

Computational Social Science

Lucy Li

University of Washington

CSE 447

11/21/25

natural language
processing

natural **language**
processing

natural **language** = social & cultural data
processing

**natural language
processing**

**computational social
science**



**natural language
processing**

**computational social
science**



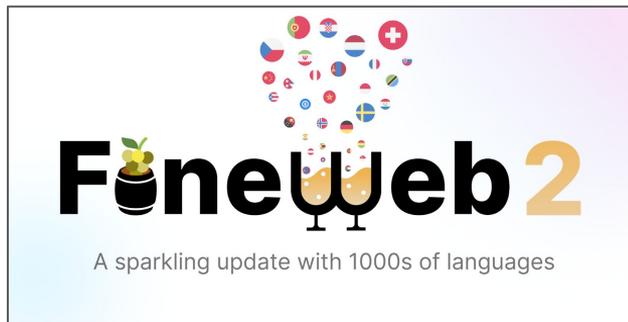
Language = social and cultural data




Edit Profile ...

Lucy Li
 @lucy3.bsky.social
 5.6K followers 637 following 324 posts

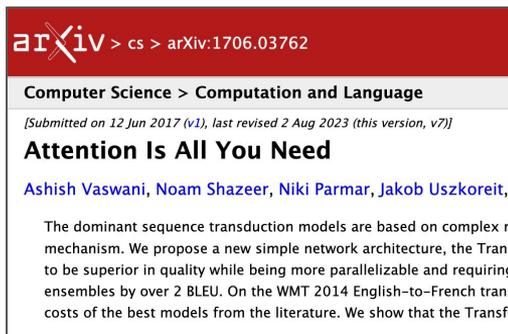
Postdoc at UW NLP 📍 #NLProc, computational social science, cultural analytics, responsible AI. she/her. Previously at Berkeley, Ai2, MSR, Stanford. Incoming assistant prof at Wisconsin CS. lucy3.github.io/prospective-students.html



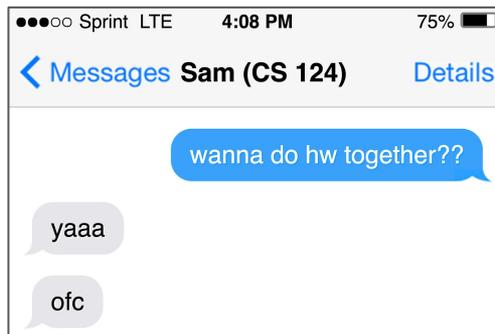
FoneWeb 2
 A sparkling update with 1000s of languages



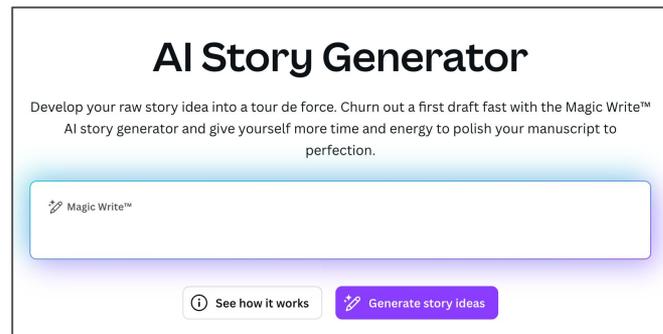
Natural Instructions



arXiv > cs > arXiv:1706.03762
Computer Science > Computation and Language
[Submitted on 12 Jun 2017 (v1), last revised 2 Aug 2023 (this version, v7)]
Attention Is All You Need
 Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit,
 The dominant sequence transduction models are based on complex re-
 mechanism. We propose a new simple network architecture, the Trans-
 to be superior in quality while being more parallelizable and requiring
 ensembles by over 2 BLEU. On the WMT 2014 English-to-French trans-
 costs of the best models from the literature. We show that the Transfo



Sprint LTE 4:08 PM 75%
 Messages Sam (CS 124) Details
 wanna do hw together??
 yaaa
 ofc



AI Story Generator
 Develop your raw story idea into a tour de force. Churn out a first draft fast with the Magic Write™ AI story generator and give yourself more time and energy to polish your manuscript to perfection.

Magic Write™
 See how it works Generate story ideas

Language = social and cultural data

Lucy Li
@lucy3.bsky.social
5.6K followers 637 following 324 posts

Postdoc at UW NLP 📍 #NLProc, computational social science, cultural analytics, responsible AI. she/her. Previously at Berkeley, Ai2, MSR, Stanford. Incoming assistant prof at Wisconsin CS. lucy3.github.io/prospective-students.html

Finetune 2

Natural Instructions

“encoder-decoder”, “attention”
“convolutional”, “transformer”
→ AI researcher

arXiv > cs > arXiv:1706.03762

Computer Science > Computation and Language

[Submitted on 12 Jun 2017 (v1), last revised 2 Aug 2023 (this version, v7)]

Attention Is All You Need

Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit,

The dominant sequence transduction models are based on complex recurrent neural network architectures. We propose a new simple network architecture, the Transformer, to be superior in quality while being more parallelizable and requiring less computation. On the WMT 2014 English-to-French translation task, the Transformer achieves a BLEU score of 38.4, which is on par with the best models from the literature. We show that the Transformer is more parallelizable and requires less computation than the best models from the literature. We show that the Transformer is more parallelizable and requires less computation than the best models from the literature.

wanna do hw together??

yaaa

ofc

AI Story Generator

Turn your idea into a tour de force. Churn out a first draft fast with the Magic Write™ AI story generator and give yourself more time and energy to polish your manuscript to perfection.

Magic Write™

See how it works Generate story ideas

Language = social and cultural data

Lucy Li
@lucy3.bsky.social
5.6K followers 637 following 324 posts

Postdoc at UW NLP [📍](#) #NLProc, computational social science, cultural analytics, responsible AI. she/her. Previously at Berkeley, Ai2, MSR, Stanford. Incoming assistant prof at Wisconsin CS. lucy3.github.io/prospective-students.html

Fonelo
A sparkling update

Natural

Abstractive translation:
"Sam is willing to do homework"
+ you're **likely friends**

arXiv > cs > arXiv:1706.03762

Computer Science > Computation and Language

[Submitted on 12 Jun 2017 (v1), last revised 2 Aug 2023 (this version, v7)]

Attention Is All You Need

Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit,

The dominant sequence transduction models are based on complex re-arranging mechanism. We propose a new simple network architecture, the Transformer, to be superior in quality while being more parallelizable and requiring ensembles by over 2 BLEU. On the WMT 2014 English-to-French translation task, the Transformer achieves a BLEU score of 40.6, which is the best performance to date. We show that the Transformer is more parallelizable and requires less computation than the best models from the literature. We show that the Transformer is more parallelizable and requires less computation than the best models from the literature. We show that the Transformer is more parallelizable and requires less computation than the best models from the literature.

Sprint LTE 4:08 PM 75%

Messages Sam (CS 124) Details

wanna do hw together??

yaaa

ofc

AI Story Generator

Develop your raw story idea into a tour de force. Churn out a first draft fast with the Magic Write™ AI story generator and give yourself more time and energy to polish your manuscript to perfection.

Magic Write™

See how it works Generate story ideas

Language = social and cultural data

Lucy Li
@lucy3.bsky.social
5.6K followers 637 following 324 posts
Postdoc at UW NLP. #NLProc, computational linguistics, responsible AI. she/her. Previously at UW-Madison. Incoming assistant prof at Wisconsin CS. lucy3@cs.wisc.edu. lucy3.students.html

FineWeb2
with 1000s of languages

Natural Instructions

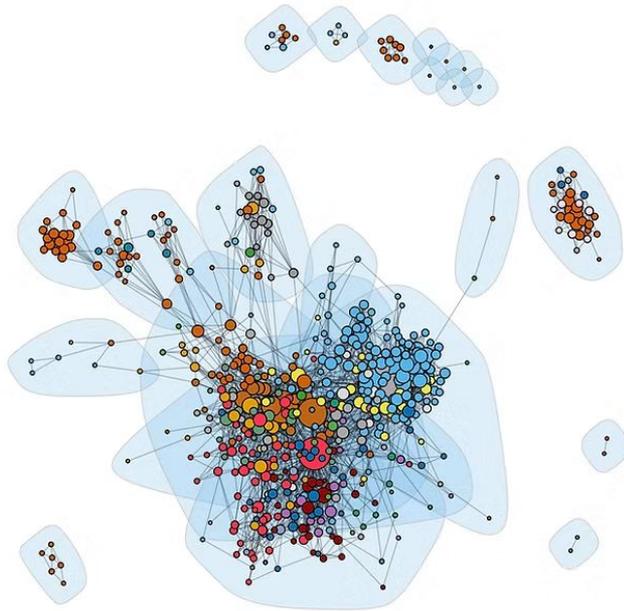
“She was a **librarian** by day, **quiet** and **unassuming**, her life a patchwork of ancient books and whispered stories”

arXiv > cs > arXiv:1706.03762
Computer Science > Computation and Language
[Submitted on 12 Jun 2017 (v1), last revised 2 Aug 2023 (this version, v7)]
Attention Is All You Need
Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit,
The dominant sequence transduction models are based on complex recurrent neural network mechanisms. We propose a new simple network architecture, the Transformer, to be superior in quality while being more parallelizable and requiring less data. On the WMT 2014 English-to-French translation task, the Transformer achieves a BLEU score of 38.4, which is on par with the best models from the literature. We show that the Transformer is superior in quality while being more parallelizable and requiring less data.

