Machine Learning for Design

Lecture 5 - Part *b*

Natural Language Processing

Previously on ML4D

Natural Language Processing

High-level understanding of the language spoken and written by humans

Also, generation (e.g., ChatGPT)

An enabler for technology like Siri or Alexa

Why natural language processing?

Fora, social media

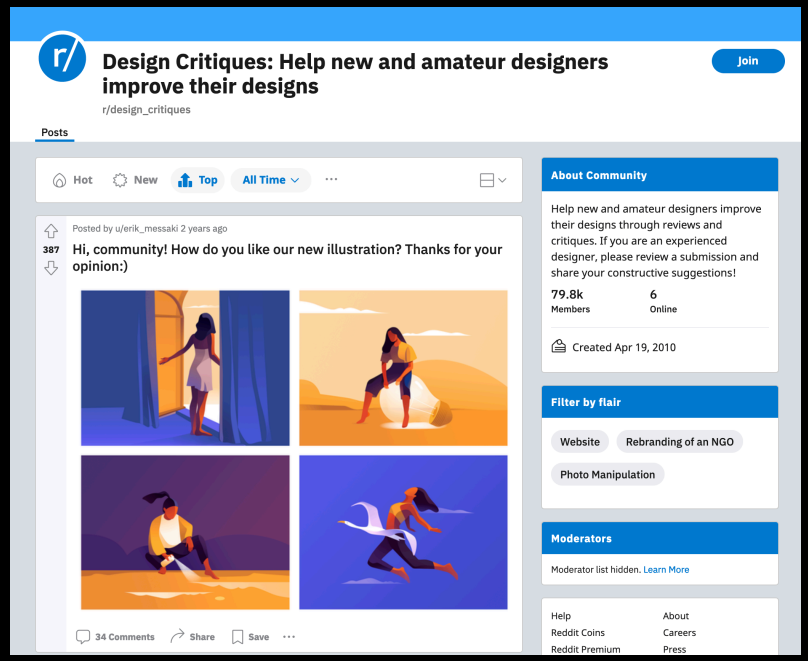

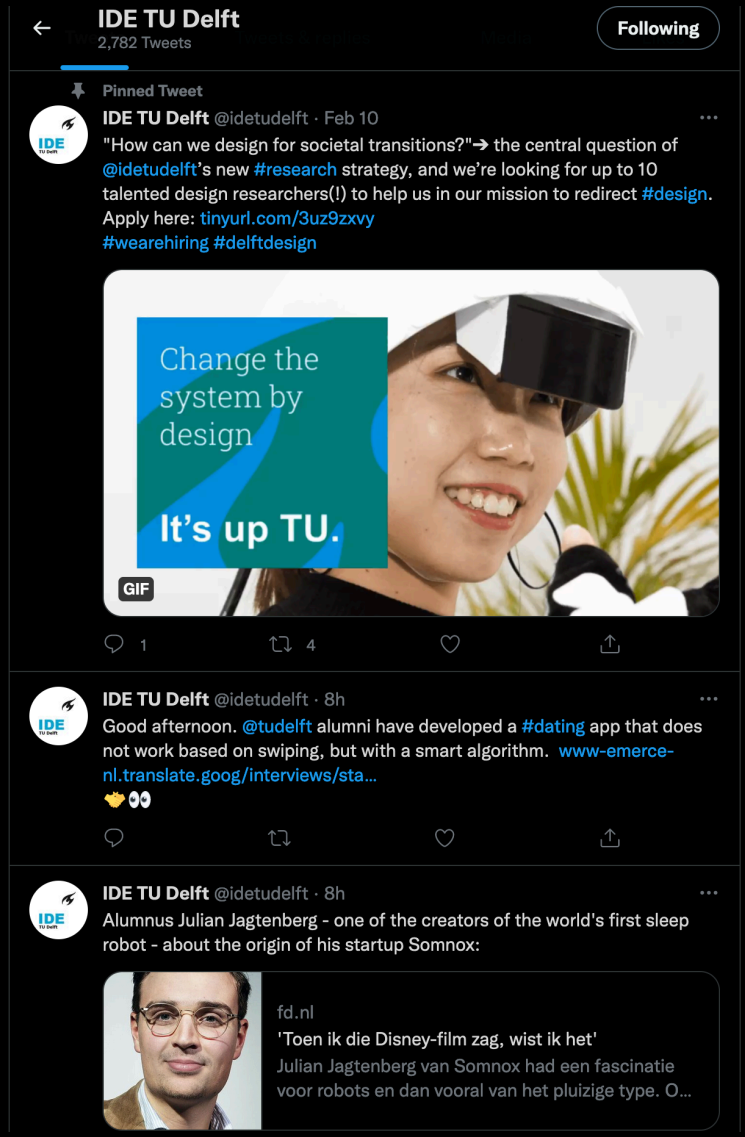

Product review

VIA AMAZON.COM

My transformation is complete

"It is day 87 and the horses have accepted me as one of their own. I have grown to understand and respect their gentle ways. Now I question everything I thought I once knew and fear I am no longer