "noble mission" of the NSA
- 6:38 Roma: jousting between Brussels e

**DE
FURSAC**

automne-hiver 13/14

2013 Google Translate: Russian



Поиск

Например: [Большой Кавказ](#)

[Мир](#) | [Наука](#) | [Общество](#) | [Здоровье](#) | [Красота](#) |

■ [Новости](#)

20:09

[В Шри-Ланке хотели перевезти золото в желудках](#)

20:00

[Выходец из России может получить "Нобеля" по химии](#)

19:46

[В США установили стандарты торговли оружием](#)

19:35

[Директор Эрмитажа: Обыски нанесли ущерб музею](#)

19:25

[Мозгу ребенка полезен послеобеденный сон](#)

19:24

[Роликом с водителями-детьми заинтересовалась петербургская полиция](#)

19:15

[К Марсу приближается "комета века"](#)

18:55

[Выявлено более 160 нарушений на судостроительных предприятиях](#)

18:44

[Астахов назначен на новый срок в Европейской сети детских омбудсменов](#)

■ [Главное](#)

["Обиженные люди работают, а иностранцы к нам не поедут"](#)

25.09.2013 19:48



Ректор "Бауманки" Анатолий "Правде.Ру", какие шаги над чиновникам и ученым в связи реформе РАН.

[Фотосессия](#)



[Наводнение в Индии: 40 жителей эвакуированы](#)

Найроби. Газета The Independent "Уэстгейт" во время захвата.

■ [Мир](#)

[Израиль не планирует](#)



Поиск

For example, [the Greater Caucasus](#)

[World](#) | [Science](#) | [Society](#) | [Health](#) | [Beauty](#) | [Regions](#) | [Photo](#) | [Video](#) | [Forums](#) | [archive](#)

■ [News](#)

20:09

[In Sri Lanka, wanted to carry the gold in the stomachs](#)

20:00

[A native of Russia can get the "Nobel" in Chemistry](#)

19:46

[In the United States set the standard arms trade](#)

19:35

[Director of the Hermitage: The searches have damaged the museum](#)

19:25

[The child's brain is useful afternoon nap](#)

19:24

[The roller with the drivers, children become interested in the St. Petersburg Police](#)

19:15

[To Mars is approaching "comet of the century"](#)

18:55

[There are over 160 violations at shipyards](#)

18:44

[Astakhov appointed for a new term in the European Network of Ombudsmen for children](#)

18:24

■ [Point](#)

["Mentally ill people are working, and foreign scholars to us will not go"](#)

25/09/2013 19:48




The Rector, "Bauman" Anatoly Alexandrov told with "Pravda.Ru" what steps need to be taken to officials and scientists in connection with the adoption of the law on the reform of the RAS.

[Photoshoot](#)



[World through the lens: September 25.](#)

2466 photos



In Kenya - mourning for the victims of the terrorist attack in Nairobi. The newspaper The Independent said about the people who were at the mall, "Westgate" during capture.

[World](#)

[Belief](#)

[Economy](#)



Enjoy over 1,400 channels



Enjoy over 1,400 channels of music, movies and more on "ice"

Hello Tomorrow



Why NLP

- To access information & knowledge
- To communicate

Human-Machine Interactions



Will this Be Part of All Our Home Devices?



FROM: AMAZON.COM

Will it rain tomorrow?

Set an alarm for eight a.m.

*Play music by
Bruno Mars*

*How many teaspoons
are in a tablespoon?*

*Add gelato to my
shopping list*

*Wikipedia: Abraham
Lincoln*

*When is
Thanksgiving?*

*Play my "dinner party"
playlist*

*What's the weather in
Los Angeles this weekend?*

*Add "make hotel reservations"
to my to-do list*



[< PREV TEAM](#) | [VIEW ALL](#)

University of Washington Sounding Board



Sounding Board




Location: Seattle, WA, USA
Faculty Advisor: [Mari Ostendorf](#)

Speech Recognition

- Automatic Speech Recognition (ASR)

- Audio in, text out
- SOTA: 0.3% error for digit strings, 5% dictation, 50%+ TV



“Speech
Lab”

- Text to Speech (TTS)

- Text in, audio out
- SOTA: totally intelligible (if sometimes unnatural)

