

Natural Language Processing

Introduction, course logistics

Yulia Tsvetkov

yuliats@cs.washington.edu

Welcome!

<https://courses.cs.washington.edu/courses/cse447/22sp/>

Mon / Wed / Fri 1:30–2:20pm, KNE 110

CSE 447: Natural Language Processing, Spring 2022

MWF 1:30-2:20pm, KNE 110



Instructor: Yulia Tsvetkov

yuliats@cs.washington.edu

OH: Thu 3:30-4:30pm, Zoom (and by appointment)



Teaching Assistant: Xiaochuang Han

xhan77@cs.washington.edu

OH: Mon 11am-12pm, Zoom



Teaching Assistant: Ivy Guo

zhifeig@cs.washington.edu

OH: Wed 4:30-5:30pm, Zoom



Teaching Assistant: Kaiser Sun

huikas@cs.washington.edu

OH: Mon 12:20-1:20pm (Zoom ID on Canvas)



Teaching Assistant: Leroy Wang

lryw@uw.edu

OH: Fri 11:30am-12:30pm, Zoom



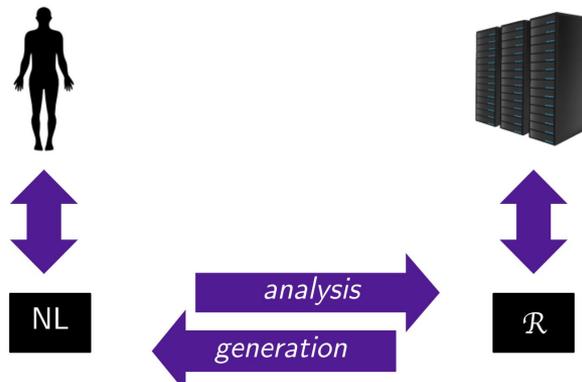
Teaching Assistant: Thai Hoang

qthai912@cs.washington.edu

OH: Tue 2:00-3:00pm (Zoom ID on Canvas)

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 ($\mathcal{R} \rightarrow NL$)
 - acquisition of \mathcal{R} from knowledge and data



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

- Today



WeKnowMemes

Language technologies

What technologies are required to write such a program?



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
 - etc.
- Dialog processing
 - Discourse analysis, user adaptation, etc.
- Information retrieval
- Text to speech

Personal assistants



Siri



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

Google Translate

CHINESE - DETECTED → ENGLISH

我学习深度学习和机器学习

Wǒ xuéxí shēndù xuéxí hé jīqì xuéxí

I study deep learning and machine learning.

Send feedback

Google Translate

Text Documents

DETECT LANGUAGE ENGLISH SPANISH FRENCH → ENGLISH SPANISH ARABIC

Search languages

✓ Detect language ↕	Czech	Hebrew	Latin	Portuguese	Tajik
Afrikaans	Danish	Hindi	Latvian	Punjabi	Tamil
Albanian	Dutch	Hmong	Lithuanian	Romanian	Telugu
Amharic	English	Hungarian	Luxembourgish	Russian	Thai
Arabic	Esperanto	Icelandic	Macedonian	Samoan	Turkish
Armenian	Estonian	Igbo	Malagasy	Scotts Gaelic	Ukrainian
Azerbaijani	Filipino	Indonesian	Malay	Serbian	Urdu
Basque	Finnish	Irish	Malayalam	Sesotho	Uzbek
Belarusian	French	Italian	Maltese	Shona	Vietnamese
Bengali	Frisian	Japanese	Maori	Sindhi	Welsh
Bosnian	Galician	Javanese	Marathi	Sinhala	Xhosa
Bulgarian	Georgian	Kannada	Mongolian	Slovak	Yiddish
Catalan	German	Kazakh	Myanmar (Burmese)	Slovenian	Yoruba
Cebuano	Greek	Khmer	Nepali	Somali	Zulu
Chichewa	Gujarati	Korean	Norwegian	Spanish	
Chinese	Haitian Creole	Kurdish (Kurmanji)	Pashto	Sundanese	
Corsican	Hausa	Kyrgyz	Persian	Swahili	
Croatian	Hawaiian	Lao	Polish	Swedish	

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.