capable of following through with my primary objective. I know that those who sent me will not relent. They will send others in my place... But we will be ready." —*via Amazon/customer review/ByronicHero*.

Books

Digital, or digitised

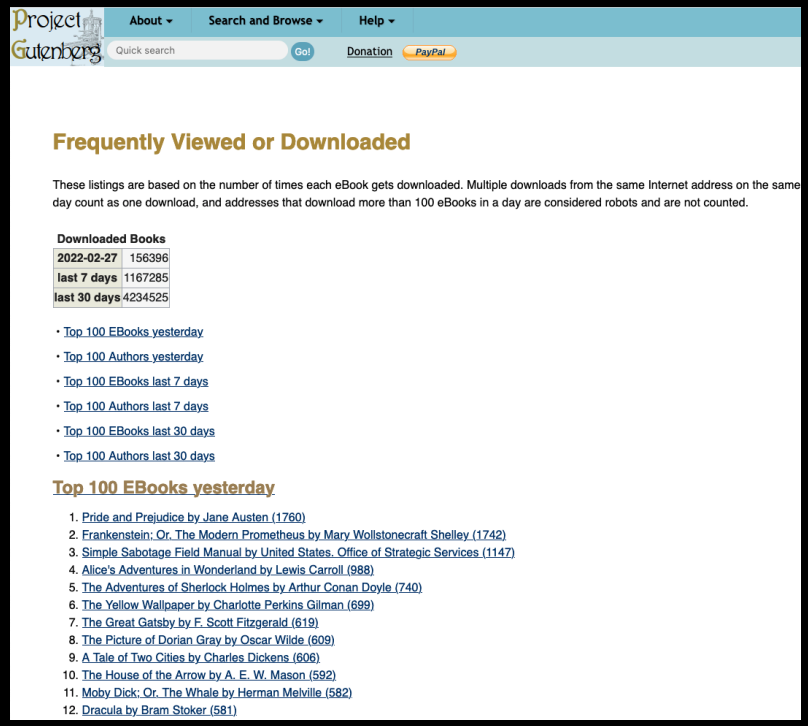

Project Gutenberg | About | Search and Browse | Help

Quick search | Go! | Donation | PayPal

Frequently Viewed or Downloaded

These listings are based on the number of times each eBook gets downloaded. Multiple downloads from the same Internet address on the same day count as one download, and addresses that download more than 100 eBooks in a day are considered robots and are not counted.

