

Natural Language Processing

Introduction, course logistics.

Yulia Tsvetkov

yuliats@cs.washington.edu

Welcome!

<https://courses.cs.washington.edu/courses/cse447/25au/>

Mon / Wed / Fri 11:30–12:20pm, CSE2 G10

CSE 447: Natural Language Processing, Fall 2025

MWF 11:30-12:20pm, CSE2 G10 (Gates, ground floor)



Instructor: [Yulia Tsvetkov](#)

yuliats@cs.washington.edu

OH: available on Zoom by appointment.



Teaching Assistant: [Melanie Sclar](#)

msclar@cs.washington.edu

OH: Fri 9:30-10:30am, CSE1 220, Zoom



Teaching Assistant: [Kabir Ahuja](#)

kahuja@cs.washington.edu

OH: Wed 10:30-11:30am, CSE1 220, Zoom



Teaching Assistant: [Min Jang](#)

minjang@cs.washington.edu

OH: Tue 1:30-2:30pm, CSE2 150, Zoom



Teaching Assistant: [Anagha Rao](#)

uma23@cs.washington.edu

OH: Thu 9:30-10:30am, CSE1 220, Zoom



Teaching Assistant: [Leonardo Chen](#)

clx@cs.washington.edu

OH: Mon 2:00-3:00pm, CSE1 218, Zoom

Instructors



Teaching Assistant: [Melanie Sclar](#)

msclar@cs.washington.edu

OH: Fri 9:30-10:30am, CSE1 220, [Zoom](#)



Teaching Assistant: [Anagha Rao](#)

uma23@cs.washington.edu

OH: Thu 9:30-10:30am, CSE1 220, [Zoom](#)



Teaching Assistant: [Kabir Ahuja](#)

kahuja@cs.washington.edu

OH: Wed 10:30-11:30am, CSE1 220, [Zoom](#)



Teaching Assistant: [Leonardo Chen](#)

clx@cs.washington.edu

OH: Mon 2:00-3:00pm, CSE1 218, [Zoom](#)



Teaching Assistant: [Min Jang](#)

minjang@cs.washington.edu

OH: Tue 1:30-2:30pm, CSE2 150, [Zoom](#)

Communication with machines

- ~1950s-1970s



Communication with machines

- ~1980s

```

File Edit Edit_Settings Menu Utilities Compilers Test Help
EDIT BS9U.DEVT3.CLIPPAU(TIMMIES) - 01.31 Columns 00001 000
Command ==> | Scroll ==> H
***** Top of Data *****
000001 /* REXX EXEC *****
000002 /*
000003 /* TIMMIES FACTOR - COMPOUND INTEREST CALCULATOR
000004 /*
000005 /* AUTHOR: PAUL GAMBLE
000006 /* DATE: OCT 1/2007
000007 /*
000008 /*
000009 /******
000010
000011
000012 say '*****'
000013 say 'Welcome Coffee drinker.'
000014 say '*****'
000015 DO WHILE DATATYPE(CoffeeAmt) \= 'NUM'
000016 say ""
000017 say "What is the price of your coffee?",
000018 say "(e.g. 1.58 = $1.58)"
000019 parse pull CoffeeAmt
000020 END
000021
000022 DO WHILE DATATYPE(CoffeeWk) \= 'NUM'
000023 say ""
000024 say "How many coffees a week do you have?"
000025 parse pull CoffeeWk
000026 END
000027
000028 DO WHILE DATATYPE(Rate) \= 'NUM'
000029 say ""
000030 say "What annual interest rate would you like to see on that money?",
000031 say "(e.g. 8 = 8%)"
000032 parse pull Rate
000033 END
000034 Rate = Rate * 0.01 /* CHG TO DECIMAL NUMBER */
000035

```

NLP: Communication with machines

- 2022



WeKnowMemes

NLP: Communication with machines

- Today

explain quantum theory to a child in the style of snoop dogg

Aight lil homie, let me break it down for ya Snoop style:

So you know how the world feels all solid and chill? Like your toys, your juice box, or even your doggy? Well, quantum theory is like peekin' behind the curtain and seein' that deep down, everything's actin' wild and funky.

- **Tiny vibes:** The world is made of teeny-tiny things called particles. They so small you can't see 'em, even with your grandma's glasses. But they ain't sittin' still—nah, they dancin', spinnin', and groovin' all the time.
- **Wave game:** Sometimes these little dudes act like marbles, sometimes like waves in the ocean. It's like your toy car suddenly decides it's also a puddle of water. That's that quantum double life.
- **Rollin' dice:** You can't always know exactly what they're gonna do. It's like the universe playin' dice, throwin' surprises left and right. You think the cookie's in the jar—but maybe it ain't. Gotta check to know for real.
- **Spooky homies:** And here's the trippiest part—two particles can be far apart, but still act like twins who finish each other's sentences. They got that *entanglement*, always vibin' together no matter the distance. That's some long-distance friendship, fo' shizzle.

So quantum theory is basically the rulebook for this cosmic dance party. It tells us how the tiniest pieces of reality move, groove, and surprise us.

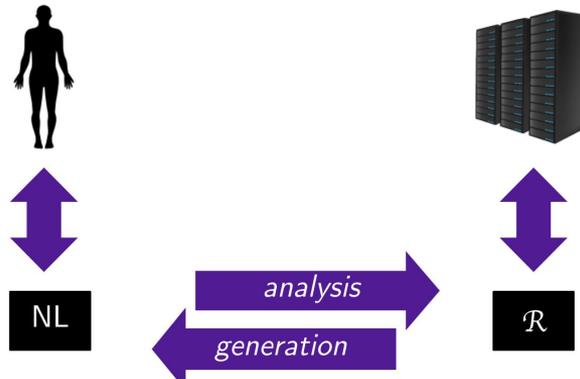
🧐 Translation: The universe ain't as simple as it looks. Deep down, it's mysterious, playful, and a little gangster.