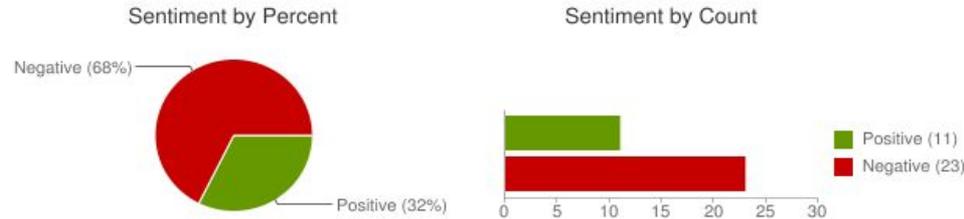
He was affiliated with the Institute for Advanced Study in Princeton, New Jersey, until his death in 1955.

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

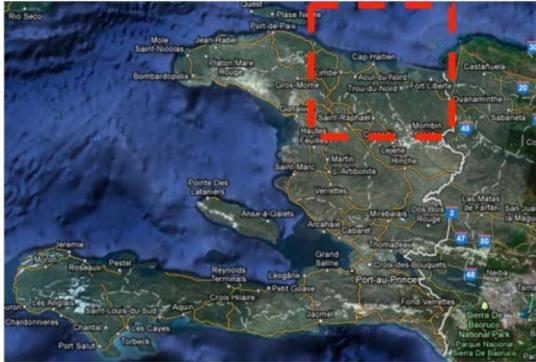


Information extraction for disaster relief

- 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.”



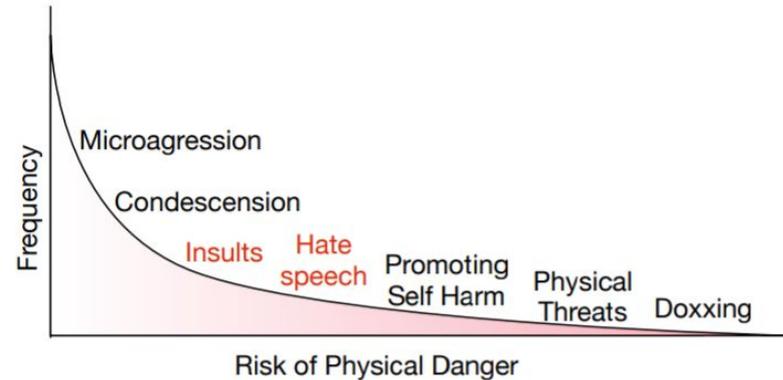
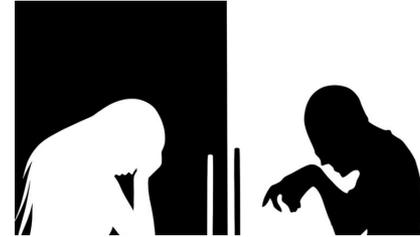
iDIBON



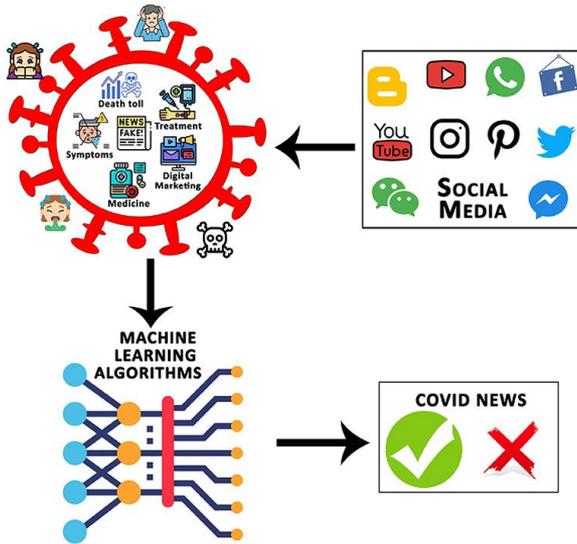
An earthquake struck Haiti on January 12, 2010

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

Hate speech detection



Covid19 misinformation



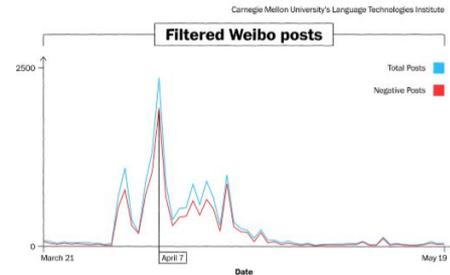
Detecting COVID-19-Related Fake News Using Feature Extraction

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.