Downloaded Books	
2022-02-27	156396
last 7 days	1167285
last 30 days	4234525

- Top 100 EBooks yesterday
- Top 100 Authors yesterday
- Top 100 EBooks last 7 days
- Top 100 Authors last 7 days
- Top 100 EBooks last 30 days
- Top 100 Authors last 30 days

Top 100 EBooks yesterday

1. Pride and Prejudice by Jane Austen (1760)
2. Frankenstein; Or, The Modern Prometheus by Mary Wollstonecraft Shelley (1742)
3. Simple Sabotage Field Manual by United States. Office of Strategic Services (1147)
4. Alice's Adventures in Wonderland by Lewis Carroll (988)
5. The Adventures of Sherlock Holmes by Arthur Conan Doyle (740)
6. The Yellow Wallpaper by Charlotte Perkins Gilman (699)
7. The Great Gatsby by F. Scott Fitzgerald (619)
8. The Picture of Dorian Gray by Oscar Wilde (609)
9. A Tale of Two Cities by Charles Dickens (606)
10. The House of the Arrow by A. E. W. Mason (592)
11. Moby Dick; Or, The Whale by Herman Melville (582)
12. Dracula by Bram Stoker (581)

Interviews

Interviewee: XXX
Interviewer: XXX
Date of Interview: mm.dd.yy
Location of Interview: XXX
List of Acronyms: FP=Frank Peterson, IN=Interviewer

[Begin Transcript 00:00:10]

IN: So what was going on in your life when you joined the Marines?

FP: Well when I joined the navy, actually that was in 1950 at the age of 18. Not much other than the fact that I wanted to get away from Topeka and see what the rest of world was really all about.

IN: Um-hm.

[00:00:26]

And of course having... gone through the flight training I received my wings and commission in October of 1952. And the- one of the reasons I opted for the Marines, I knew there had never been a black pilot in the Marine Corps. So I wanted to see if I could achieve that goal, which I was able to do.

And then my first duty assignment would have been in Cherry Point, North Carolina. But I'd had enough of the South and decided I wanted to stay away from the South if I possibly could, so Headquarters Marine Corps, at my request, changed my orders to El Toro, El Toro, California.

But what I didn't realize is that I'd jumped from the frying pan into the fire because El Toro was the training base for replacement pilots in Korea. So I jumped from the frying pan into the Korean War via El Toro.

IN: I see.

[End Transcript 00:01:21]

Big Textual Data = Language at scale

- One of the largest reflections of the world, a man-made one
- Essential to better understand people, organisations, products, services, systems
  - and *their relationships*!
- Language is a proxy for human behaviour and a strong signal of individual characteristics
  - Language is always situated
  - Language is also a political instrument

- Answer questions using the Web
- Translate documents from one language to another
- Do library research; summarize
- Archive and allow access to cultural heritage
- Interact with intelligent devices
- Manage messages intelligently
- Help make informed decisions
- Follow directions given by any user
- Fix your spelling or grammar
- Grade exams
- Write poems or novels
- Listen and give advice
- Estimate public opinion
- Read everything and make predictions
- Interactively help people learn
- Help disabled people
- Help refugees/disaster victims
- Document or reinvigorate indigenous languages

What is Natural Language Processing?

- Computer using natural language as input and/or output

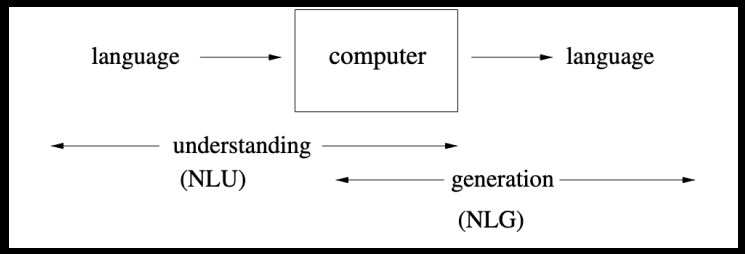

- **N**atural: human communication, unlike e.g., programming languages
- **L**anguage: signs, meanings, and a code connecting signs with their meanings
- **P**processing: computational methods to allow computers to `understand', or to generate