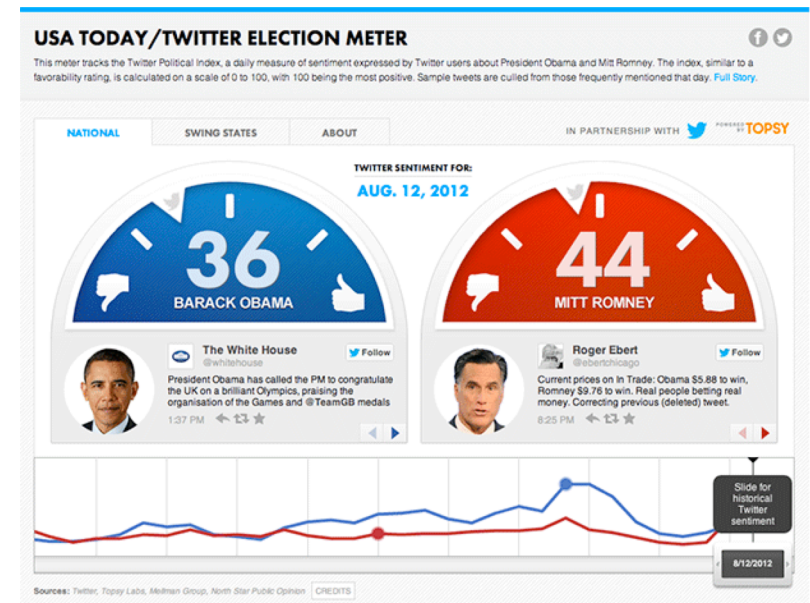


Why NLP

- To access information & knowledge
- To communicate
- To understand our society

Analyzing public opinion, making political forecasts

- Today: In 2012 election, automatic sentiment analysis actually being used to complement traditional methods (surveys, focus groups)
- Past: "Sentiment Analysis" research started in 2002
- Future: **computational social science** and NLP for digital humanities (psychology, communication, literature and more)
- Challenge: Need statistical models for deeper semantic understanding --- subtext, intent, nuanced messages



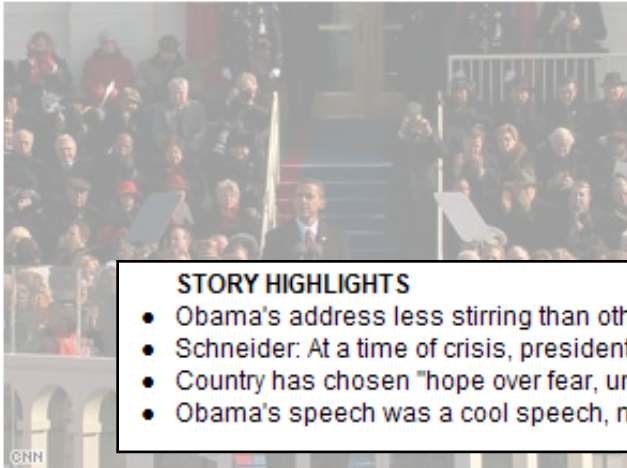
Why NLP

- To access information & knowledge
- To communicate
- To understand our society
- And to make our lives easier

Summarization

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

WASHINGTON (CNN) -- President Obama's inaugural address was cooler, more measured and reassuring than that of other presidents making it, perhaps, the right speech for the times.



Some inaugural addresses are known for their soaring, inspirational language. Like John F. Kennedy's in 1961: "Ask not what your country can do for you. Ask what you can do for your country."

Obama's address was less stirring, perhaps, but it was also more candid and down-to-earth.

"Starting today," the new president said, "we must begin

STORY HIGHLIGHTS

- Obama's address less stirring than others but more candid, analyst says
- Schneider: At a time of crisis, president must be reassuring
- Country has chosen "hope over fear, unity of purpose over ... discord," Obama said
- Obama's speech was a cool speech, not a hot one, Schneider says

President Obama renewed his call for a massive plan to stimulate economic growth.

[more photos >](#)

his first inaugural in 1933, "The only thing we have to fear is fear itself." Or Bill Clinton, who took office during the economic crisis of the early 1990s. "There is nothing wrong with America that cannot be fixed by what is right with America," Clinton declared at his first inaugural.

[Obama](#), too, offered reassurance.

"We gather because we have chosen hope over fear, unity of purpose over conflict and discord," Obama said.