Want me to turn this into a lil' bedtime rap verse for a kid, Snoop style?



What is Natural Language Processing (NLP)?

- $NL \in \{\text{Mandarin Chinese, Hindi, Spanish, Arabic, English, ... Inuktitut, Njerep}\}$
- Automation of NLPs:
 - **analysis** of (“understanding”) what a text means, to some extent ($NL \rightarrow \mathcal{R}$)
 - **generation** of fluent, meaningful, context-appropriate text ($\mathcal{R} \rightarrow NL$)
 - **representation learning** – acquisition of \mathcal{R} from knowledge and data



Language technologies

What technologies are required to write such a program?



A conversational agent contains

-
-
-
-
-

Language Technologies



A conversational agent contains

- Speech recognition
- Language analysis
- Dialog processing
- Information retrieval
- Text to speech

Natural Language Processing



A conversational agent contains

- Speech recognition
- Language analysis
 - Language modelling, spelling correction
 - Syntactic analysis: part-of-speech tagging, syntactic parsing
 - Semantic analysis: named-entity recognition, event detection, word sense disambiguation, semantic role labelling
 - Longer range semantic analysis: coreference resolution, entity linking
 - Deeper semantic analysis: reasoning, knowledge, ect.
- Dialog processing
 - Discourse analysis, user adaptation, etc.
- Information retrieval
- Text to speech

Syllabus

<https://courses.cs.washington.edu/courses/cse447/25au/>

- **Introduction**
 - Overview of NLP as a field
- **Modeling (ML fundamentals)**
 - Text classification: linear models (perceptron, logistic regression), non-linear models (FF NNs, CNNs)
 - Language modeling: n-gram LMs, neural LMs, RNNs
 - Representation learning: word vectors, contextualized word embeddings, Transformers
- **Large language models**
 - Pretraining, in-context learning, fine-tuning, decoding, prompting, reasoning
- **Applications & research**
 - Text classification: Sentiment analysis, toxicity detection
 - Language generation: Machine translation, summarization, QA
 - AI safety, AI Ethics
 - NLP in Industry

Course structure

Will be updated on the course website's syllabus

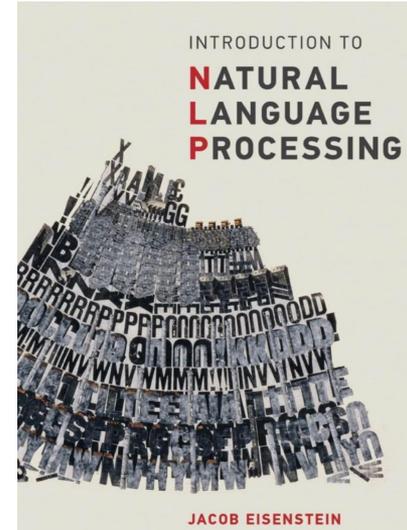
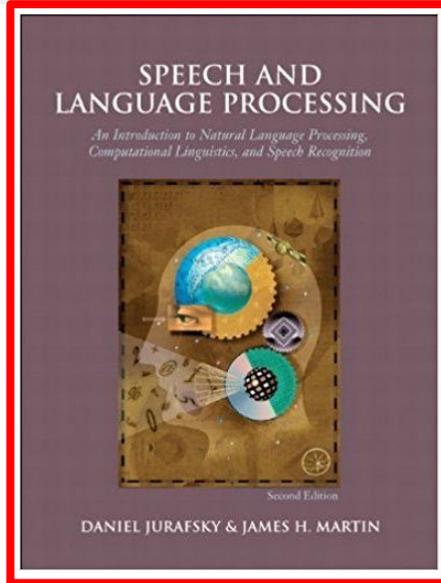
<https://courses.cs.washington.edu/courses/cse447/25au/>

Calendar

Calendar is tentative and subject to change. More details will be added as the quarter continues.

Week	Date	Topics	Readings	Homeworks
------	------	--------	----------	-----------

Readings



- <https://web.stanford.edu/~jurafsky/slp3/>
- <https://github.com/jacobeisenstein/gt-nlp-class/blob/master/notes/eisenstein-nlp-notes.pdf>
- +additional readings posted weekly, lots of resources online

Deliverables & grading

- **Homework projects - 90%**
 - 3 programming assignments, 30% each
 - “Semi-autograded” – Most of the grades (~70-80%) come from evaluating if the submission passes the hidden test cases. Sample test cases will also be provided for students to check their implementations. The rest of the grades would involve writeups on algorithm details, performance trends, and other conceptual questions.
 - We’ll discuss the setup in detail when HW1 is released
- **Quizzes - 10%**
 - 7 (relatively simple, but closed book) quizzes weekly
 - 10 minutes at the beginning or end of the class
 - Starting from the 3rd week, often on Mondays
 - 5 best quizzes, 2% each
- **Participation in course discussions - 6% bonus**
 - **A helpful response to HW questions** and discussions from your classmates on Ed
 - Contribute “insightful” discussions on Ed - 2% extra credit per response, 6% max

Homework assignments

- Project 0: [Optional] **Python and Pytorch Tutorial / Review**
- Project 1: **Text Classification and N-gram language models**
 - Implementing Naive Bayes and Logistic Regression for text classification
 - Training, evaluating, and sampling from n-gram Language Models
- Project 2: **Neural Text Classification and Neural Language Modeling***
 - Training feed-forward neural networks for text classification using word2vec and sentence transformers representations
 - Training a transformer-based language model from scratch
- Project 3: **Fine-tuning and Prompting Pre-trained Language Models***
 - Fine-tuning pre-trained model for text classification
 - Prompting LLMs for reasoning / QA. Will cover different prompting methods like In-context learning, CoT, and self-consistency as well as other tricks such as RAG

*Subject to change based on factors like class performance, compute feasibility, and topics covered during the course.