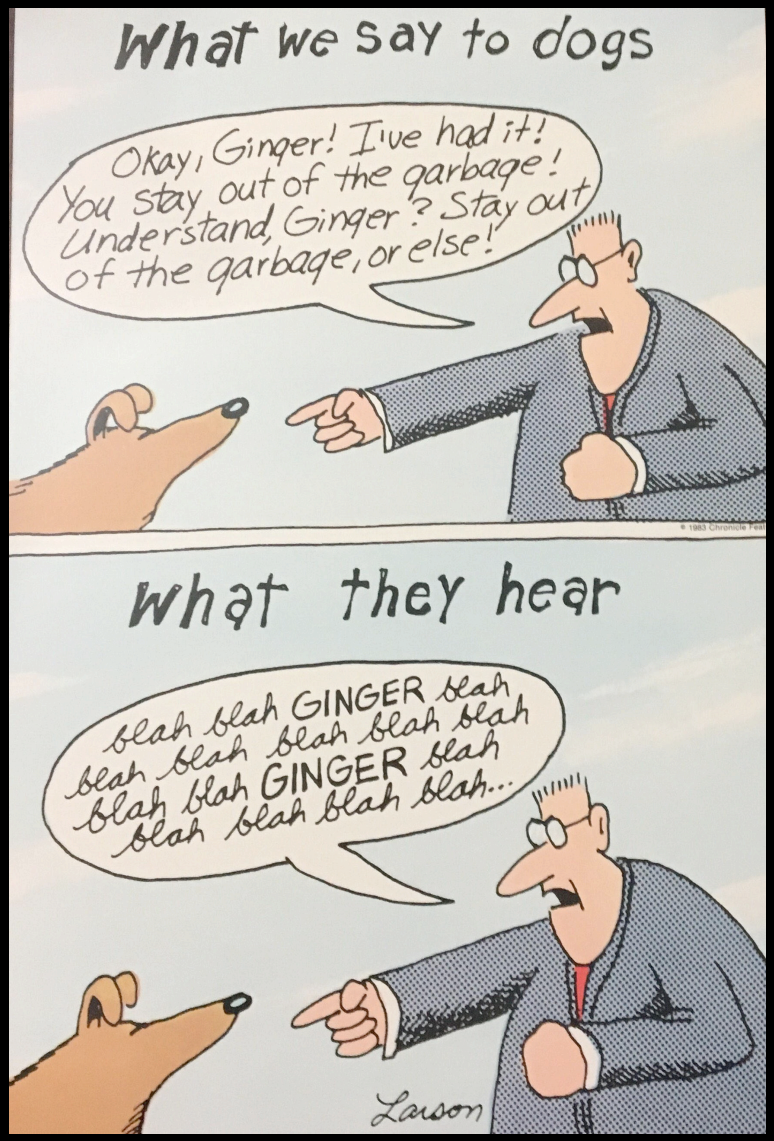

Beyond keyword matching

- Identify the **structure** and **meaning** of **words**, **sentences**, **texts** and **conversations**
- Deep understanding of broad language

Why is NLP Hard?

Human languages are **messy**, **ambiguous**, and **ever-changing**

A string may have many possible interpretations at every level

The correct resolution of the ambiguity will depend on the *intended meaning*, which is often inferable from the *context*

There is tremendous **diversity** in human languages

Languages express the same kind of meaning in different ways

Some languages express some meanings more readily/often

Knowledge Bottleneck

About language

About the world: Common sense and Reasoning

Ambiguity and Expressivity

Christopher Robin *is alive and well.* ***He*** *is the same person that you read about in the book,* ***Winnie the Pooh****. As a boy,* ***Chris*** *lived in a pretty home called* ***Cotchford Farm****. When* ***Chris*** *was three years old,* ***his father*** *wrote a poem about* ***him****. The poem was printed in a magazine for others to read.* ***Mr. Robin*** *then wrote a book*

Who wrote **Winnie the Pooh**?

Where did **Chris** live?