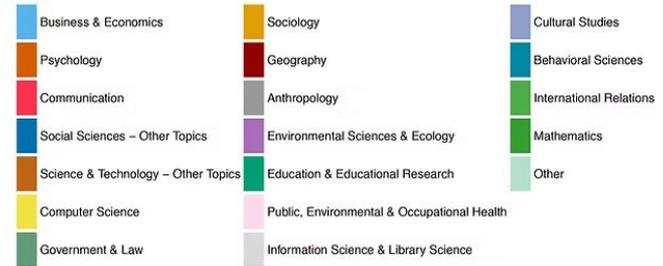
wanna do hw together??
yaaa
ofc

AI Story Generator
Develop your raw story idea into a tour de force. Churn out a first draft fast with the Magic Write™ AI story generator and give yourself more time and energy to polish your manuscript to perfection.
Magic Write™
See how it works Generate story ideas

NLP for **social science**

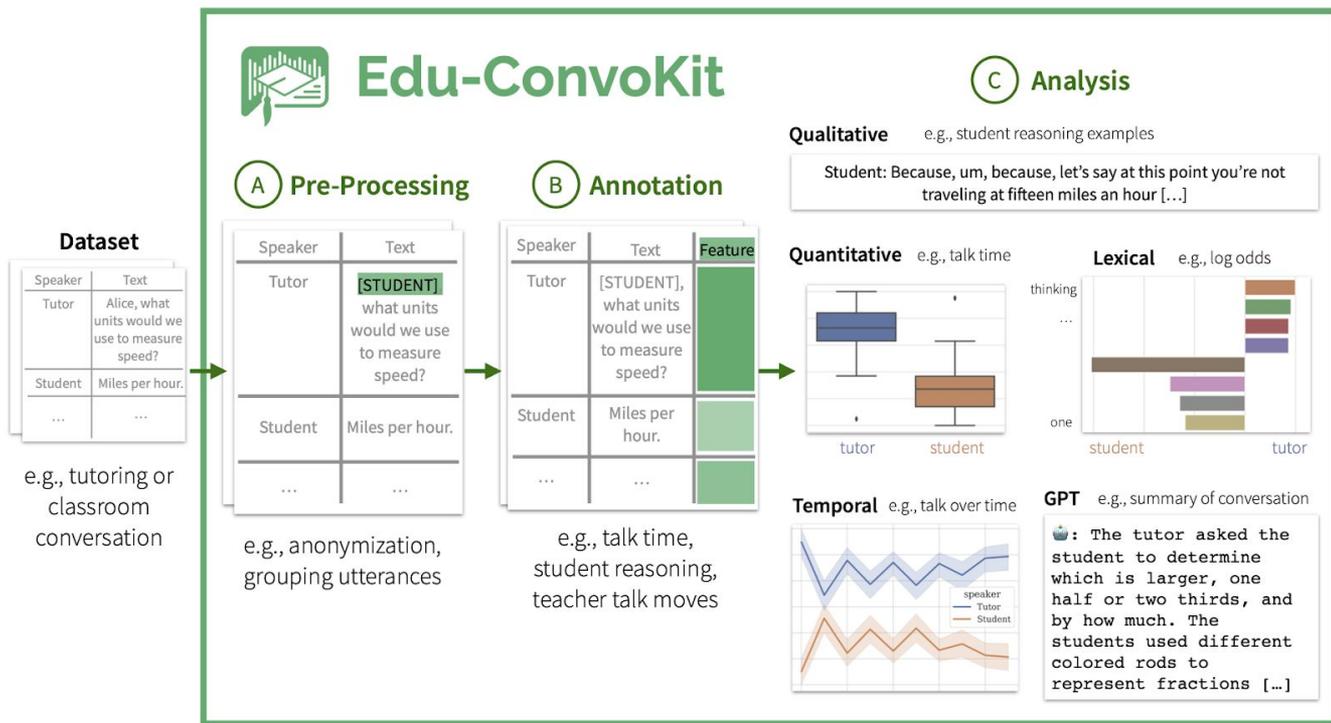


Nodes colored by first-listed discipline



[Chris Bail's "Mapping Computational Social Science"](#)

Examples: NLP for education



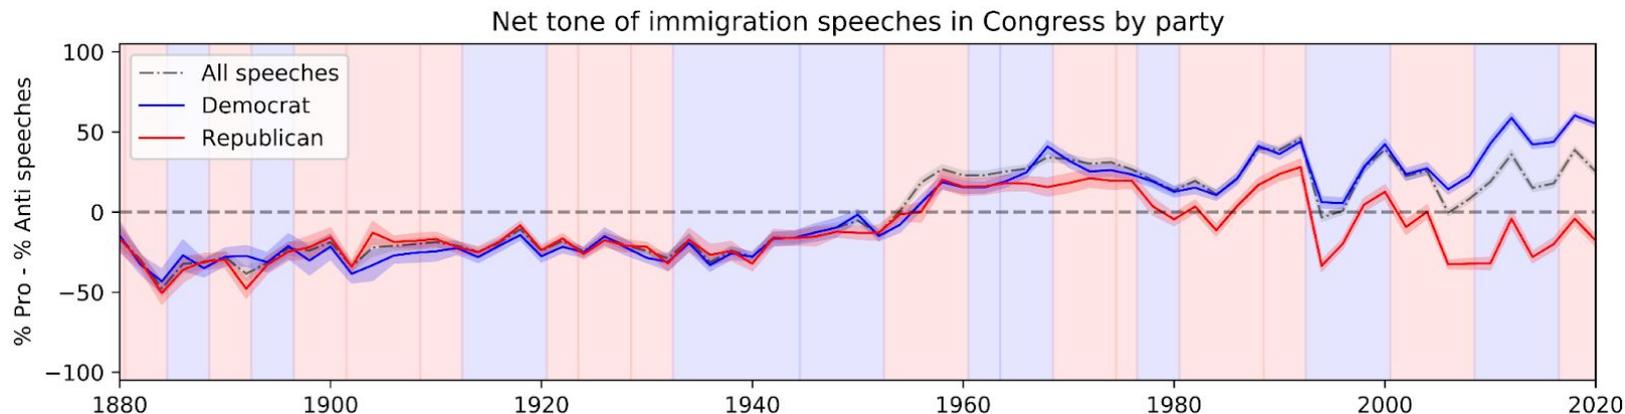
"Do you agree with Juan that the answer is 7/10?"

→ Getting students to relate to another's ideas

"Why could I argue that the slope should be increasing?"