Obama's call to unity after decades of political division echoed Abraham Lincoln's first inaugural address in 1861. Even though he delivered it at the onset of a terrible civil war, Lincoln's speech was not a call to battle. It was a call to look beyond the war, toward reconciliation based on what he called "the better angels of our nature."

Some presidents used their [inaugural address](#) to set out a bold agenda.

Start-up Summly → Yahoo!

CEO Marissa Mayer announced an update to the app in a blog post, saying, "The new Yahoo! mobile app is also smarter, using Summly's natural-language algorithms and machine learning to deliver quick story summaries. We acquired Summly less than a month ago, and we're thrilled to introduce this game-changing technology in our first mobile application."



Launched 2011, Acquired 2013 for \$30M

Can a robot write news?

Despite an expected dip in profit, analysts are generally optimistic about **Steelcase** as it prepares to reports its third-quarter earnings on Monday, December 22, 2014. The consensus earnings per share estimate is 26 cents per share.

The consensus estimate remains unchanged over the past month, but it has decreased from three months ago when it was 27 cents. Analysts are expecting earnings of 85 cents per share for the fiscal year. Revenue is projected to be 5% above the year-earlier total of \$784.8 million at \$826.1 million for the quarter. For the year, revenue is projected to come in at \$3.11 billion.

The company has seen revenue grow for three quarters straight. The less than a percent revenue increase brought the figure up to \$786.7 million in the most recent quarter. Looking back further, revenue increased 8% in the first quarter from the year earlier and 8% in the fourth quarter.

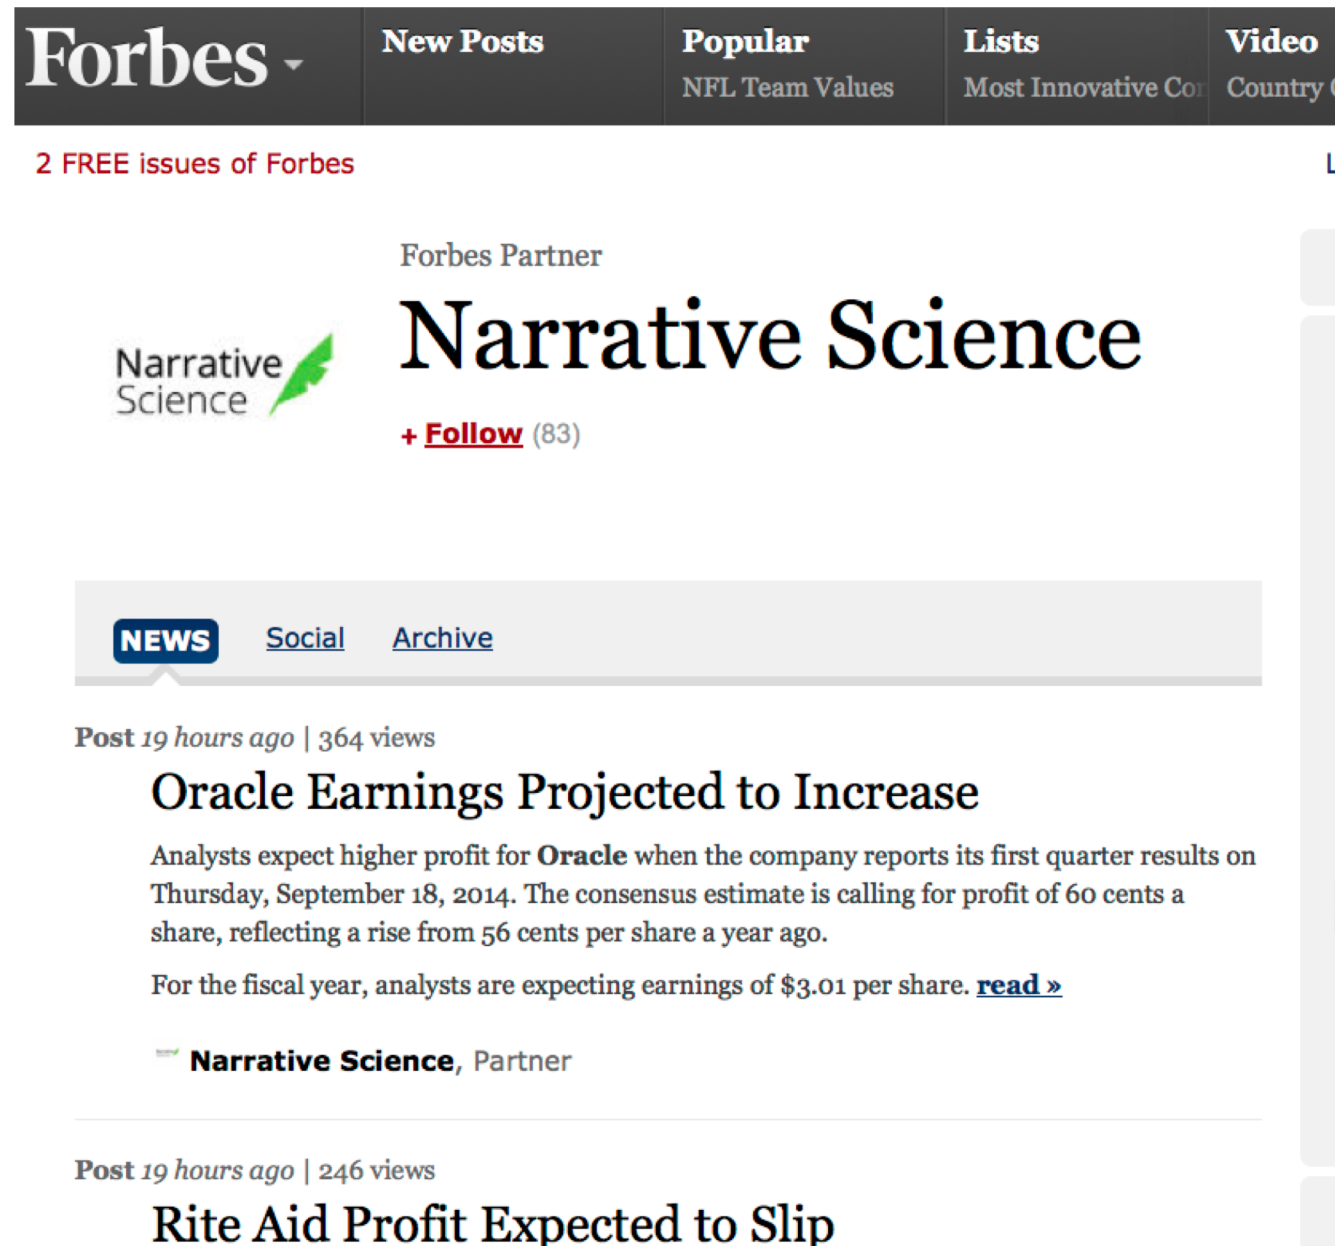
The majority of analysts (100%) rate Steelcase as a buy. This compares favorably to the analyst ratings of three similar companies, which average 57% buys. Both analysts rate Steelcase as a buy.