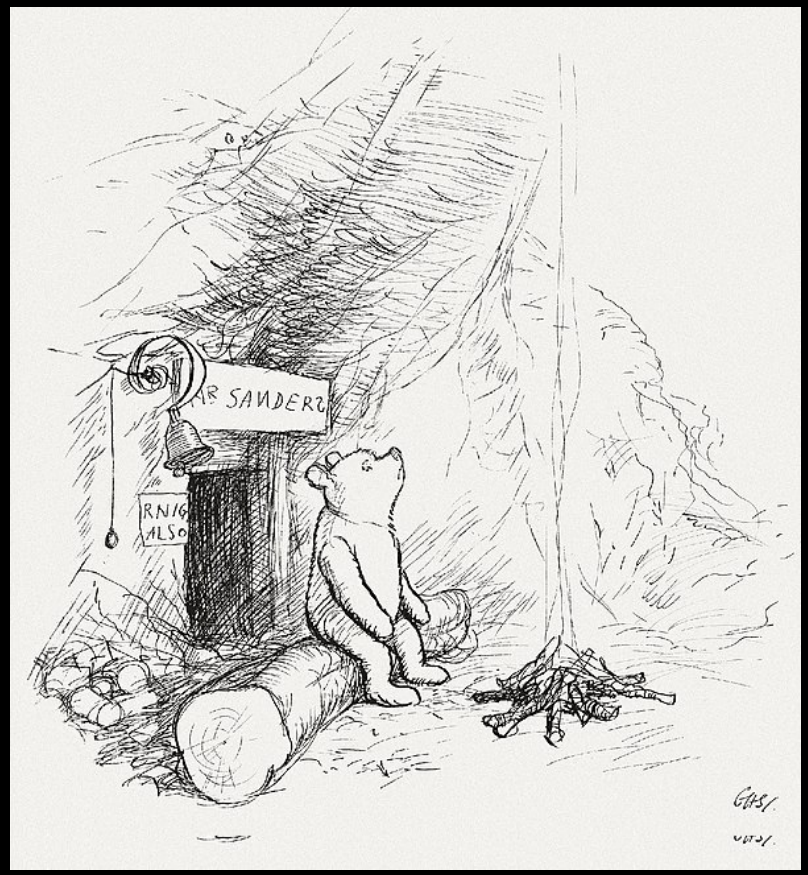

Lexical ambiguity (Word sense ambiguity)