Learning goals

At the end of this course, you will be able to:

- Build a supervised classifier to solve problems like sentiment classification
- Build a neural network and train it using stochastic gradient descent
- Build tools for extracting linguistic knowledge from raw text, e.g. names
- Learn ML fundamentals for text processings (including state-of-the-art methods)
- Learn important algorithms for text processings (that are useful also in other fields)
- Learn methodological tools (training/test sets, cross-validation)
- Build a toy “large” language model, fine-tune it for text classification, prompt it for question answering

- It's gentle (my goal is to explain everything) and broad (covering many many topics)
- Mastery independent learning, quizzes and programming homeworks
- No research project, but fun research-oriented lectures towards the end of the course

Late submissions

- **Late policy**

- Each student will be granted **5 late days** to use over the duration of the quarter.
- You can use a **maximum of 3 late days on any one project**.
- Weekends and holidays are also counted as late days.
- Late submissions are automatically considered as using late days.
- Using late days will not affect your grade.
- However, **projects submitted late after all late days have been used will receive no credit**. Be careful!

- **Additional late days**

- We allocate an extra week for each homework assignment
 - E.g. if we believe that the homework will take you 2 weeks to complete, we set a deadline in 3 weeks
 - Start early!

- **We will not grant any extensions beyond these**

Class participation

- **In-person** instruction!
- Lectures and homework assignments complement each other
- Lecture materials are broader
- Homework assignments will go deeper into important topics
- Try to attend the lectures
- Quizzes are designed to encourage you to do so
- But if you miss a lecture – you can read assigned book chapters, read slides
- Participate in class discussions on Ed, 6% bonus is an incentive
 - But don't just provide code solutions to questions on homework projects– those are for individual work!
 - Provide insights, theoretical background, references to readings
- **Your questions are always welcome!**

Communication with instructors

- <https://courses.cs.washington.edu/courses/cse447/25au/>
 - Office hours, announcements, calendar, etc.
 - But check most recent announcements on Ed
- You should be able to see yourselves be added to the Ed discussion board of CSE 447 / CSE M 547 25 au. **Please contact the staff if you are not.**
- **Discussion Board (EdSTEM)** will be used to answer questions related to lectures and assignments
 - We really encourage you to ask/discuss higher level questions on the discussion board.
 - We encourage that generic questions should be posted as “Public” so that other classmates would also get benefited from it.
 - Please do not post detail about your solutions (detail ideas, codes, etc.) on public threads. Private discussion should be used for these posts.
- For grading issues, please email the instructor team directly.

Office hours

Questions about homework assignments, grading, course logistics:

- **Melanie** – head TA; NLP research, outstanding personal issues, logistics, not grading
- **Kabir** – HW lead; HW or NLP research questions, not grading/logistics
- **Min, Anagha, Leo, Melanie, Kabir** – HW support
- **Yulia** – CSE 566 only by appointment, questions about lectures, research, NLP in general
- **Teaching sections** – We'll announce when we will have a teaching section; not held by default

	A	B	C	D	E	F
		Monday	Tuesday	Wednesday	Thursday	Friday
9:00 AM						
9:30 AM					Anagha CSE 220	Melanie CSE 220
10:00 AM						
10:30 AM				Kabir CSE 220		
11:00 AM						
11:30 AM		LECTURE		LECTURE		LECTURE
12:00 PM						
12:30 PM						
1:00 PM						
1:30 PM			Min CSE2 150			
2:00 PM		Leo CSE 218				
2:30 PM						
3:00 PM						
3:30 PM						

Quizzes

- 7 quizzes, students can drop 3
- Each quiz has ~5 multiple-choice questions, autograded
- Quizzes are on Canvas, during the lecture time
- Quiz time - 10 minutes in the beginning of the class
- Starting from the 3rd week
- On Mondays unless we announce otherwise
- Grading on 5 best quizzes, 2% each
- Important: only Canvas window should be open during the quiz. We autograde the quiz but then check report from Canvas if you left the window during the quiz (e.g. switched to Chrome). We will zero-out all reported quizzes.

ChatGPT, Copilot, and other AI assistants

- **Quizzes:** canvas functionalities
- **Homework assignments**
 - You can “consult” with ChatGPT like you’d do with another student in the class
 - You cannot feed HW questions and paste solutions
 - We’ll run automated plagiarism checks
 - In the assignments you’ll be asked to clarify whether/how you used generative AI

Session Information

Started at Thu Dec 10 2020 11:53:10 GMT-0700 (Mountain Sta
Attempt **1**

Action Log

- 00:03 ○ Session started
- 00:33 ● Resumed.
- 00:48 ○ Viewed (and possibly read) question #1
- 00:48 ○ Viewed (and possibly read) question #2
- 00:48 ○ Answered question: #1
- 00:53 ○ Answered question: #2
- 01:03 ○ Viewed (and possibly read) question #3
- 02:29 ○ Answered question: #3
- 02:50 ○ Answered question: #4
- 03:03 ○ Viewed (and possibly read) question #4
- 04:18 ○ Viewed (and possibly read) question #5
- 04:33 ○ Answered question: #5
- 05:33 ○ Viewed (and possibly read) question #6
- 05:49 ⊗ Stopped viewing the Canvas quiz-taking page...
- 06:03 ● Resumed.
- 07:03 ○ Answered question: #6