Steelcase is a designer, marketer and manufacturer of office furniture. Other companies in the furniture and fixtures industry with upcoming earnings release dates include: HNI and Knoll.

Writer-bots for earthquake & financial reports

Some of the formulaic news articles are now written by computers.

- Definitely far from "Op-ed"
- Can we make the generation engine statistically learned rather than engineered?



The screenshot shows the Forbes website interface. At the top, there is a navigation bar with the Forbes logo and several menu items: "New Posts", "Popular" (with a sub-item "NFL Team Values"), "Lists" (with a sub-item "Most Innovative Co"), and "Video" (with a sub-item "Country"). Below the navigation bar, there is a promotional banner for "2 FREE issues of Forbes". The main content area features a post by "Narrative Science", which is identified as a "Forbes Partner". The post title is "Oracle Earnings Projected to Increase". The post text states: "Analysts expect higher profit for **Oracle** when the company reports its first quarter results on Thursday, September 18, 2014. The consensus estimate is calling for profit of 60 cents a share, reflecting a rise from 56 cents per share a year ago. For the fiscal year, analysts are expecting earnings of \$3.01 per share. [read »](#)". The post is attributed to "Narrative Science, Partner". Below this post, there is another post snippet with the title "Rite Aid Profit Expected to Slip".

Forbes

New Posts Popular NFL Team Values Lists Most Innovative Co Video Country

2 FREE issues of Forbes

Forbes Partner

Narrative Science

+ Follow (83)

NEWS Social Archive

Post 19 hours ago | 364 views

Oracle Earnings Projected to Increase

Analysts expect higher profit for **Oracle** when the company reports its first quarter results on Thursday, September 18, 2014. The consensus estimate is calling for profit of 60 cents a share, reflecting a rise from 56 cents per share a year ago.

For the fiscal year, analysts are expecting earnings of \$3.01 per share. [read »](#)

Narrative Science, Partner

Post 19 hours ago | 246 views

Rite Aid Profit Expected to Slip

Why NLP

- To access information & knowledge
- To communicate
- To understand our society
- To make our lives easier
- NLP and AI


Language Comprehension?

"The rock was still wet. The animal was glistening, like it was still swimming," recalls Hou Xianguang. Hou discovered the unusual fossil while surveying rocks as a paleontology graduate student in 1984, near the Chinese town of Chengjiang. "My teachers always talked about the Burgess Shale animals. It looked like one of them. My hands began to shake." Hou had indeed found a *Naraoia* like those from Canada. However, Hou's animal was 15 million years older than its Canadian relatives.

It can be inferred that Hou Xianguang's "hands began to shake", because he was:



- (A) afraid that he might lose the fossil
- (B) worried about the implications of his finding
- (C) concerned that he might not get credit for his work
- (D) uncertain about the authenticity of the fossil
- (E) excited about the magnitude of his discovery

Language and Vision




"Imagine, for example, a computer that could look at an arbitrary scene anything from a sunset over a fishing village to Grand Central Station at rush hour and produce a verbal description. This is a problem of overwhelming difficulty, relying as it does on finding solutions to both vision and language and then integrating them. I suspect that scene analysis will be one of the last cognitive tasks to be performed well by computers"

-- David Stork (HAL's Legacy, 2001) on A. Rosenfeld's vision



What begins to work (e.g., Kuznetsova et al. 2014)



The flower was so
vivid and attractive.



Blue flowers are **running rampant** in my garden.

We sometimes do well: 1 out of 4 times, machine captions were preferred over the original Flickr captions:



Spring in a white dress.

Blue flowers have no scent. Small white flowers have no idea what they are.



Scenes around the lake on my bike ride.



This horse walking along the road as we drove by.



Table of Content

- Definition of NLP
- Historical account of NLP

NLP History: pre-statistics

(1) Colorless green ideas sleep furiously.

(2) Furiously sleep ideas green colorless.

- It is fair to assume that neither sentence (1) nor (2) (nor indeed any part of these sentences) had ever occurred in an English discourse. Hence, in any statistical model for grammaticality, these sentences will be ruled out on identical grounds as equally "remote" from English. Yet (1), though nonsensical, is grammatical, while (2) is not." (Chomsky 1957)
- 70s and 80s: more linguistic focus
 - Emphasis on deeper models, syntax and semantics
 - Toy domains / manually engineered systems
 - Weak empirical evaluation

NLP: machine learning and empiricism

“Whenever I fire a linguist our system performance improves.” –Jelinek, 1988

- 1990s: Empirical Revolution
 - Corpus-based methods produce the first widely used tools
 - Deep linguistic analysis often traded for robust approximations
 - *Empirical evaluation* is essential
- 2000s: Richer linguistic representations used in statistical approaches, scale to more data!

NLP: deep learning / neural networks

“The idea of what an internal representation would look like was it would be some kind of symbolic structure. That has completely changed with these big neural nets.”


–Hinton, 2016

- ~2014-now: Neural networks
 - Big models, more data, less and less linguistic bias
 - Can be brittle to adversarial inputs
 - Can be difficult to interpret
- 2020s: What comes next?
 - Hybrid models? Just deeper networks?
 - You decide!!!

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?

- Mapping audio signals to text
- Traditionally separate from NLP, converging?
- Two components: acoustic models and language models
- Language models in the domain of stat NLP