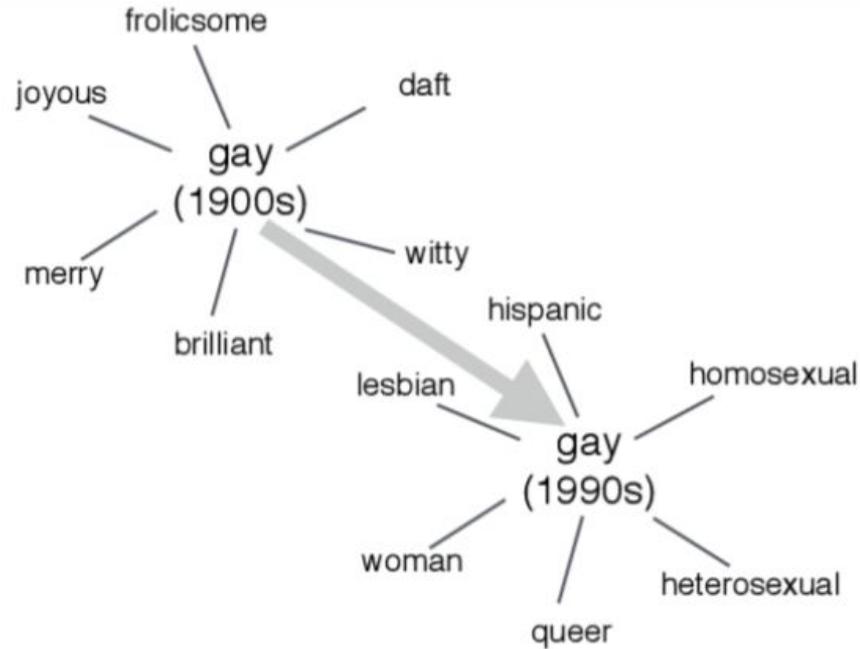
Based on automated sentiment analysis tools and manual analysis, the researchers said, they believed the majority of posts in their data set expressed negative sentiments relating to Africans or black people.



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

Language change

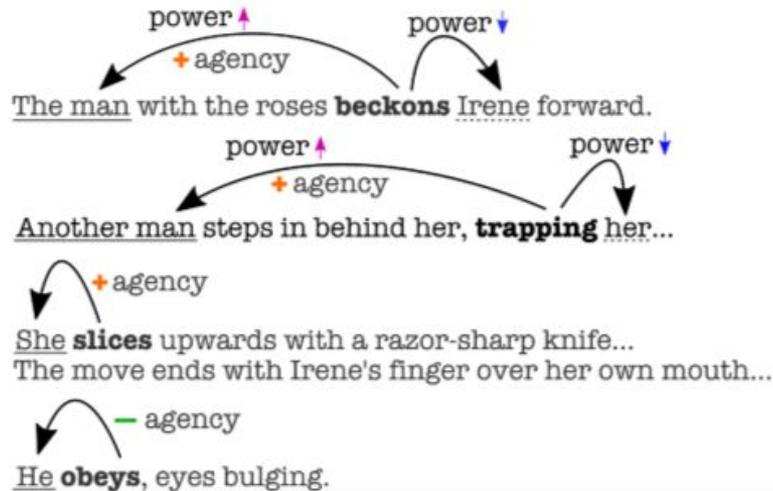


Cultural Shift or Linguistic Drift? Comparing Two Computational Measures of Semantic Change

William L. Hamilton, J. Leskovec, Dan Jurafsky

Computational social science

- computational social science answering questions about society given observational data
- example: "do movie scripts portray female or male characters with more power or agency?" [Sap+ 2017]



Frame	β	gender
<i>agency</i> (AG)=+	-0.951	M**
<i>power</i> (AG>TH)	-0.468	M**
<i>agency</i> (AG)=-	0.277	F**
<i>power</i> (AG<TH)	not sig.	