Your Brain on ChatGPT: Accumulation of Cognitive Debt when Using an AI Assistant for Essay Writing Task

Nataliya Kosmyna, Eugene Hauptmann, Ye Tong Yuan, Jessica Situ, Xian-Hao Liao, Ashly Vivian Beresnitzky, Iris Braunstein, Pattie Maes

This study explores the neural and behavioral consequences of LLM-assisted essay writing. Participants were divided into three groups: LLM, Search Engine, and Brain-only (no tools). Each completed three sessions under the same condition. In a fourth session, LLM users were

What background do I need to have?

- 447/547 prerequisite courses
- Python programming
- ML is not a prerequisite but we very strongly suggest to take the course only if you have some ML background
- Prior experience in linguistics or natural languages is helpful, but not required
- There will be a lot of algorithms and coding in this class, some statistics, probabilities, linear algebra

More course logistics

We care that you learn!

Your questions are always welcome.

The screenshot shows an Ed Discussion forum for the course "CSE 447 / CSE M 547 - 24au". The forum is titled "Switch to 547 #6".

Question: An anonymous user asks, "Hello, I have two short questions. Is 547 offered, and is it the graduate version of this class? Does it entail B problems like in past offerings of 446/546? Can we switch to it for a greater challenge like past offerings of 446?"

Answer: Yulia Tsvetkov (STAFF) replies, "Yes 547 is offered. The class will be the same for 447 and 547, with additional questions in homework assignments that will be required for 547 and optional for 447. Up to you to switch to 547 if you're eligible!"

447 vs 547

- Same course content
- Same quizzes
- Additional question in HW assignments for 547 that will be a bonus question for 447

Course registration

- The instructor cannot generate an Add Code
- ~~If you wish to register to the course and have completed prerequisite courses~~
 - ~~Fill out the [500 level course enrollment request form from \(managed by the grad advisers\)](#)~~
 - ~~<https://docs.google.com/forms/d/e/1FAIpQLSc9IbYwpg4KmbiGMmYSA7Ju11G8HZiSbnazwn9M4DNf1UGZOW/viewform>~~
 - ~~Email Pim Lustig <pl@cs.washington.edu> and Ugrad Adviser <ugrad-adviser@cs.washington.edu> to request an Add Code~~
 - ~~Cc Yulia~~
- Unfortunately we cannot add new students to the course, sorry :(

Questions?

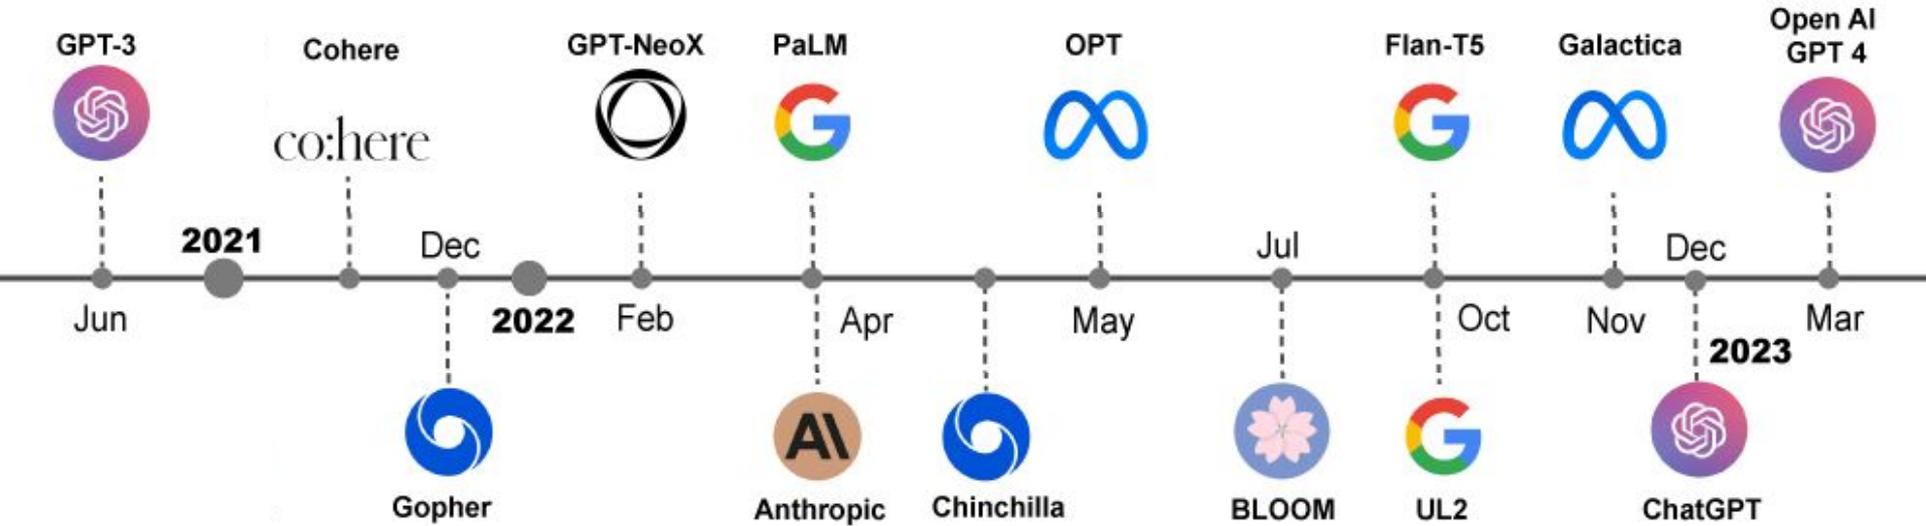
Examples of NLP tools: Personal assistants

amazon alexa

Siri



Large Language Models (LLMs)