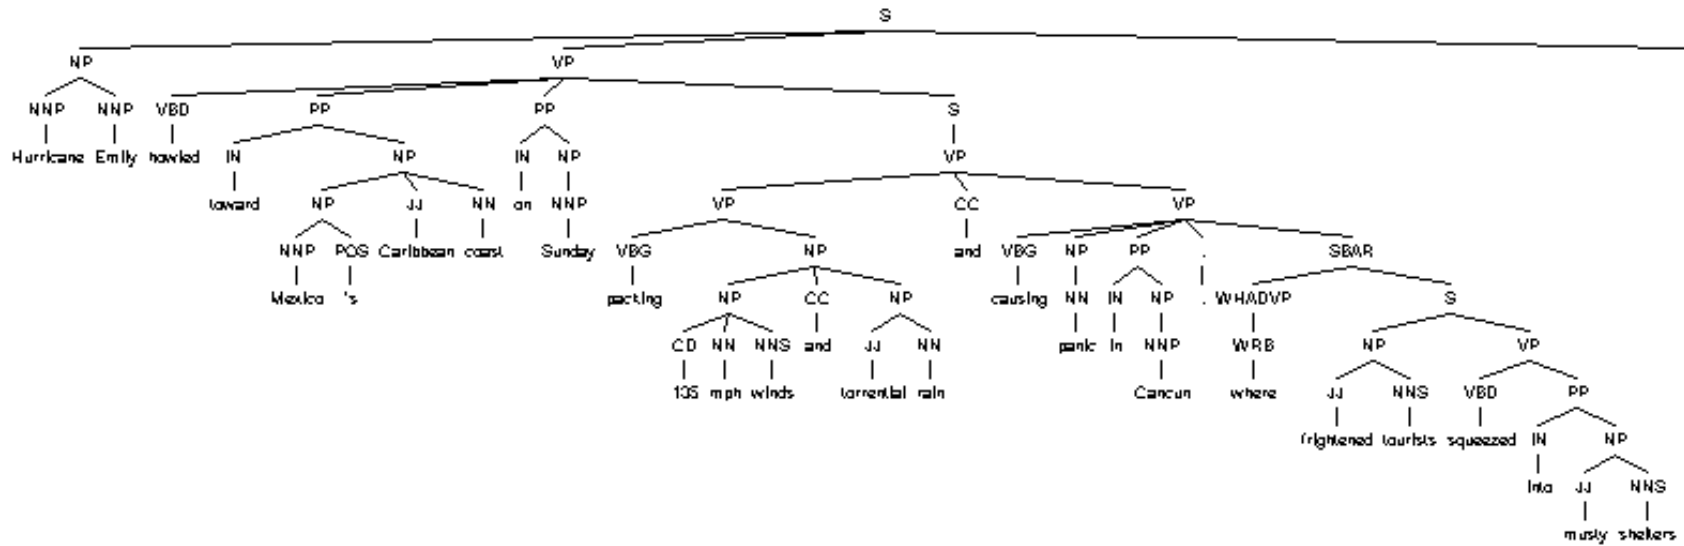
Table of Content

- Definition of NLP
- Historical account of NLP
- Unique challenges of NLP

Problem: Ambiguities

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

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 .

- **SOTA:** ~90% accurate for many languages when given many training examples, some progress in analyzing languages given few or no examples

Semantic Ambiguity

At last, a computer that understands you like your mother.

- Direct Meanings:

- It understands you like your mother (does) [presumably well]
- It understands (that) you like your mother
- It understands you like (it understands) your mother

- But there are other possibilities, e.g. mother could mean:

- a woman who has given birth to a child
- a stringy slimy substance consisting of yeast cells and bacteria; is added to cider or wine to produce vinegar

- Context matters, e.g. what if previous sentence was:

- Wow, Amazon predicted that you would need to order a big batch of new vinegar brewing ingredients. 😊

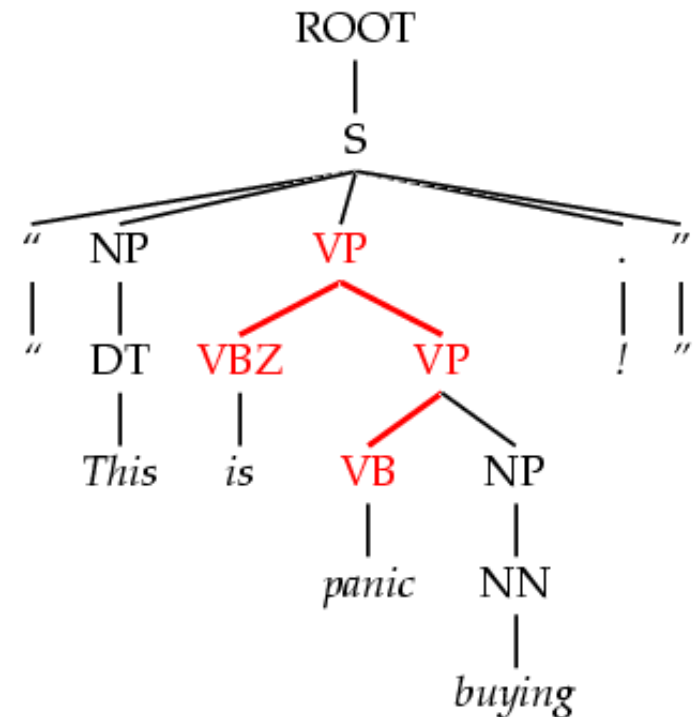
[Example from L. Lee]