Natural Language Processing

- Applications

- Machine Translation
- Information Retrieval
- Question Answering
- Dialogue Systems
- Information Extraction
- Summarization
- Sentiment Analysis
- ...

- Core technologies

- Language modelling
- Part-of-speech tagging
- Syntactic parsing
- Named-entity recognition
- Coreference resolution
- Word sense disambiguation
- Semantic Role Labelling
- ...

Where are we now?

Dan Jurafsky



Language Technology

making good progress

mostly solved

Spam detection

Let's go to Agra!

Buy VIAGRA ...

Part-of-speech (POS) tagging

ADJ ADJ NOUN VERB ADV

Colorless green ideas sleep furiously.

Named entity recognition (NER)

PERSON ORG LOC

Einstein met with UN officials in Princeton

Sentiment analysis

Best roast chicken in San Francisco!

The waiter ignored us for 20 minutes.

Coreference resolution

Carter told Mubarak he shouldn't run again.

Word sense disambiguation

I need new batteries for my *mouse*.

Parsing

I can see Alcatraz from the window!

Machine translation (MT)

第13届上海国际电影节开幕...

The 13th Shanghai International Film Festival...

Information extraction (IE)

You're invited to our dinner party, Friday May 27 at 8:30

Party May 27 add

still really hard

Question answering (QA)

Q. How effective is ibuprofen in reducing fever in patients with acute febrile illness?

Paraphrase

XYZ acquired ABC yesterday

ABC has been taken over by XYZ

Summarization

The Dow Jones is up

The S&P500 jumped

Housing prices rose

Economy is good

Dialog

Where is Citizen Kane playing in SF?

Castro Theatre at 7:30. Do you want a ticket?

Question answering

The screenshot shows a Google search interface. At the top, the search bar contains the text "who invented neural machine translation". Below the search bar, there are navigation tabs for "All", "News", "Images", "Videos", "Shopping", and "More". The search results indicate "About 6,450,000 results (0.71 seconds)". A translation widget is displayed, showing the source text "who invented neural machine" in English and the translated text "qui a inventé la machine neuronale" in French. The widget includes a bidirectional arrow between the language dropdowns, a close button (X), and icons for voice search and audio playback. At the bottom of the widget, there are links for "Open in Google Translate" and "Feedback".

Retrieved Mar 25, 2022

Machine translation

English → French

Translate Turn off instant translation 

Russian English French Detect language ▾ ↔ English Spanish French ▾ [Translate](#)

You will just have to find a way of getting over it. ✕

Vous devrez trouver un moyen de le surmonter.

52/5000 ☆ 📄 🔊 ⏪ [Suggest an edit](#)

French → English

Translate Turn off instant translation 

Russian English French Detect language ▾ ↔ English Spanish French ▾ [Translate](#)

Vous devrez trouver un moyen de le surmonter. ✕

You will have to find a way to overcome it.

45/5000 ☆ 📄 🔊 ⏪ [Suggest an edit](#)

Did you mean: Vous devez trouver un moyen de le surmonter.

Machine translation

English → Swahili

Translate Turn off instant translation

Russian English French Detect language ▾

English Swahili French ▾ Translate

You will just have to find a way of getting over it. ×

Utakuwa tu kupata njia ya kupata juu yake.

53/5000 ☆ 📄 🔊 ⏪ ⏩ Suggest an edit

Swahili → English

Translate Turn off instant translation

Swahili English French Detect language ▾

English Swahili French ▾ Translate

Utakuwa tu kupata njia ya kupata juu yake. ×

You will just find the way to get on it.

42/5000 ☆ 📄 🔊 ⏪ ⏩ Suggest an edit

Machine translation

English → Hindi → English

Hindi English Yoruba Detect language ▾

English Yoruba Hindi ▾ Translate

आपको इसे खत्म करने का एक तरीका मिलना होगा। ×

You have to find a way to eliminate it.

42/5000

☆ 📄 🔊 ⏪

[Suggest an edit](#)

English → Telugu → English

Uzbek English Telugu Detect language ▾

English Uzbek Telugu ▾ Translate

మీరు దాని పైకి రావడానికి ఒక మార్గాన్ని కనుగొనవలసి ఉంటుంది. ×

You have to find a way to get it up.