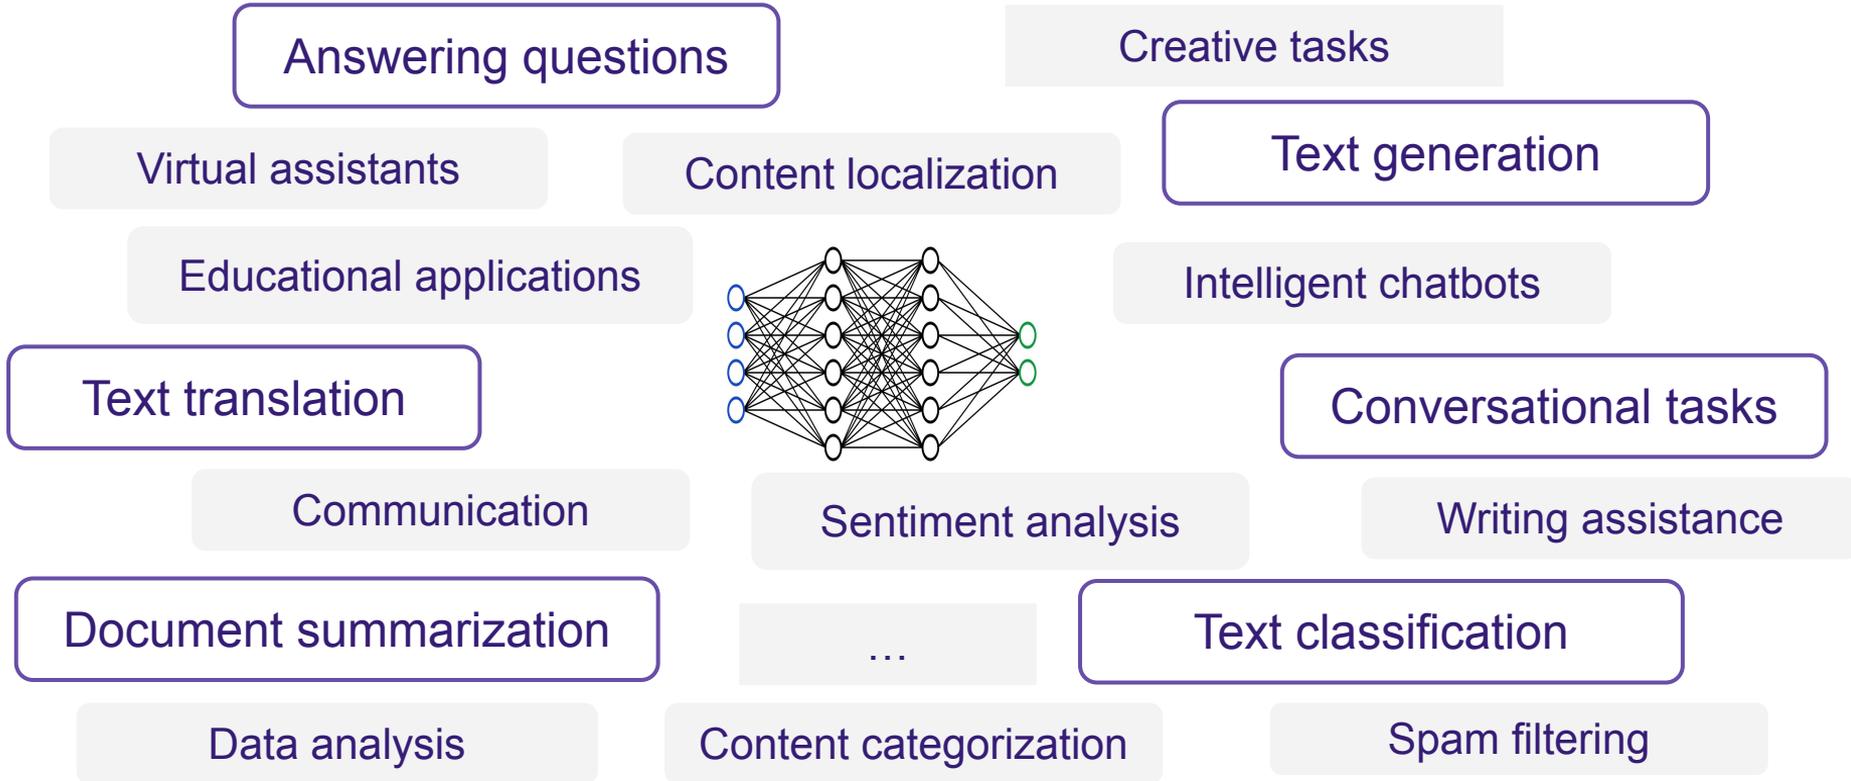
Natural Language Processing



A conversational agent contains

- Speech recognition
- Language analysis
 - Language modelling, spelling correction
 - Syntactic analysis: part-of-speech tagging, syntactic parsing
 - Semantic analysis: named-entity recognition, event detection, word sense disambiguation, semantic role labelling
 - Longer range semantic analysis: coreference resolution, entity linking
 - Deeper semantic analysis: reasoning, knowledge, ect.
- Dialog processing
 - Discourse analysis, user adaptation, etc.
- Information retrieval
- Text to speech