Dark Ambiguities

- *Dark ambiguities*: most structurally permitted analyses are so bad that you can't get your mind to produce them

This analysis corresponds to the correct parse of

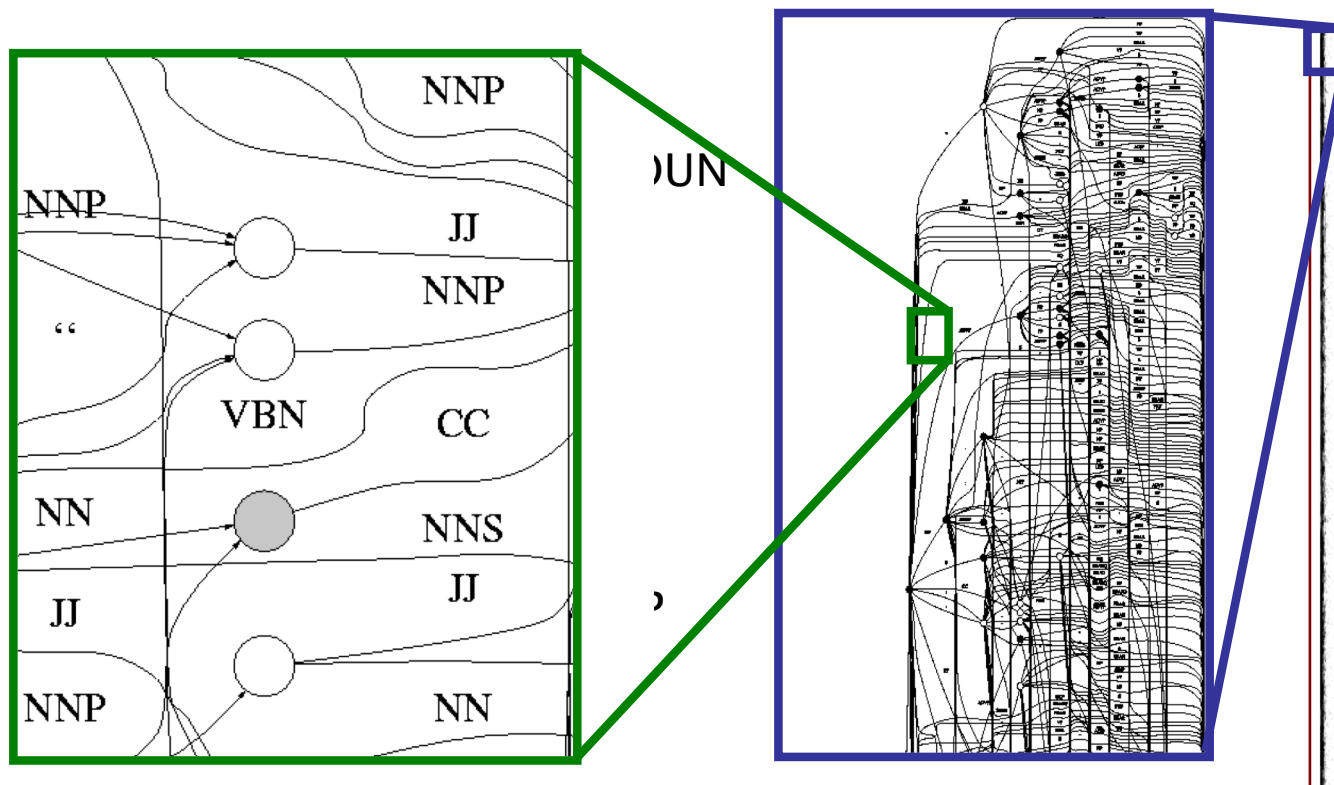
“This will panic buyers ! ”