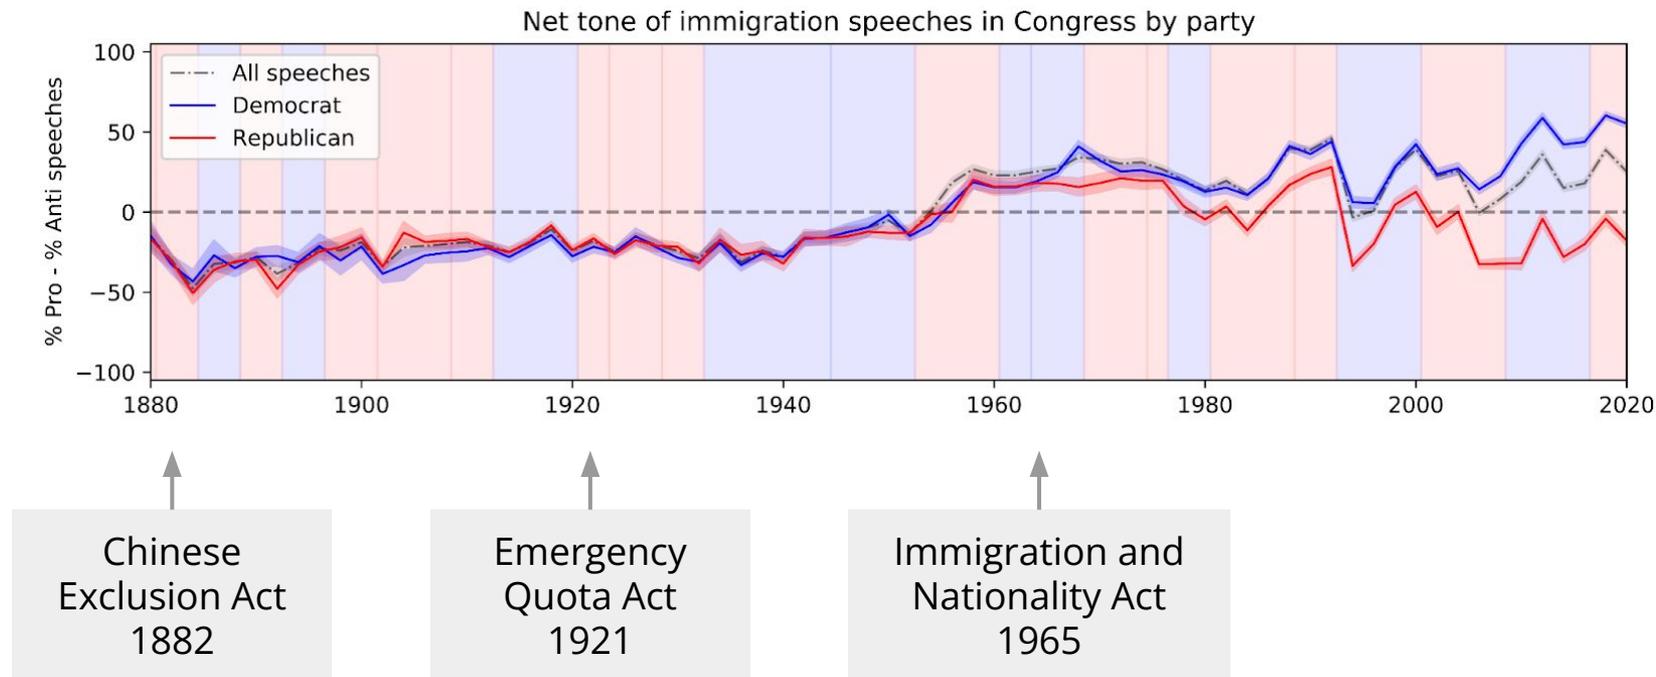
→ Pressing for reasoning

Example: NLP for history & political science

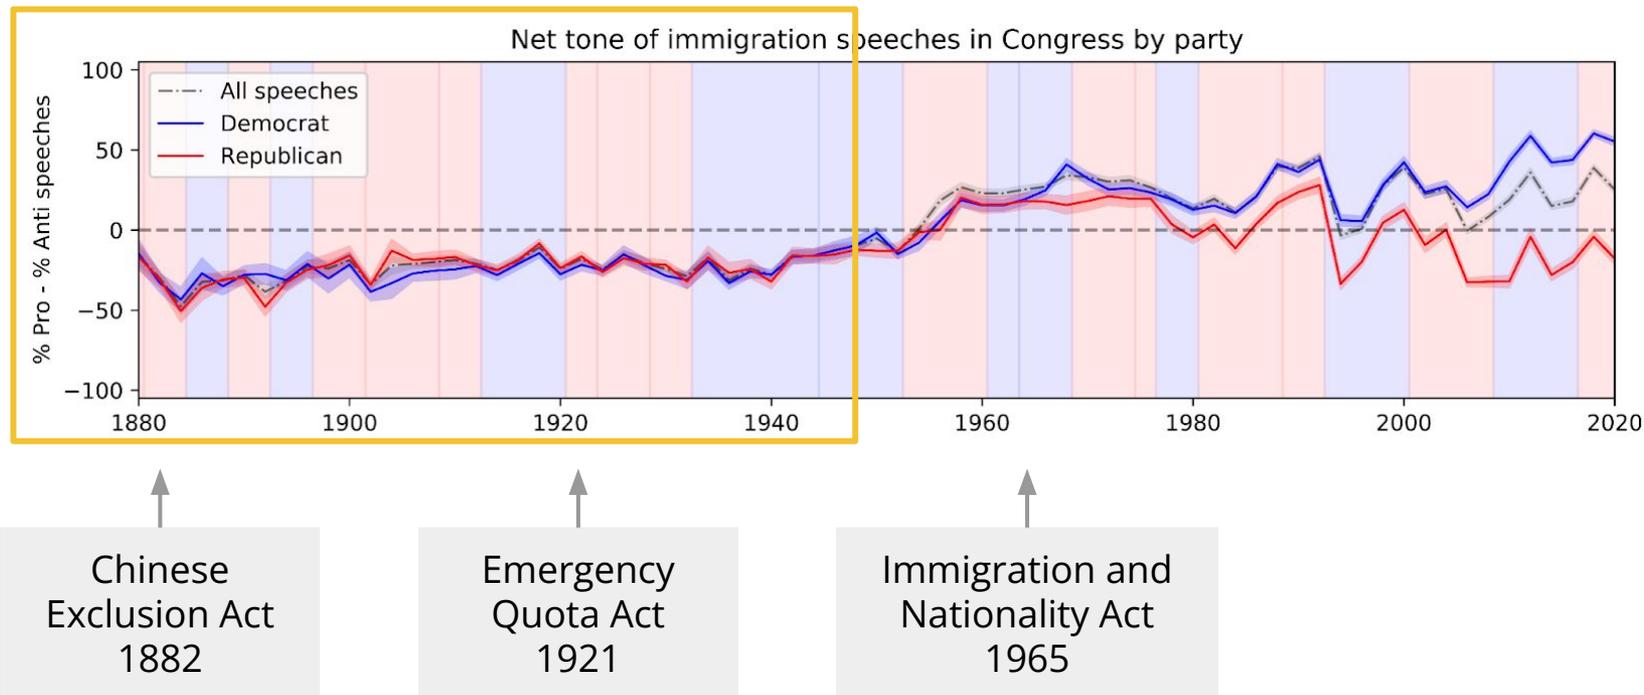


1. Adapt RoBERTa to U.S. congressional speeches using self-supervised finetuning
2. Then, finetune on human annotators' labels of tone

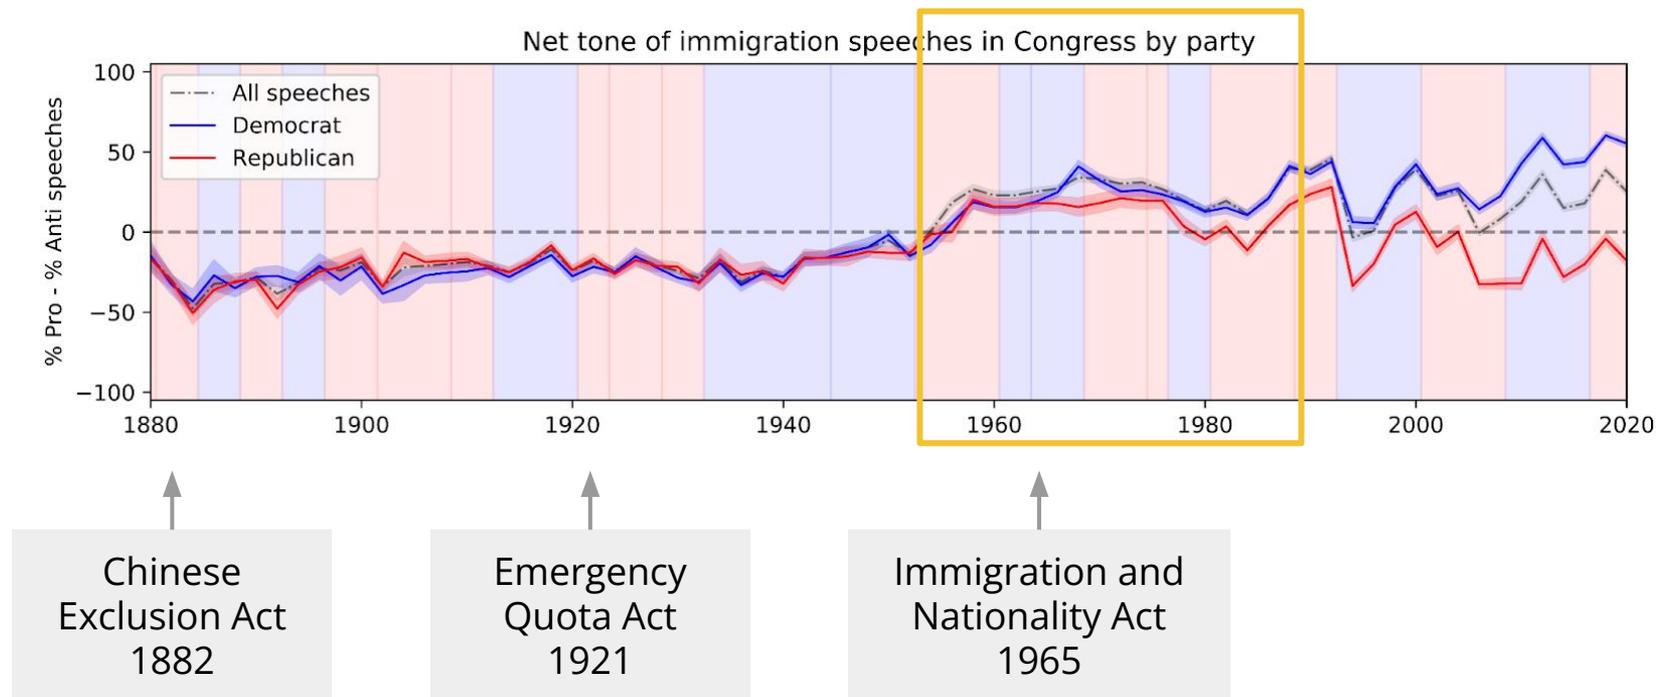
Example: NLP for history & political science