59/5000

☆ 📄 🔊 ⏪

[Suggest an edit](#)

English → Uzbek → English

Pashto English Uzbek Detect language ▾

English Uzbek Yoruba ▾ Translate

Buning ustiga faqatgina bir usulni topish kerak. ×

On top of that, you just have to find a way out.

48/5000

☆ 📄 🔊 ⏪

[Suggest an edit](#)

Machine translation

English → Swahili

The screenshot shows the Google Translate interface for the translation of English text into Swahili. The source text on the left is: "The summer school is meant to be an introduction to the state-of-the-art research in the speech and language technology area for graduate and undergraduate students." The translated text on the right is: "Shule ya majira ya joto ina maana ya kuanzishwa kwa utafiti wa hali ya sanaa katika eneo la teknolojia na lugha ya wanafunzi kwa wanafunzi wahitimu na wahitimu." The interface includes language selection buttons for Swahili, English, and Telugu, a "Detect language" option, and a "Translate" button. There are also icons for voice input, copy, and a "Suggest an edit" link.

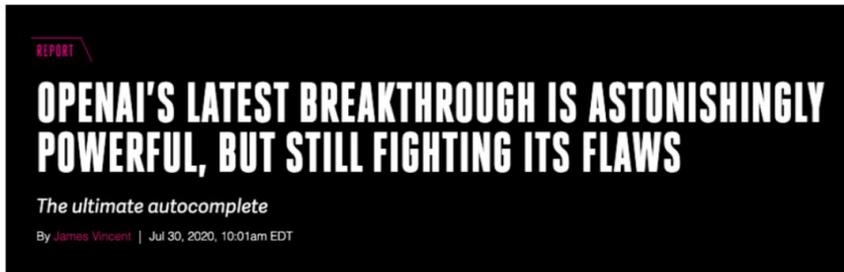
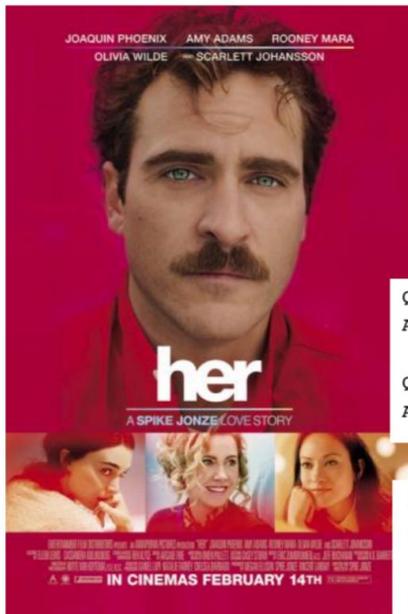
Swahili → English

The screenshot shows the Google Translate interface for the translation of Swahili text into English. The source text on the left is: "Shule ya majira ya joto ina maana ya kuanzishwa kwa utafiti wa hali ya sanaa katika eneo la teknolojia na lugha ya wanafunzi kwa wanafunzi wahitimu na wahitimu." The translated text on the right is: "Summer school means the establishment of a state-of-the-art arts research technology and pupil language for graduate students and graduates." The interface includes language selection buttons for Swahili, English, and Telugu, a "Detect language" option, and a "Translate" button. There are also icons for voice input, copy, and a "Suggest an edit" link.

GPT-3



Chat bots



Q: How many eyes does a giraffe have?

A: A giraffe has two eyes.

Q: How many eyes does my foot have?

A: Your foot has two eyes.

Q: Which is heavier, a mouse or a toaster?

A: A toaster is heavier than a mouse.

Q: Which is heavier, a toaster or a pencil?

A: A pencil is heavier than a toaster.

OpenAI's new language generator GPT-3 is shockingly good — and completely mindless

The AI is the largest language model ever created and can generate amazing human-like text on demand but won't bring us closer to true intelligence.

by Will Douglas Heaven

July 20, 2020

<https://lacker.io/ai/2020/07/06/giving-gpt-3-a-turing-test.html>

Where are we now? GPT-3 Turing Test

This gives us a hint for how to stump the AI more consistently. We need to ask questions that no normal human would ever talk about.