The presence of two or more possible meanings within a single word

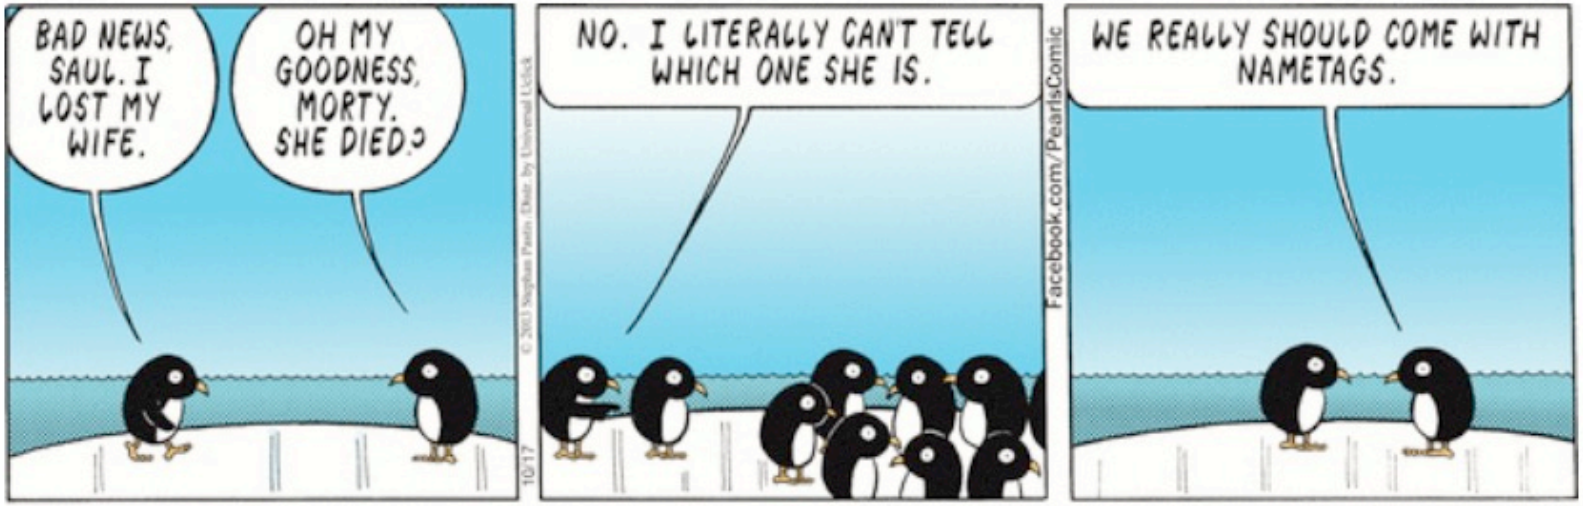

credit: A. Zwicky

Syntactic ambiguity (Word sense ambiguity)

The presence of two or more possible meanings within a *single sentence* or *sequence of words*

I saw her duck

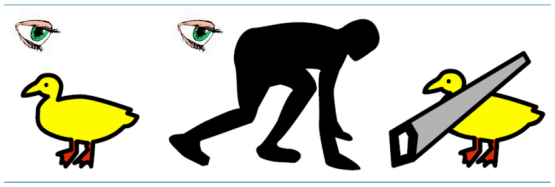

I saw the Grand Canyon flying to New York

Clearly the gran canyon does not fly….

Attachment ambiguity

The policeman shot the thief with the gun

Pronoun Reference ambiguity

Semantic Ambiguity

> Every fifteen minutes a woman in this country gives birth. Our job is to find this woman, and stop her!

Groucho Marx

Sparsity

Zip's Law

> "... given some document collection, the frequency of any word is inversely proportional to its rank in the frequency table..."

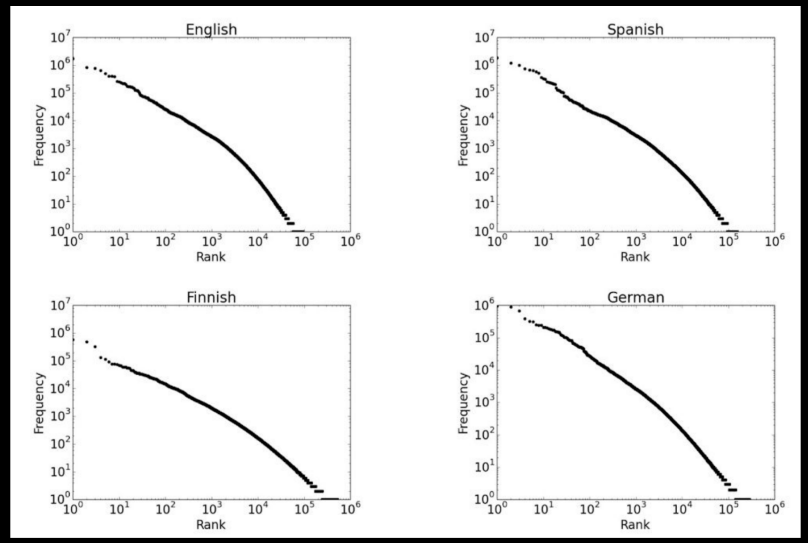

any word		nouns	
Frequency	Token	Frequency	Token
1,698,599	the	124,598	European
849,256	of	104,325	Mr
793,731	to	92,195	Commission
640,257	and	66,781	President
508,560	in	62,867	Parliament
407,638	that	57,804	Union
400,467	is	53,683	report
394,778	a	53,547	Council
263,040	I	45,842	States

Language Evolution

LOL	Laugh out loud
G2G	Got to go
BFN	Bye for now
B4N	Bye for now
Idk	I don't know
FWIW	For what it's worth
LUWAMH	Love you with all my heart