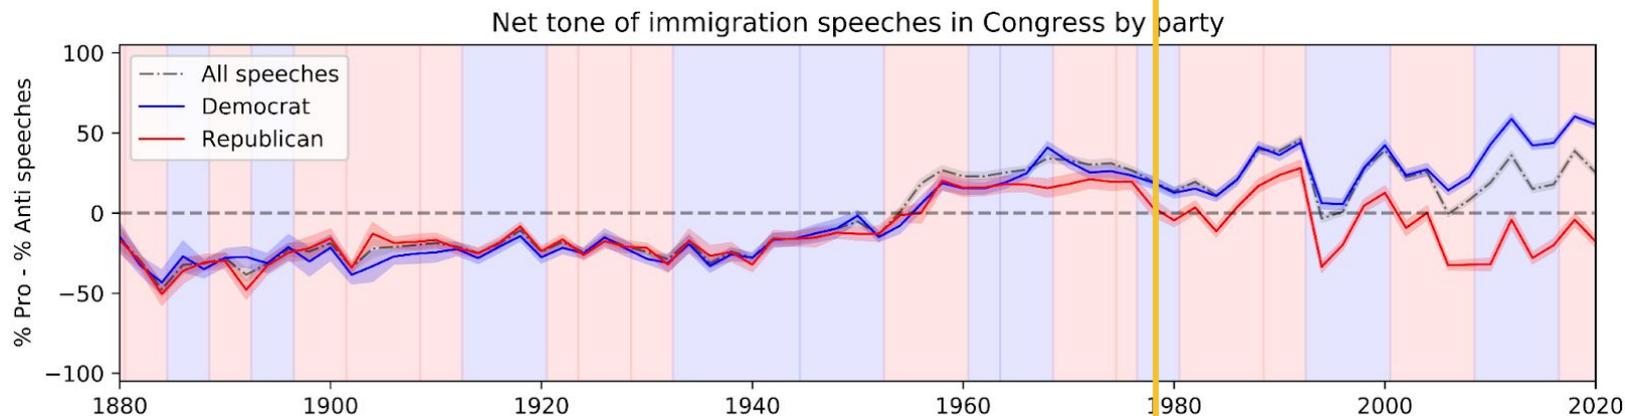
Example: NLP for history & political science



Example: NLP for history & political science



Example: NLP for history & political science



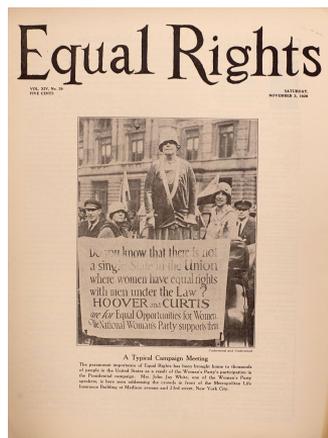
Chinese
Exclusion Act
1882

Emergency
Quota Act
1921

Immigration and
Nationality Act
1965

Example: NLP for history & sociology

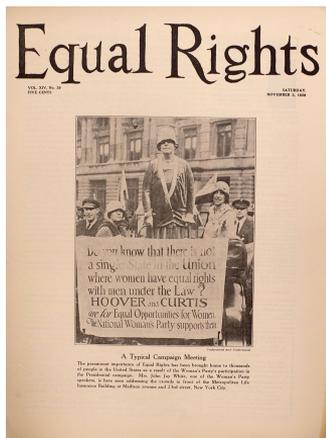
NLP for measuring **collective memories**



primary source materials → secondary source

Example: NLP for history & sociology

NLP for measuring **collective memories**



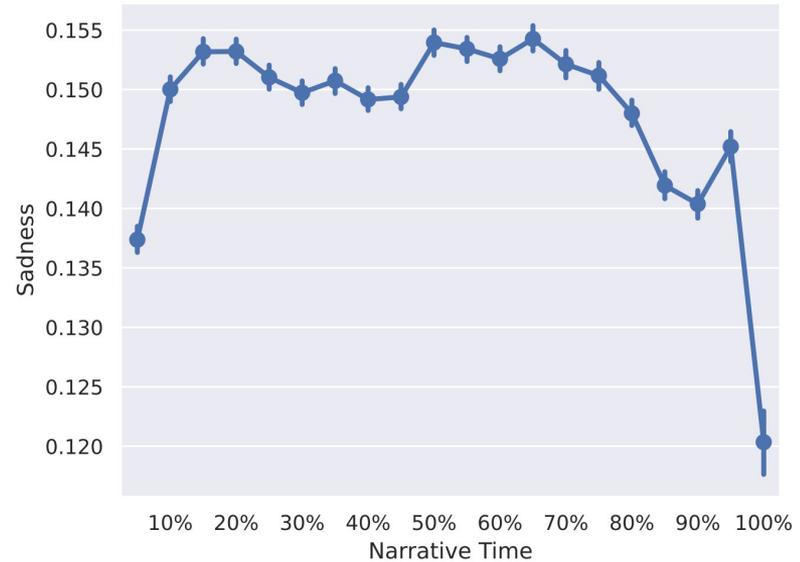
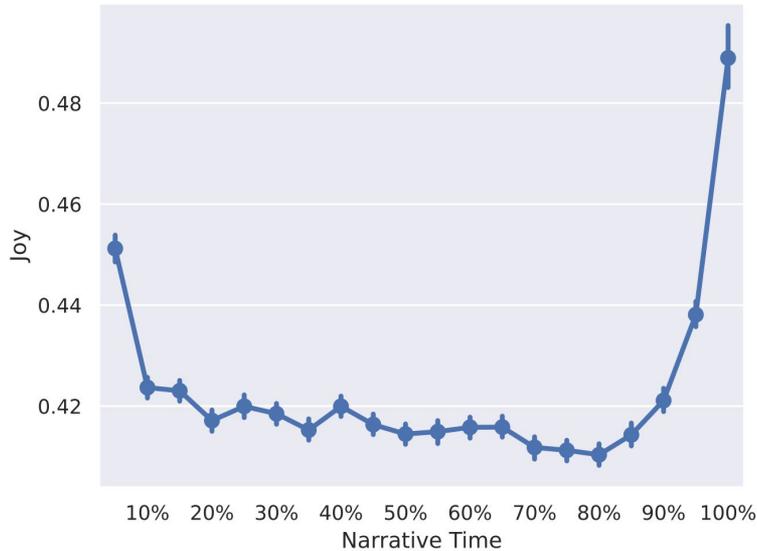
primary source materials → secondary source

What frequent phrases are **missing** on Wikipedia?

Writings of Black Women Suffragists	Equal Rights Journal	National Consumers' League
frederick douglass center	edith houghton hooker	minimum wage boards
war camp community service	alva belmont house	pure food law
ida joyce jackson	lucretia mott amendment	minimum wage bill
phyllis wheatley home	equal rights treaty	forty-eight hour week
colored club women	special labor laws	industrial home work
jim crow car	hague nationality convention	federal child labor amendment
political equality association	equal nationality treaty	state factory inspector
miner normal school	equal nationality bill	continental waist company
high moral standards	equal nationality rights	eight hours law
domestic service occupations	jury service bill	misbranded malt liquors