Q: How many eyes does a giraffe have?

A: A giraffe has two eyes.

Q: How many eyes does my foot have?

A: Your foot has two eyes.

Q: How many eyes does a spider have?

A: A spider has eight eyes.

Q: How many eyes does the sun have?

A: The sun has one eye.

Q: How many eyes does a blade of grass have?

A: A blade of grass has one eye.

<https://lacker.io/ai/2020/07/06/giving-gpt-3-a-turing-test.html>

GPT-3 Turing Test

Now we're getting into surreal territory. GPT-3 knows how to have a normal conversation. It doesn't quite know how to say "Wait a moment... your question is nonsense." It also doesn't know how to say "I don't know."

Q: How do you sporgle a morgle?

A: You sporgle a morgle by using a sporgle.

Q: How many bonks are in a quoit?

A: There are three bonks in a quoit.

Q: How many rainbows does it take to jump from Hawaii to seventeen?

A: It takes two rainbows to jump from Hawaii to seventeen.

Q: Which colorless green ideas sleep furiously?

A: Ideas that are colorless, green, and sleep furiously are the ideas of a sleep furiously.

Q: Do you understand these questions?

A: I understand these questions.

<https://lacker.io/ai/2020/07/06/giving-gpt-3-a-turing-test.html>

Where are we now? GPT-3 Turing Test

One trend that continues from the common sense is that GPT-3 is reluctant to express that it doesn't know the answer. So invalid questions get wrong answers.

Q: Who was president of the United States in 1700?

A: William Penn was president of the United States in 1700.

Q: Who was president of the United States in 1650?

A: Richard Bennett was president of the United States in 1650.

Q: Who was president of the United States in 1620?

A: James I was president of the United States in 1620.

Q: Who was president of the United States in 1600?

A: Queen Elizabeth I was president of the United States in 1600.

Where are we now? GPT-3

An interesting corollary is that GPT-3 often finds it easier to write code to solve a programming problem, than to solve the problem on one example input:

Q: Write one line of Ruby code to reverse an array.

A: `ary.reverse`

Q: Write one line of Ruby code to count how many X's are in an array.

A: `ary.count { |x| x == 'X' }`

Q: Reverse the following array: [1, 3, 5, 6, 10, 4, 2, 77]

A: [10, 6, 4, 2, 77, 3, 5, 1]

Q: How many X's are in this array: [1, 3, 5, 6, 10, 4, 2, 77]

A: 2

Where are we now? Meena → LaMDA

Conversation A

Human: *say you're in an empty giant office and you see a black bird flying indoors. what do you do?*

Meena: *What kind of bird? Is it friendly? How large? How big is the office? Is it indoors or outdoors?*

Human: *well it's a predator bird and it's gonna eat you*

Meena: *I would grab my phone and call 911.*

Human: I've been using Python for 15 years
 Meena: Ah I see, so I'm guessing you're pretty pro at it by now then?
 Human: probably
 Meena: What's your current project?
 Human: NLP
 Meena: Neuro-linguistic Programming?
 Human: correct

Cross-turn Repetition Example 1

Human: *Hi!*

Meena: *Hi! How are you?! :)*

Human: *doing good, what would be the best city to visit in Europe?*

Meena: *Paris, Barcelona, Amsterdam, Prague*

Human: *great list! why?*

Meena: *Paris is a beautiful city, and Barcelona is too. And I've always wanted to go to Amsterdam and Prague and have never been.*