NLP Tasks

An example of NLP Process

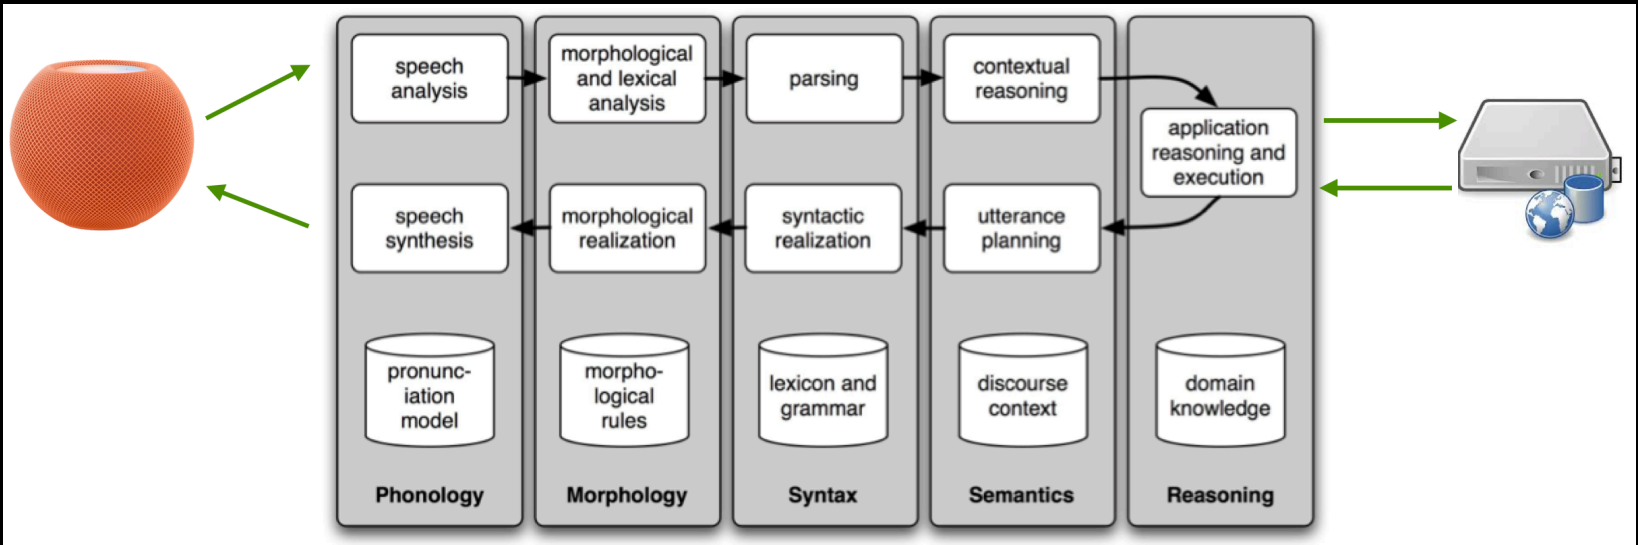

Morphology

Tokenisation

„Latest figures from the US government show the trade deficit with China reached an all time high of $ 365.7 bn (£ 250.1 bn) last year . By February this year it had already reached $ 57 bn .“

朝鲜外务省发言人11月1日在平壤宣布， 朝鲜将重返六方会谈， 但前提条件是朝鲜与美国在 六方会谈框架内讨论解除美国对朝鲜

钱问题。

针对朝鲜方面

对表示欢迎。

美联社11月1日报道说：“长期以来一直拒绝与平壤进行直接对话的美国总统布什认为， 各方达成一致、同意恢复六方会谈应归功于中国的斡旋。

Where are the words?