Common tasks LLMs are trained to solve

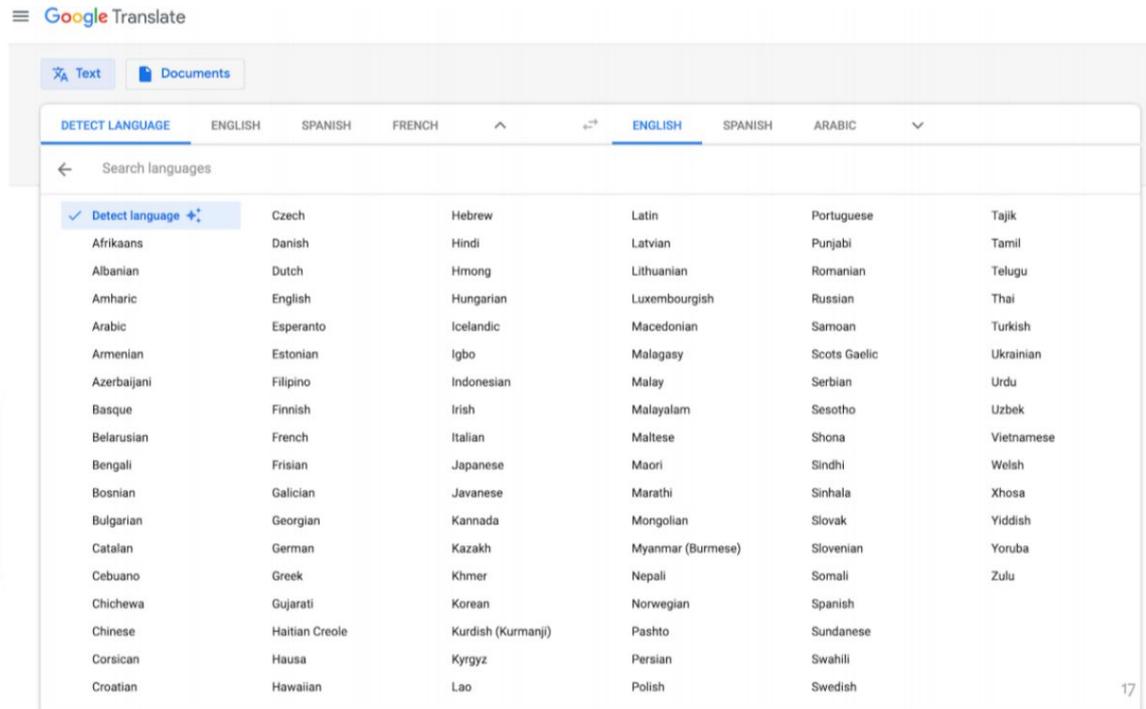


Question answering

- What does “divergent” mean?
- What year was Abraham Lincoln born?
- How many states were in the United States that year?
- How much Chinese silk was exported to England in the end of the 18th century?
- What do scientists think about the ethics of human cloning?



Machine translation



Positive or negative movie review?



- unbelievably disappointing



- Full of zany characters and richly applied satire, and some great plot twists



- this is the greatest screwball comedy ever filmed



- It was pathetic. The worst part about it was the boxing scenes.

Sentiment analysis



HP Officejet 6500A Plus e-All-in-One Color Ink-jet - Fax / copier / printer / scanner
\$89 online, \$100 nearby ★★★★★ 377 reviews
 September 2010 - Printer - HP - Inkjet - Office - Copier - Color - Scanner - Fax - 250 sh

Reviews

Summary - Based on 377 reviews



What people are saying

ease of use		"This was very easy to setup to four computers."
value		"Appreciate good quality at a fair price."
setup		"Overall pretty easy setup."
customer service		"I DO like honest tech support people."
size		"Pretty Paper weight."
mode		"Photos were fair on the high quality mode."
colors		"Full color prints came out with great quality."

Information extraction

In 1933, while Einstein was visiting the United States, Adolf Hitler came to power.

Because of his Jewish background, Einstein did not return to Germany.

He settled in the United States and became an American citizen in 1940.

Einstein supported the Allied forces, but he generally denounced the idea of using nuclear fission as a weapon.

He signed the Russell -- Einstein Manifesto with British philosopher Bertrand Russell, which highlighted the danger of nuclear weapons.

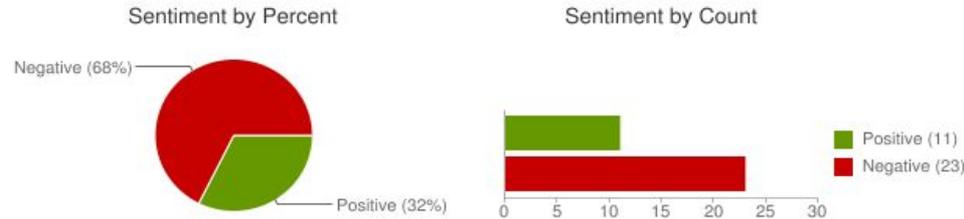
He was affiliated with the Institute for Advanced Study in Princeton, New Jersey, until his death in 1955.