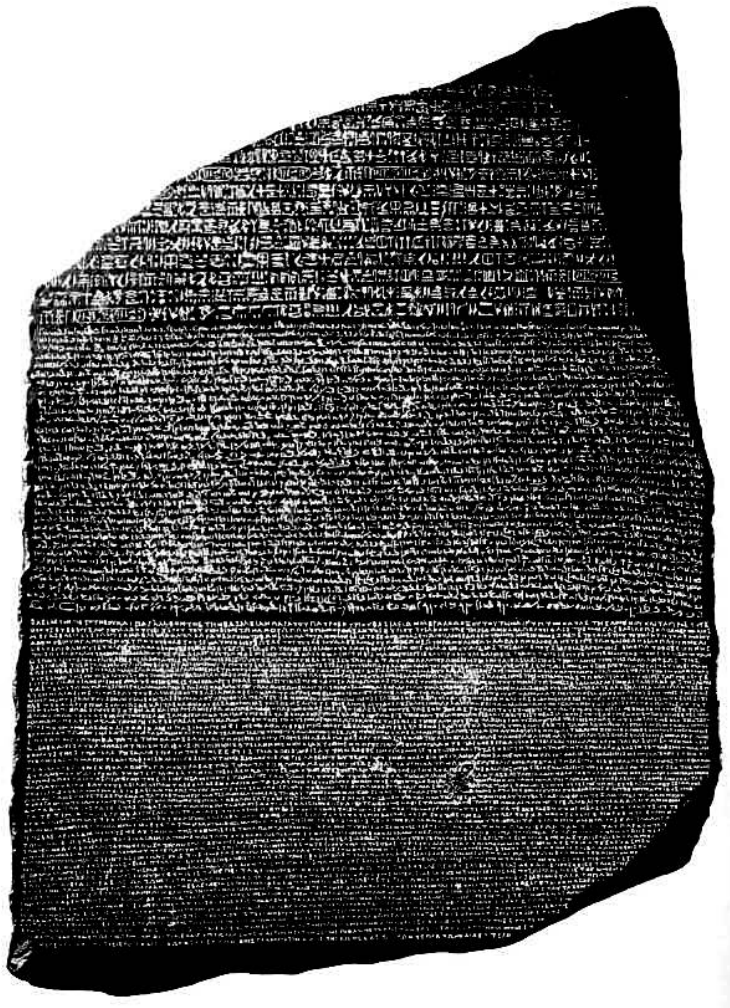
- Unknown words and new usages
- **Solution**: We need mechanisms to focus attention on the best ones, probabilistic techniques do this

Problem: Scale

- People *did* know that language was ambiguous!
 - ...but they hoped that all interpretations would be “good” ones (or ruled out pragmatically)
 - ...they didn't realize how bad it would be



Corpora



- A corpus is a collection of text
 - Often annotated in some way
 - Sometimes just lots of text
 - Balanced vs. uniform corpora
- Examples
 - Newswire collections: 500M+ words
 - Brown corpus: 1M words of tagged “balanced” text
 - Penn Treebank: 1M words of parsed WSJ
 - Canadian Hansards: 10M+ words of aligned French / English sentences
 - The Web: billions of words of who knows what

Problem: Sparsity

- However: sparsity is always a problem
 - New unigram (word), bigram (word pair)




Table of Content

- Definition of NLP
- Historical account of NLP
- Unique challenges of NLP
- Class administrivia / discussion

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?
 - Statistical Modeling Methods
 - Increasingly complex 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...

Comparisons with Other Classes

- Compared to ML
 - Typically multivariate, dynamic programming everywhere
 - Structural Learning & Inference
 - Insights into language matters (a lot!)
 - DL: RNNs, LSTMs, Seq-to-seq, Attention, ...
- Compared to undergrad NLP
 - Faster paced
 - Stronger engineering skills & higher degree of independence assumed
- Compared to CompLing classes
 - More focus on core algorithm design, technically more demanding in terms of math, algorithms, and programming

Class Requirements and Goals

- Class requirements

- Uses a variety of skills / knowledge:
 - Probability and statistics
 - Basic linguistics background
 - Decent coding skills
- Most people are probably missing one of the above
- You will often have to work 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!