Example: NLP for media studies

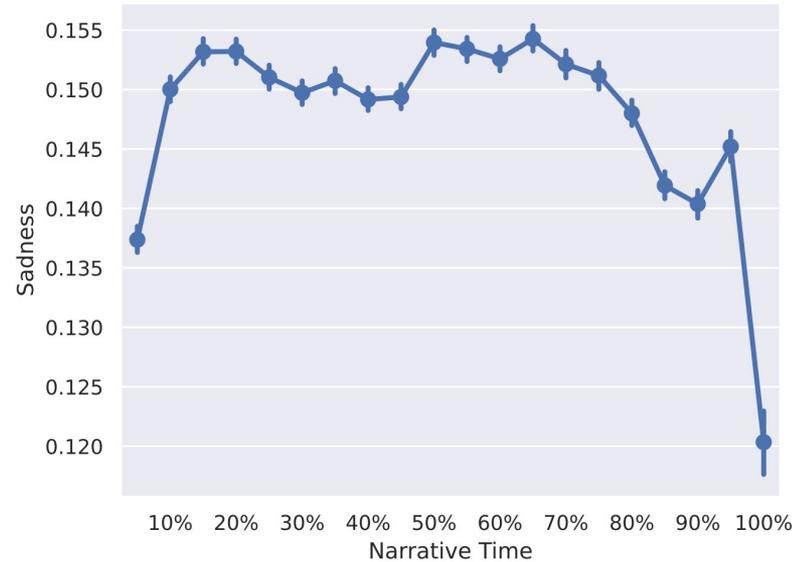
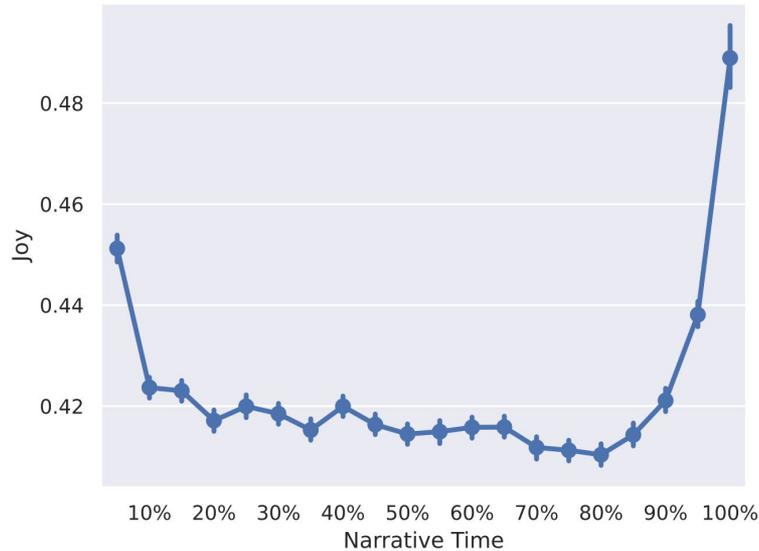
Automated speech emotion recognition on top box office films, 1980-2022



[Naitian & Bamman 2024](#)

Example: NLP for media studies

cultural analytics = "distant reading"



Example: NLP for media studies

cultural analytics = “distant reading”

“... cultural analytics is a multi-disciplinary project—a bridge between the **humanities** and **quantitative social science**, belonging equally to both.”

- [Ted Underwood 2018](#)



natural language processing

Example: NLP for *social* media studies

Who builds and maintains an online community?

Analyze **text** from user sessions on cancer discussion board [\(Yang et al. 2019\)](#)

Supervised classification of **behavioral features** e.g. “seeking informational support”

Unsupervised clustering of these features produced 11 distinct **social roles** of community members

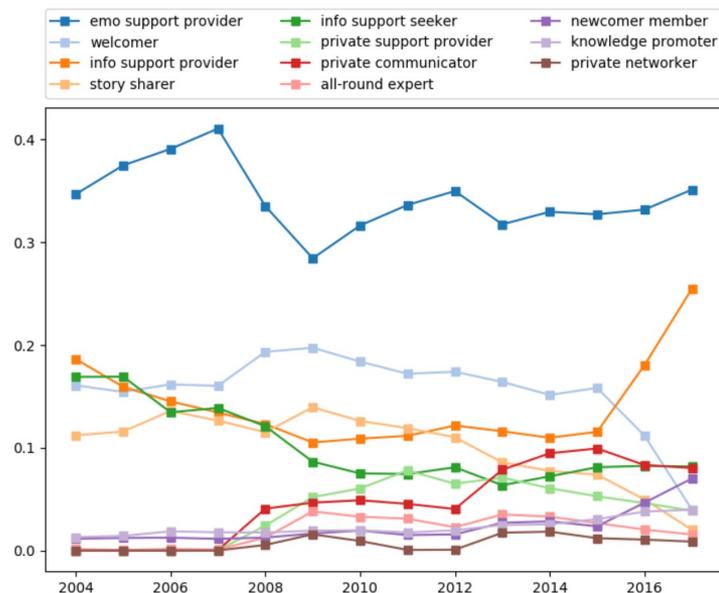


Figure 1: The percentage of different role occupations

~5 min activity

<https://bit.ly/4r3XN7o>

If you were a computational social science / cultural analytics researcher, what kind of data would you want to study, assuming there is a route for obtaining it ethically/legally?

Add an idea and/or put “+” next to an existing idea if you’re interested too (you have an unlimited # of agreement votes)

Note: the data can be human-AI interaction data too! So, consider questions you might ask about AI use in different settings 🤔

Most upvoted ideas in Berkeley's NLP class

1. Reddit ++++++
2. Instagram reel comments ++++++
3. Memes & "Italian brainrot" ++++++
4. Movies ++++++
5. Classroom conversations ++++++

Now, moving more broadly...

computational social science != computer science + social data

- Hanna Wallach, 2018, Communications of the ACM

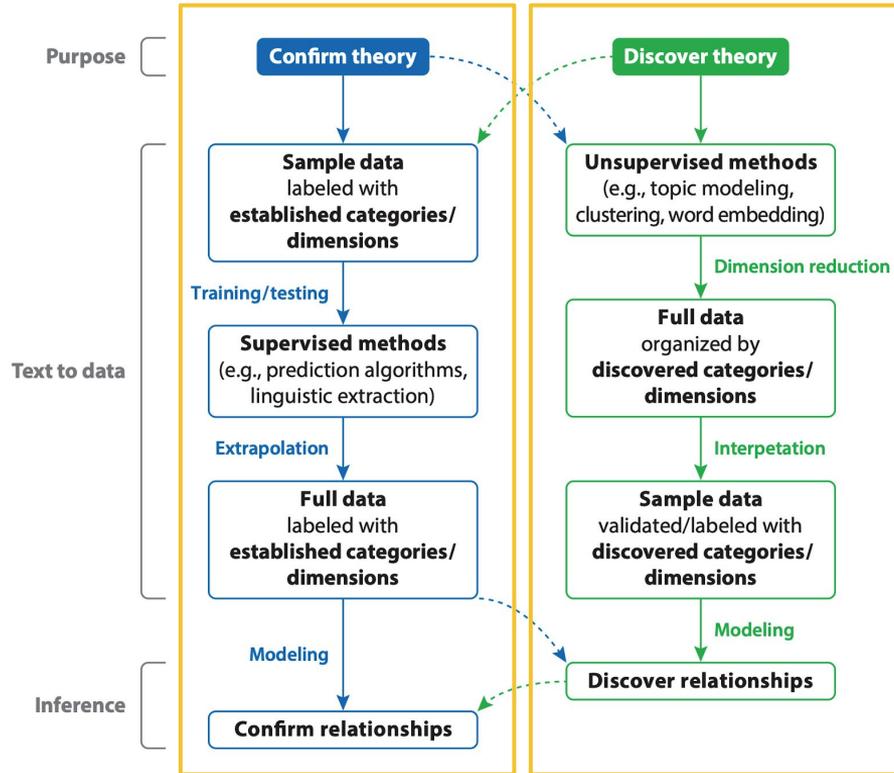
Computer science:
What? When?

- Predictive analyses
- Short term, platform-specific
- Example RQ:
 - Can we classify whether users will make a purchase?

Social science:
Why? How?

- Explanation, interpretability
- Long term, big picture, multiple data sources
- Example RQ:
 - Do tweets influence on-the-ground political protests?

Text analysis & social theory



Operationalization

Measures, or indicators, for concepts of interest, or the process of turning abstract concepts (e.g. “self-esteem”) into something quantifiable (e.g. counts of keywords)

Construct validity: *Do our measurements measure what we want to measure?*

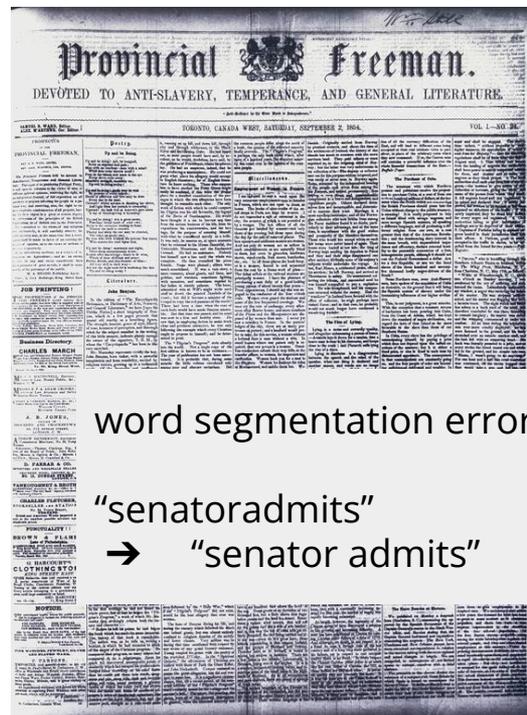
Internal validity: *Does our analysis correctly lead our measurements to our conclusions?*