Sentiment analysis + information extraction

Type in a word and we'll highlight the good and the bad

[Save this search](#)

Sentiment analysis for "united airlines"



iljacobson: OMG... Could @United airlines have worse customer service? W8g now 15 minutes on hold 4 questions about a flight 2DAY that need a human.
Posted 2 hours ago

12345clumsy6789: I hate United Airlines Ceiling!!! Fukn impossible to get my conduit in this damn mess! ?
Posted 2 hours ago

EMLandPRGbelgiu: EML/PRG fly with Q8 united airlines and 24seven to an exotic destination. <http://t.co/Z9QloAjF>
Posted 2 hours ago

CountAdam: FANTASTIC customer service from United Airlines at XNA today. Is tweet more, but cell phones off now!
Posted 4 hours ago

Information extraction for disaster relief

- Haiti Earthquake 2010
- About 3 million people were affected by the quake
- Classifying SMS messages



Information extraction for disaster relief

- SMS messages start streaming in

- Fanmi mwen nan Kafou, 24 Cote Plage, 41A bezwen manje ak dlo
- Moun kwense nan Sakre Kè nan Pòtoprens
- Ti ekipman Lopital General genyen yo paka minm fè 24 è
- Fanm gen tranche pou fè yon pitit nan Delmas 31



Information extraction for disaster relief

- Translation

- Fanmi mwen nan Kafou, 24 Cote Plage, 41A bezwen manje ak dlo
- Moun kwense nan Sakre Kè nan Pòtoprens
- Ti ekipman Lopital General genyen yo paka minm fè 24 è
- Fanm gen tranche pou fè yon pitit nan Delmas 31
- My family in Carrefour, 24 Cote Plage, 41A needs food and water
- People trapped in Sacred Heart Church, PauP
- General Hospital has less than 24 hrs. supplies
- Undergoing children delivery Delmas 31

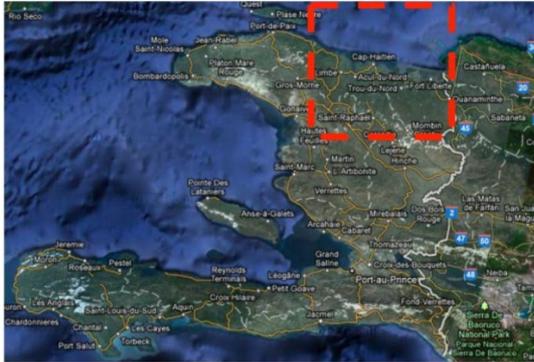


Information extraction for disaster relief

- Translation + information extraction

Lopital Sacre-Coeur ki nan vil Okap, pre pou li resevwa moun malad e lap mande pou moun ki malad yo ale la.

“Sacre-Coeur Hospital which located in this village of **Okap** is ready to receive those who are injured. Therefore, we are asking those who are sick to report to that hospital.”



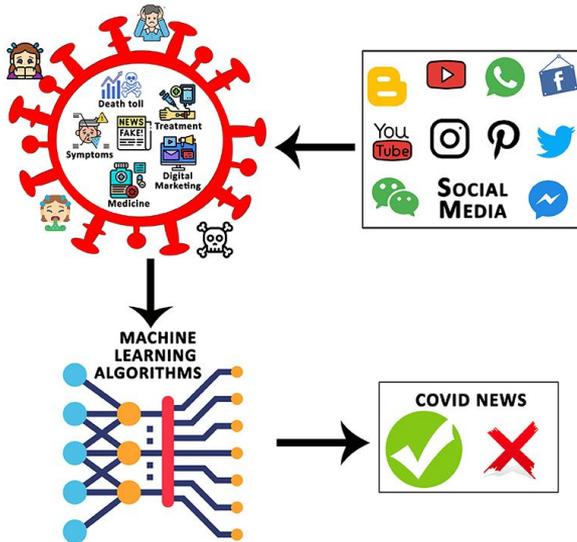
iDIBON



An earthquake struck Haiti on January 12, 2010

Most local services failed, but most cell-towers remained functional.

Covid19 misinformation



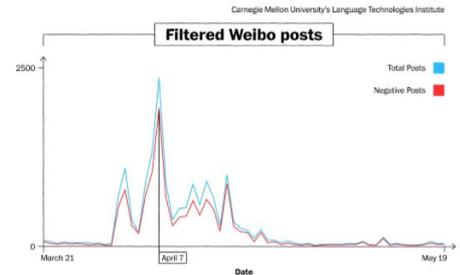
Detecting COVID-19-Related Fake News Using Feature Extraction

Suleman Khan, Saqib Hakak, N. Deepa, B. Prabadevi, Kapal Dev and Silvia Trelova

<https://www.washingtonpost.com/politics/2020/06/18/video-evidence-anti-black-discrimination-china-over-coronavirus-fears/>

The Fact Checker worked with researchers at professor Yulia Tsvetkov's lab at Carnegie Mellon University's Language Technologies Institute and the Center for Human Rights Science to track what happened on social media during this period. Researchers collected about 16,000 Weibo posts, filtered from a larger data set of 200,000 posts, containing at least one Guangzhou location tag and one "African-related" keyword from late March through May. Weibo is a Chinese social media platform.