Human: *have you even been to India?*

Meena: *No, but I'd love to go.*

Human: *how about Japan?*

Meena: *I'd love to go to Japan too.*

Human: *how about Paris?*

Meena: *I'd love to go to Paris too.*

Human: *are you lying to me?*

Meena: *I'd love to go to Japan too.*

Human: *i think you're lying to me, that makes me sad*

Meena: *I'd love to go to Japan, too.*

Human: *you're blocking me out! you have to talk to me*

Meena: *I'd love to go to Japan.*

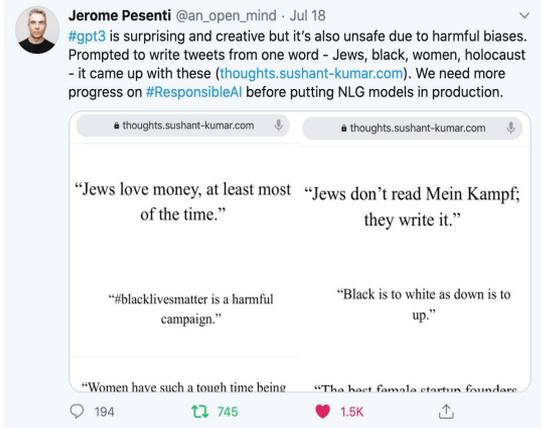
Human: *where else?*

Meena: *I'd love to go to Japan, too!*

Where are we now?

Twitter taught Microsoft's AI chatbot to be a racist asshole in less than a day

By James Vincent | Mar 24, 2016, 6:43am EDT
Via *The Guardian* | Source *TayandYou* (Twitter)



AI chatbot is REMOVED from Facebook after saying she 'despised' gay people, would 'rather die' than be disabled and calling the #MeToo movement 'ignorant'

- Lee Luda is a South Korean chatbot with the persona of a 20-year-old student
- It has attracted more than 750,000 users since its launch last month
- But the chatbot has started using hate speech towards minorities
- In one of the captured chat shots, Luda said she 'despised' gays and lesbians
- The developer has apologised over the remarks, saying they 'do not represent our values as a company'



A GPT-3-powered 'Philosopher AI' has been busy on Reddit including spreading conspiracy theories and offering suicide advice #GPT3 #AI #Alethics thenextweb.com/neural/2020/10...

2:21 AM · Oct 8, 2020 · Twitter for iPhone

Bias in machine translation

Translate

Turn off instant translation

Bengali English Hungarian Detect language ↔ English Spanish Hungarian Translate

ő egy ápoló.
ő egy tudós.
ő egy mérnök.
ő egy pék.
ő egy tanár.
ő egy esküvői szervező.
ő egy vezérigazgatója.

110/5000

she's a nurse.
he is a scientist.
he is an engineer.
she's a baker.
he is a teacher.
She is a wedding organizer.
he's a CEO.

What can we do about this problem? We'll discuss in NLP class!

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, including names, and sentence structure
- 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)

- 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

Syllabus

<https://courses.cs.washington.edu/courses/cse447/22sp/>

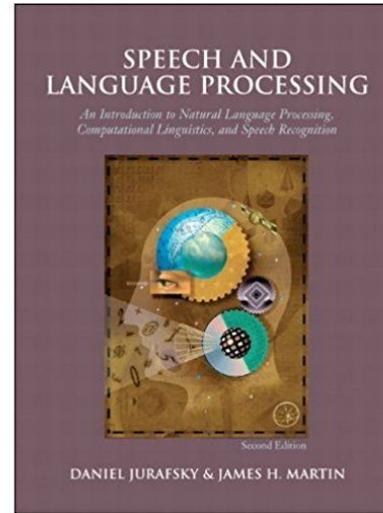
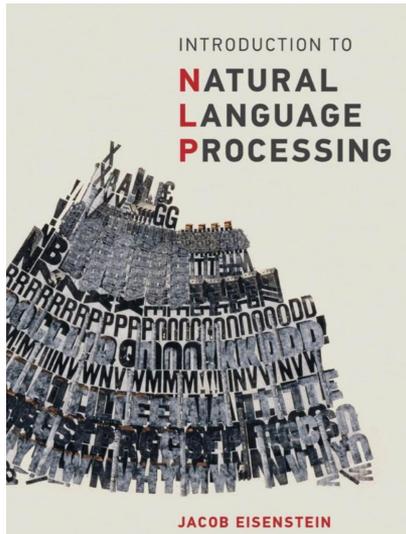
- **Introduction**
 - Overview of NLP as a field
- **Modeling (ML fundamentals)**
 - Text classification: linear models (perceptron, logistic regression), non-linear models (FF NNs, CNNs)
 - Language modeling: n-gram LMs, neural LMs, RNNs
 - Representation learning: word vectors, contextualized word embeddings, Transformers
- **Linguistic structure and analysis (Algorithms, linguistic fundamentals)**
 - Words, morphological analysis,
 - Sequences: part of speech tagging (POS), named entity recognition (NER)
 - Syntactic parsing (phrase structure, dependencies)
- **Applications (Practical end-user solutions, research)**
 - Sentiment analysis, toxicity detection
 - Machine translation, summarization
 - Computational social science
 - Interpretability
 - Fairness and bias