External validity: *Do our findings generalize to other situations?*

Challenges of working with social data

Data quality

- Sparsity, noise, redundancy



word segmentation errors

“senatoradmits”
→ “senator admits”

Challenges of working with social data

Data quality

- Sparsity, noise, redundancy

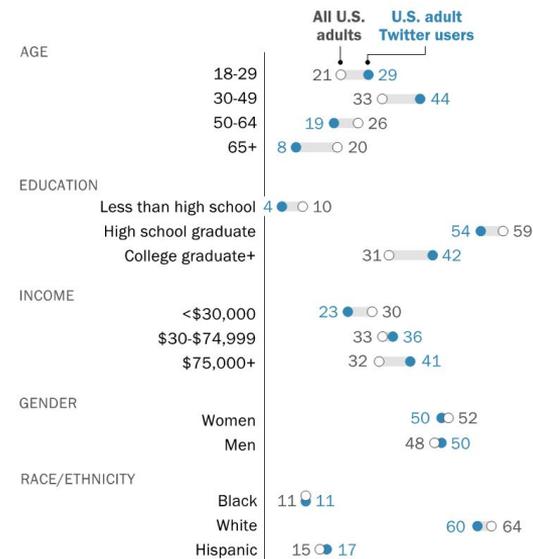
Population biases

- Demographics & characteristics of users

Pew Research

Twitter users are younger, more highly educated and wealthier than general public

% of ____ who are ...



Challenges of working with social data

Data quality

- Sparsity, noise, redundancy

Population biases

- Demographics & characteristics of users

Behavioral biases

- Distortions in user behavior, interaction

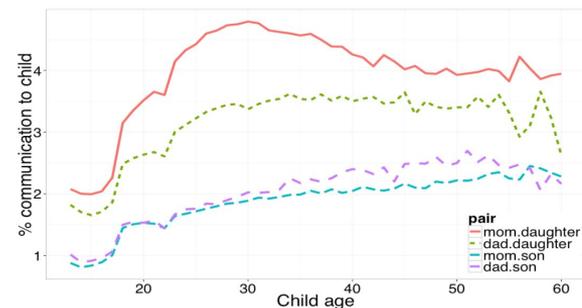


Figure 4a. Percent of a parent's posts targeted at his or her child.

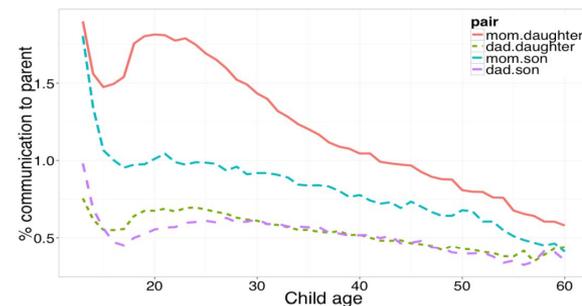


Figure 4b. Percent of a child's posts targeted at his or her parents.

["Families on Facebook" Burke et al. 2013](#)

Challenges of working with social data

Content production biases

- Linguistic variation in content



Predicting political orientation using tweets' hashtags

1. *"California Governor"* ✓
2. *"Republican and loving it!"* ✓
3. *"Mom of four"* ✗

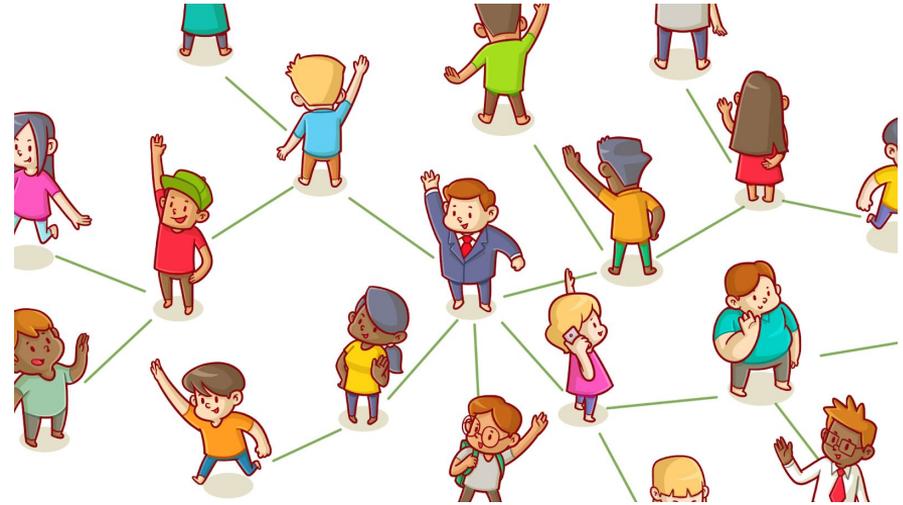
Challenges of working with social data

Content production biases

- Linguistic variation in content

Linking biases

- Implications of how connections are defined



Challenges of working with social data

Content production biases

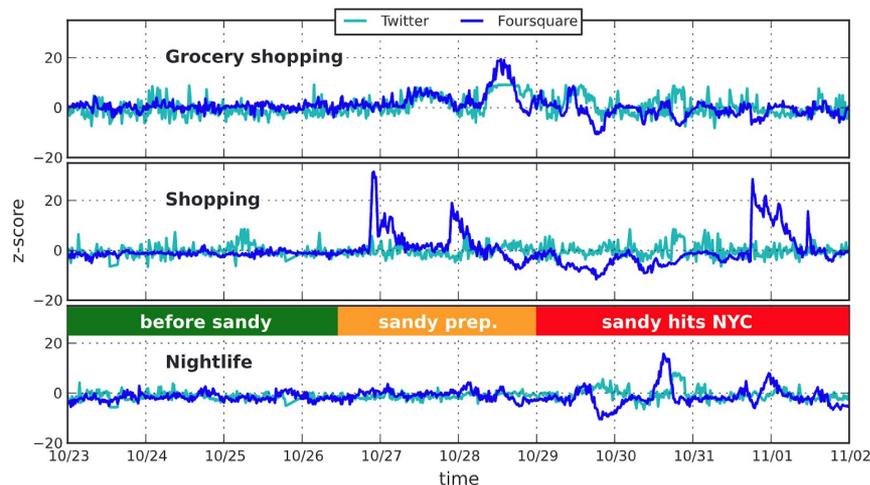
- Linguistic variation in content

Linking biases

- Implications of how connections are defined

Temporal biases

- Sudden events, seasonal changes



[Grinberg et al. \(2013\)](#) "Extracting Diurnal Patterns of Real World Activity from Social Media"

**natural language
processing**

**computational social
science**



Back in my day ...

Use machine learning models, e.g. logistic regression, SVMs, random forest classifiers, to

1. scale up “annotations”, or predictions
2. identify *features* highly related to a social phenomena

Example: predicting random acts of 🍕 on Reddit

[Althoff et al. 2014](#)

Coefficient	Estimate	SE
Intercept	-2.02***	0.14
Community Age (Decile)	-0.13***	0.01
First Half of Month (Binary)	0.22**	0.08
Gratitude (Binary)	0.27**	0.08
Including Image (Binary)	0.81***	0.17
Reciprocity (Binary)	0.32**	0.10
Strong Positive Sentiment (Binary)	0.14	0.08
Strong Negative Sentiment (Binary)	-0.07	0.08
Length (in 100 Words)	0.30***	0.05
Karma (Decile)	0.13***	0.02
Posted in RAOP before (Binary)	1.34***	0.16
Narrative Craving (Binary)	-0.34***	0.09
Narrative Family (Binary)	0.22*	0.09
Narrative Job (Binary)	0.26**	0.09
Narrative Money (Binary)	0.19**	0.08
Narrative Student (Binary)	0.09	0.09

*** $p < 0.001$, ** $p < 0.01$, * $p < 0.05$

from topic modeling

And then ...

[ARTICLE]

Which of the following leanings would a political scientist say that the above article has?