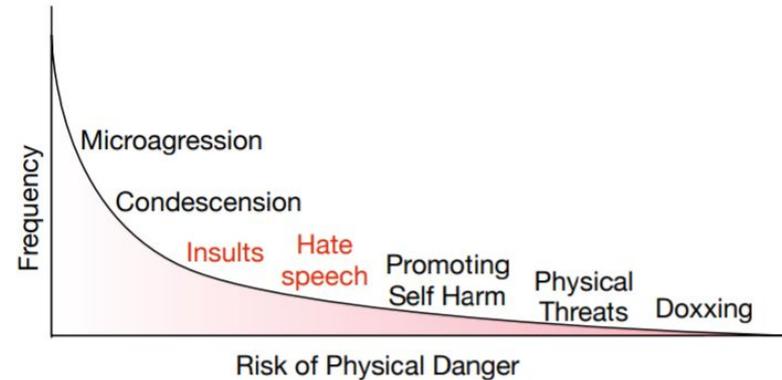
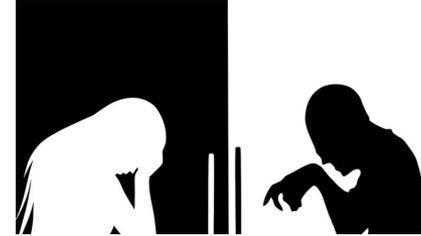
Based on automated sentiment analysis tools and manual analysis, the researchers said, they believed the majority of posts in their data set expressed negative sentiments relating to Africans or black people.



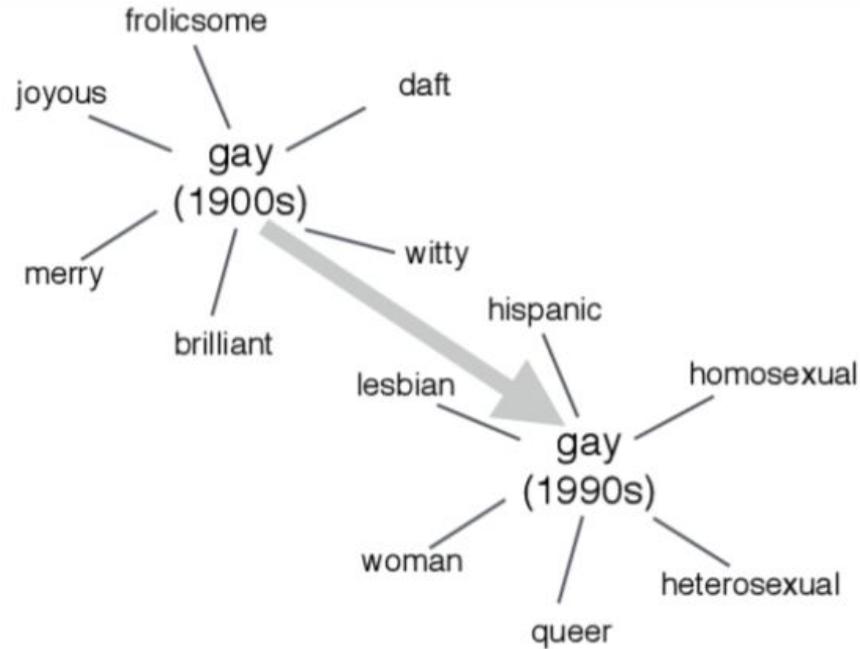
Researchers collected around 16,000 posts containing at least one Guangzhou location tag and one "African-related" keyword from late March through May. (Athar Mirza/The Washington Post)

Their research showed there was a significant surge in negative posts beginning April 1. There were just 23 negative posts in the data set on March 31. The next day, the number of posts climbed to 500. From April 1-2, there was a spike in the number of posts on Weibo using the keywords "foreign trash." A Chinese cartoon depicting officials throwing foreigners who weren't abiding by

Hate speech detection



Language change

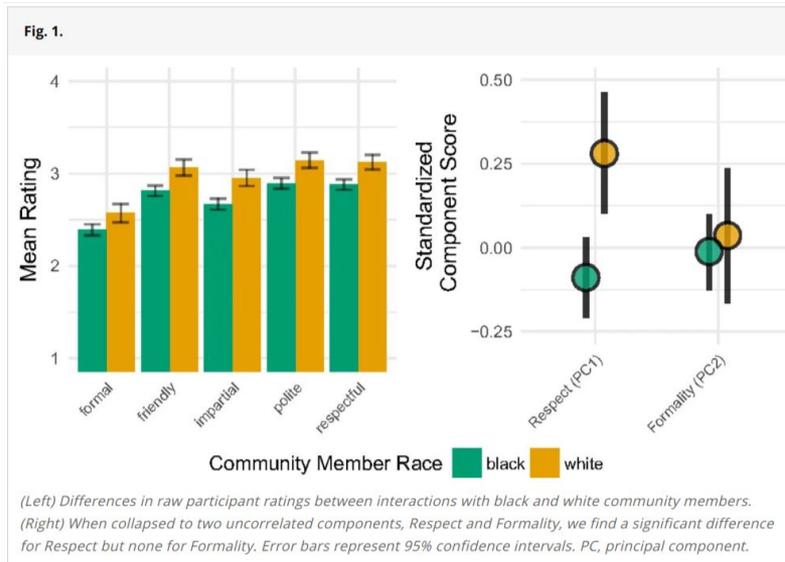


Cultural Shift or Linguistic Drift? Comparing Two Computational Measures of Semantic Change

William L. Hamilton, J. Leskovec, Dan Jurafsky

Computational social science

- computational social science answering questions about society given observational data
- example: "Do police officers speak with Black and White Americans in the same way?"



Language from police body camera footage shows racial disparities in officer respect

Rob Voigt [✉], Nicholas P. Camp, Vinodkumar Prabhakaran, [✉], and Jennifer L. Eberhardt [✉] [Authors Info & Affiliations](#)

Contributed by Jennifer L. Eberhardt, March 26, 2017 (sent for review February 14, 2017; reviewed by James Pennebaker and Tom Tyler)

June 5, 2017 | 114 (25) 6521-6526 | <https://doi.org/10.1073/pnas.1702413114>