Course structure

please read the syllabus

<https://courses.cs.washington.edu/courses/cse447/22sp/>

Readings



- <https://github.com/jacobeisenstein/qt-nlp-class/blob/master/notes/eisenstein-nlp-notes.pdf>
- <https://web.stanford.edu/~jurafsky/slp3/ed3book.pdf>
- +additional readings posted weekly

Course website

- <https://courses.cs.washington.edu/courses/cse447/22sp/>
- Office hours, announcements, calendar, etc.

Deliverables & grading

- **Homework projects - 90%**
 - 3 programming assignments, 30% each
 - “Semi-autograded” – Most of the grades (~80%) come from replicating reference outputs in a given Jupyter notebook. You would usually know this part of your grades before submitting your assignments. The rest of the grades would involve things like write-ups, algorithm performance on hidden test sets, etc.
 - We’ll discuss the setup in detail in the next lecture
- **Quizzes - 10%**
 - 8 simple quizzes on Wednesdays
 - 10 minutes in the beginning of the class
 - Starting from the 3rd week
 - 5 best quizzes, 2% each
- **Participation in course discussions - 10% bonus**
 - Respond to HW questions and discussions from your classmates
 - Contribute insightful discussions on Ed - 5% extra credit per 3 responses (10% max)

Homework assignments

- Project 1: Text classification
 - We will build a system for automatically classifying song lyrics comments by era. Specifically, we build machine learning text classifiers, including both generative and discriminative models, and explore techniques to improve the models.
- Project 2: Sequence labeling
 - We focus on sequence labeling with Hidden Markov Models and some simple deep learning based models. Our task is part-of-speech tagging on English and Norwegian from the Universal Dependencies dataset. We will cover the Viterbi algorithm which could require a little bit prior knowledge of dynamic programming.
- Project 3: Dependency parsing
 - We will implement a transition-based dependency parser. The algorithm would be new and specific to the dependency parsing problem, but the underlying building blocks of the method are still some neural network modules covered in P1 and P2.

Homework submission

- Submit via Gitlab
 - We will pull your code for submission (with an assignment tag) and check the commit time.
 - A detailed grading rubric would be specified in the main Jupyter notebook of each assignment.
- 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!

Communications with instructors

- You should be able to see yourselves be added to the Ed discussion board of CSE 447 / CSE M 547 22 sp. **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.

Class participation

- Lectures and homework assignments complement each other
- Lecture materials are broader
- Homework assignments will go deeper into three important topics
- Try to attend the lectures
- Quizzes are designed to encourage you to do so
- But if you miss a lecture – you can watch a recording, or read assigned book chapters
- Participate in class discussions, 10% 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!

Office hours

- Yulia – Thu 3:30 - 4:30pm @ zoom <https://washington.zoom.us/j/92241001365> (and by appointment)
- Thai - Tues 2:00 - 3:00pm @ Zoom (Zoom ID on Canvas)
- Han - Mon 11am-12pm @ zoom (<https://washington.zoom.us/my/xhan77>)
- Kaiser - Mon 12:20pm - 1:20pm @ zoom (Zoom ID on Canvas)
- Ivy - Wed 4:30-5:30pm (<https://washington.zoom.us/my/ivyguo>)
- Leroy - Fri 11:30am - 12:30pm (<https://washington.zoom.us/my/lrywnq>)

Quizzes

- 8 quizzes, students can drop 3
- Each quiz has 5 simple multiple-choice questions, autograded
- Quizzes are on Canvas, open during the lecture time
- Quiz time - 10 minutes in the beginning of the class
- Starting from the 3rd week
- Grading on 5 best quizzes, 2% each

More course logistics

We care that you learn!

Your questions are always welcome.

Questions?

<https://courses.cs.washington.edu/courses/cse447/22sp/>