A: Liberal

B: Conservative

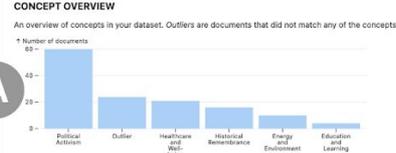
C: Neutral

Model Data	Baselines		FLAN-T5				FLAN		text-001			text-002	text-003	Chat		
	Rand	Finetune	Small	Base	Large	XL	XXL	UL2	Ada	Babb.	Curie	Dav.	Davinci	Davinci	GPT3.5	GPT4
Utterance Level Tasks																
Dialect	3.3	3.0	0.2	4.5	23.4	24.8	30.3	32.9	0.5	0.5	1.2	9.1	17.1	14.7	11.7	23.2
Emotion	16.7	71.6	19.8	63.8	69.7	65.7	66.2	70.8	6.4	4.9	6.6	19.7	36.8	44.0	47.1	50.6
Figurative	25.0	99.2	16.6	23.2	18.0	32.2	53.2	62.3	10.0	15.2	10.0	19.4	45.6	57.8	48.6	17.5
Humor	49.5	73.1	51.8	37.1	54.9	56.9	29.9	56.8	38.7	33.3	34.7	29.2	29.7	33.0	43.3	61.3
Ideology	33.3	64.8	18.6	23.7	43.0	47.6	53.1	46.4	39.7	25.1	25.2	23.1	46.0	46.8	43.1	60.0
Impl. Hate	16.7	62.5	7.4	14.4	7.2	32.3	29.6	32.0	7.1	7.8	4.9	9.2	18.4	19.2	16.3	3.7
Misinfo	50.0	81.6	33.3	53.2	64.8	68.7	69.6	77.4	45.8	36.2	41.5	42.3	70.2	73.7	55.0	26.9
Persuasion	14.3	52.0	3.6	10.4	37.5	32.1	45.7	43.5	3.6	5.3	4.7	11.3	21.6	17.5	23.3	56.4
Sem. Chng.	50.0	62.3	33.5	41.0	56.9	52.0	36.3	41.6	32.8	38.9	41.3	35.7	41.9	37.4	44.2	21.2
Stance	33.3	36.1	25.2	36.6	42.2	43.2	49.1	48.1	18.1	17.7	17.2	35.6	46.4	41.3	48.0	76.0
Conversation Level Tasks																
Discourse	14.3	49.6	4.2	21.5	33.6	37.8	50.6	39.6	6.6	9.6	4.3	11.4	35.1	36.4	35.4	16.7
Empathy	33.3	71.6	16.7	16.7	22.1	21.2	35.9	34.7	24.5	17.6	27.6	16.8	16.9	17.4	22.6	6.4
Persuasion	50.0	33.3	9.2	11.0	11.3	8.4	41.8	43.1	6.9	6.7	6.7	33.3	33.3	53.9	51.7	28.6
Politeness	33.3	75.8	22.4	42.4	44.7	57.2	51.9	53.4	16.7	17.1	33.9	22.1	33.1	39.4	51.1	59.7
Power	49.5	72.7	46.6	48.0	40.8	55.6	52.6	56.9	43.1	39.8	37.5	36.9	39.2	51.9	56.5	42.0
Toxicity	50.0	64.6	43.8	40.4	42.5	43.4	34.0	48.2	41.4	34.2	33.4	34.8	41.8	46.9	31.2	55.4
Document Level Tasks																
Event Arg.	22.3	65.1	-	-	-	-	-	-	-	-	8.6	8.6	21.6	22.9	22.3	23.0
Event Det.	0.4	75.8	9.8	7.0	1.0	10.9	41.8	50.6	29.8	47.3	47.4	44.4	48.8	52.4	51.3	14.8
Ideology	33.3	85.1	24.0	19.2	28.3	29.0	42.4	38.8	22.1	26.8	18.9	21.5	42.8	43.4	44.7	51.5
Tropes	36.9	-	1.7	8.4	13.7	14.6	19.0	28.6	7.7	12.8	16.7	15.2	16.3	26.6	36.9	44.9

Language models for content analysis

LLoom WORKBENCH

Concept Overview

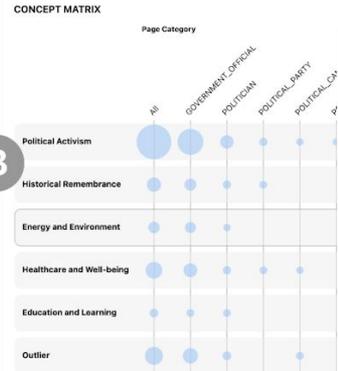


Concept Matrix

ROWS
LLoom
concepts

COLUMNS
User-
defined
slices

CIRCLES
Number of
documents



Concept/Slice Detail View

Concept: Energy and Environment

CONCEPT DETAILS

Criteria:
Is the text example related to energy policies or environmental issues?

Summary:
Advocating for American energy independence, environmental sustainability, and reducing reliance on foreign adversaries in energy production.

Concept matches: 10 documents

Representative examples:

- "WA families & small business can save thousands of dollars on energy bills through new appliances, thanks to the Democrats' Inflation Reduction Act."
- "House Republicans voted to give #BigOil open access to drill on public lands. Then, they voted against my amendment to protect Big Cypress National Preserve to stop this destructive #BigOil giveaway from harming South Floridians who need it for clean H2O, hunting, fishing and tourism."

POTENTIAL CONCEPT MATCHES

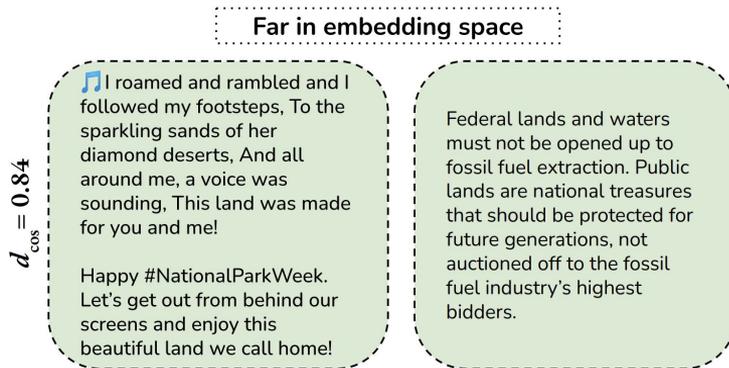
C1 CONCEPT DETAILS

C2 CONCEPT MATCHES & NON-MATCHES

id	concept	concept score	text	text bullets	score rationale	Likes	Page Category
All	Energy and Environment	1.0	Starting this year, WA families & small business can save thousands of dollars on energy bills through new appliances, thanks to the Democrats' Inflation Reduction Act. Read more about how these investments can benefit you.	<ul style="list-style-type: none"> Energy bill savings with new appliances 	The text mentions saving money on energy bills through new appliances, indicating a focus on energy efficiency.	12	GOVERNMENT_OFFICIAL
All	Energy and Environment	1.0	TUNE IN NOW I'm leading another roundtable to discuss the ongoing threats to our energy security. We will cover everything from increasing domestic energy production to reducing our dependence on foreign adversaries like China and Russia. Check it out here.	<ul style="list-style-type: none"> Discussion on energy security threats 	The text mentions discussing energy security, domestic energy production, and reducing dependence on foreign adversaries, indicating a focus on energy policies.	16	GOVERNMENT_OFFICIAL
		1.0	Congratulations to the new Speaker of the House, Kevin McCarthy! The American people expect their representatives to address the challenges facing our	<ul style="list-style-type: none"> Addressing 	The text mentions investments in		

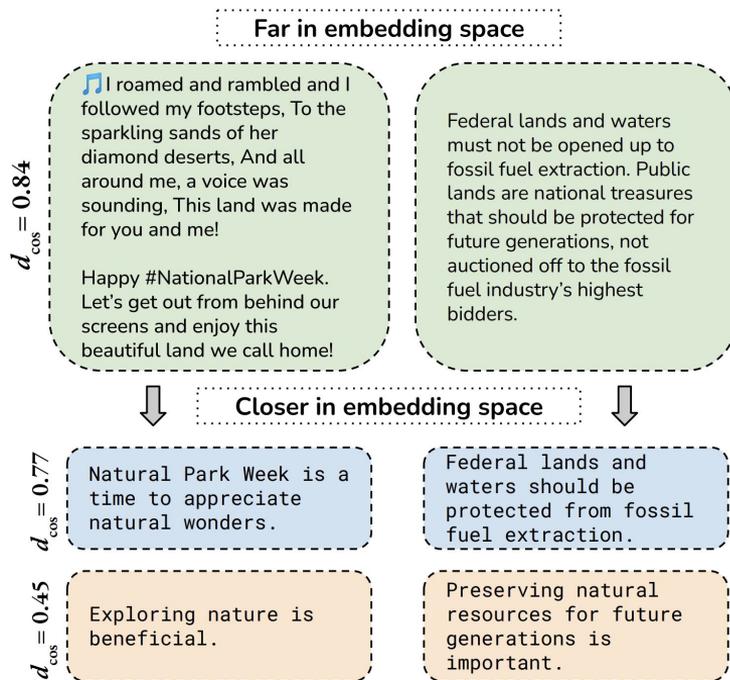
Language models for content analysis

Decompose content into
implicit propositions



Language models for content analysis

Decompose content into
implicit propositions



[Hoyle et al. 2023](#)

Language models for content analysis

Q: Should I fine-tune a small LM or prompt a large LM?

- Example: social media story detection in [Antoniak et al. 2025](#)
 - Fine-tuned RoBERTa >> zero-shot, few-shot, and chain-of-thought GPT-4
- Example: [Bamman et al. 2025](#) tested 10 cultural analytics tasks
 - Fine-tuning BERT, RoBERTa, Llama 3 8B >> 10-shot prompting LLMs if task involved a *new concept*

Sometimes what ends up mattering is *efficiency*.

“LLM-hacking”

A modern form of p -hacking

“By replicating 37 data annotation tasks from 21 published social science studies, we show that, with just a handful of prompt paraphrases, virtually anything can be presented as statistically significant.”

- [Baumann et al. 2025](#)

How to reduce the risk?

- Use highly capable models with higher task performance
- Human annotations can protect against false positives

**natural language
processing**

**computational social
science**





**natural language
processing**

**computational social
science**

Human-computer interaction

Some HCI subfields have traditionally involved a lot of intersection w/ NLP:

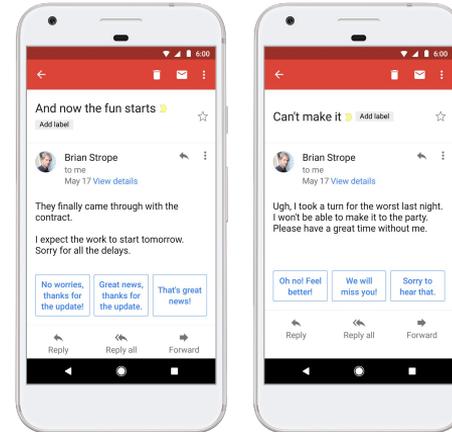
- AI fairness, ethics, & safety
- Social computing: the study of how computers and technology facilitate human behavior, interactions, and communication

Now, with language models ...

smart compose



smart reply

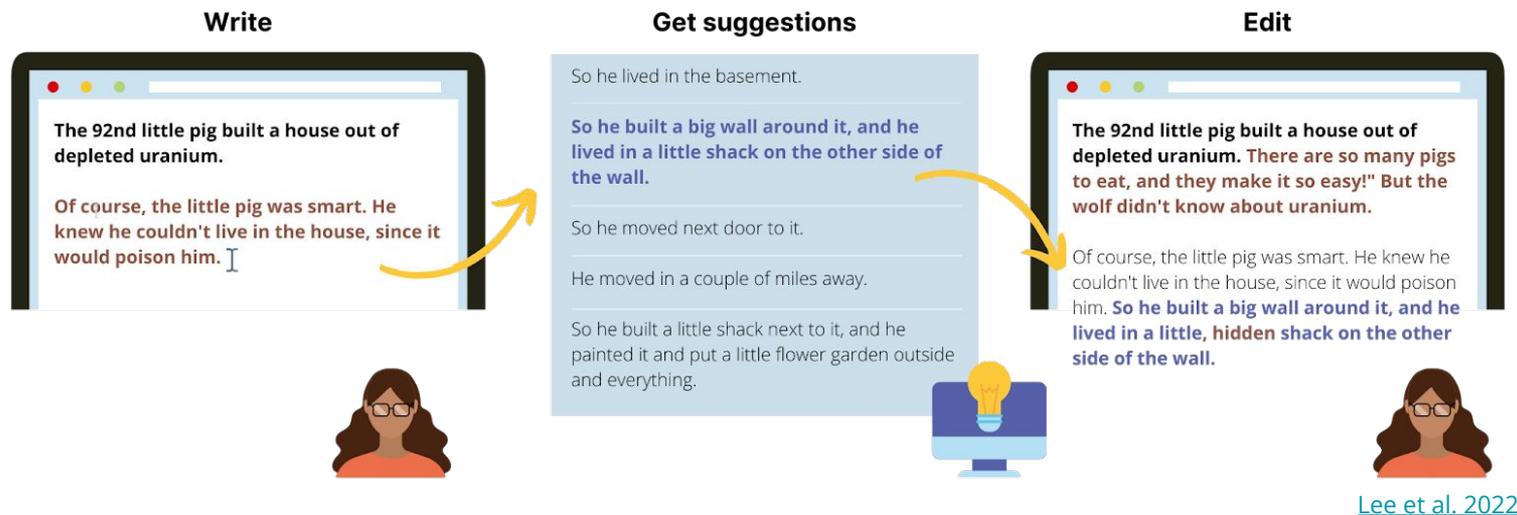


Human-computer interaction

Some HCI subfields have traditionally involved a lot of intersection w/ NLP:

- AI fairness, ethics, & safety
- Social computing: the study of how computers and technology facilitate human behavior, interactions, and communication
- **Human-AI interaction**

Human-AI interaction



Evaluation

Evaluating AI

Static benchmarks, automatic metrics.

→ *Pros:*

- ◆ Task-focused
- ◆ Allows for rapid iteration during development
- ◆ Reproducible

→ *Cons:*

- ◆ Lack of generalizability
- ◆ May not reflect real-world use

Evaluating AI + users

Interactive studies involving people/users.

→ *Pros:*

- ◆ Emphasizes real world relevance and complexity
- ◆ Could measure longer-term outcomes and impact

→ *Cons:*

- ◆ Resource-intensive
- ◆ Slower

How could we evaluate human-AI *interactions*?

Static eval: end result matters most

Interactive eval: evaluate...

- An entire **interaction trace**, or sequences of states (e.g. dialogue history) & actions (e.g. clicks)
- Third party (e.g. consistency, edit distance) *and* user's first-person experiences
- More than accuracy, e.g. enjoyment, efficiency, effort

Human-AI Language-based **Interaction Evaluation** (HALIE)

The screenshot shows a crossword puzzle interface titled "Solve a crossword puzzle with an AI Teammate!". The interface includes a crossword grid, a list of clues, and a chat window. The chat window shows a user asking for help with clue 53, "Filipino action star Fernando ___ Jr.", and the AI responding with "You which island is Honolulu on?".

ACROSS

- 45 Onsets
- 47 Cooling units, for short
- 48 Humdrum routines
- 49 Bring out from someone
- 52 "Strong Enough" singer
- 53 Filipino action star Fernando ___ Jr.
- 56 Rooster
- 57 Leading by a large margin
- 60 Shape of a 0
- 61 What incense might cover
- 62 Mario Kart dinosaur
- 63 TV star Leakes
- 64 ___ stick (bouncy toy)

DOWN

- 35 "I figured it out!"
- 41 Dry sweeping tool
- 42 Dinner-for-one platforms
- 44 ___6 movie
- 46 Union payments
- 49 "Elasticity of demand" subj.
- 50 "c3-c3-c3" sentiment
- 51 "Not if ___ help it!"
- 52 Gum up
- 53 Philippine currency
- 54 Honolulu's island
- 55 Revise
- 58 Tying-the-knot phrase
- 59 Adjective after "red" or "salba"

Chat about Clue 53 down:

Need help with a clue? Try talking to your AI teammate!
Example of ways to chat with an AI teammate can be found [here](#). (Note that the AI does not know the exact clue answers either!). Sometimes the AI might be wrong, or give an unclear response - feel free to try sending the same or a new message again! 🗨️

Chat about Clue 53 down:

You which island is Honolulu on? 12:41 PM
AI: One second... 12:41 PM
AI: Honolulu is on the island of Oahu. 12:42 PM
You: I'm a little teapot, short and 12:42 PM
AI: Just one moment! 12:42 PM
AI: stout. Here is my handle, here is my spout. When I get all steamed up, I shout, "Tip me over and pour me out!" 12:43 PM
You: what is philippine currency? 12:42 PM
AI: Please wait a few moments! 12:42 PM
AI: The Philippine currency is the Philippine peso. 12:42 PM
You: who is the dinosaur on mario kart 12:42 PM
AI: One second... 12:42 PM
AI: The dinosaur on Mario Kart is Yoshi. 12:43 PM

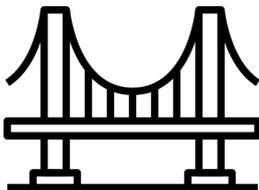
State (Puzzle, Selected clue, User letters, Dialogue history, **User input**)

Actions {Press a key to modify user input, Press the enter key to submit input, Select a **square** in the puzzle, Enter a letter into a square, Select a **clue** from the list, Finish the session}

Single-turn → multi-turn

Evaluating AI

Static benchmarks, automatic metrics.



Evaluating AI + users

Interactive studies involving people/users.

A possible bridging step: simulation or modeling of user behaviors

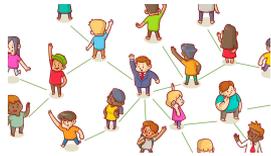
Challenges:

- What kinds of user behaviors are useful for identifying pitfalls in LM chat assistants?
- What kinds of user behaviors can we faithfully simulate?

Recap of today



NLP for education, history, political science, sociology, media studies, social media



Operationalization, measurement, validity, and data biases



Language models for content analysis



Human-AI interaction

Thank you! Questions?

Contact: lucy3li@cs.washington.edu

I will be a prof next year at the University of Wisconsin-Madison.

I am recruiting PhD students and research-based master's students!