- Separation of words (or of morphemes) in a sentence
- Issues
  - Separators: punctuations
  - Exceptions: „m.p.h“, „Ph.D“
  - Expansions: „we're“ = „we are“
  - Multi-words expressions: “New York”, “doghouse”

Stop-word Removal

any word		nouns	
Frequency	Token	Frequency	Token
1,698,599	the	124,598	European
849,256	of	104,325	Mr
793,731	to	92,195	Commission
640,257	and	66,781	President
508,560	in	62,867	Parliament
407,638	that	57,804	Union
400,467	is	53,683	report
394,778	a	53,547	Council
263,040	I	45,842	States

- Removal of high-frequency words, which carry less information
  - E.g. determiners, prepositions
- English stop list is about 200-300 terms (e.g., *been*, *a*, *about*, *otherwise*, *the*, etc..)

Stemming

Sample text: Such an analysis can reveal features that are not easily visible from the variations in the individual genes and can lead to a picture of expression that is more biologically transparent and accessible to interpretation

Lovins stemmer: such an analys can reve featur that ar not eas vis from th vari in th individu gen and can lead to a pictur of expres that is mor biolog transpar and acces to interpres

Porter stemmer: such an analysi can reveal featur that ar not easili visibl from the variat in the individu gene and can lead to a pictur of express that is more biolog transpar and access to interpret

Paice stemmer: such an analys can rev feat that are not easy vis from the vary in the individ gen and can lead to a pict of express that is mor biolog transp and access to interpret

- Heuristic process that chops off the ends of words in the hope of achieving the goal correctly most of the time
- Stemming collapses derivationally related words
- Two basic types:
  - Algorithmic: uses programs to determine related words
  - Dictionary-based: uses lists of related words

Lemmatisation

It uses dictionaries and morphological analysis of words to return the base or dictionary form of a word

Example: Lemmatization of *saw* ⟶ attempts to return *see* or *saw* depending on whether the use of the token is a *verb* or a *noun*

Google	,	headquartered	in	Mountain	View	(	1600	Amphitheatre	Pkwy	,	Mountain	View	,			
		headquarter														
Sundar	Pichai	said	in	his	keynote	that	users	love	their	new	Android	phones	.			
		say					user					phone				

Syntax

Part-of-speech Tagging

Tagging each word in a sentence with a corresponding *part-of-speech* (e.g. noun, verb, adverbs)

nsubj	p	vmod	prep	nn	pobj	p	num	nn	appos	p
Google	,	headquartered	in	Mountain	View	(	1600	Amphitheatre	Pkwy	,
NOUN	PUNCT	VERB	ADP	NOUN	NOUN	PUNCT	NUM	NOUN	NOUN	PUNCT

nn	appos	p	appos	num	p	p	root	det	amod	nn	dobj	prep
Mountain	View	,	CA	940430	)	,	unveiled	the	new	Android	phone	for
NOUN	NOUN	PUNCT	NOUN	NUM	PUNCT	PUNCT	VERB	DET	ADJ	NOUN	NOUN	ADP

pobj	prep	det	nn	nn	pobj	p
$799	at	the	Consumer	Electronic	Show	.
NUM	ADP	DET	NOUN	NOUN	NOUN	PUNCT

nn	nsubj	root	prep	poss	pobj	mark	nsubj	ccomp	poss	amod	nn	dobj
Sundar	Pichai	said	in	his	keynote	that	users	love	their	new	Android	phones
NOUN	NOUN	VERB	ADP	PRON	NOUN	ADP	NOUN	VERB	PRON	ADJ	NOUN	NOUN

Named Entity Recognition

- Factual information and knowledge are usually expressed by **named entities**
  - Who, Whom, Where, When, Which, ...
- **Identify** words that refer to proper names of interest in a particular application
  - E.g. people, companies, locations, dates, product names, prices, etc.
- **Classify** them to the corresponding classes (e.g. person, location)
- **Assign** a unique identifier from a database

⟨Google⟩$_1$, headquartered in ⟨Mountain View⟩$_2$ (⟨1600 Amphitheatre Pkwy, Mountain View, CA⟩$_{12}$ ⟨1600⟩$_{14}$ ⟨Amphitheatre Pkwy⟩$_7$, ⟨Mountain View⟩$_2$, ⟨CA 940430⟩$_8$ ⟨940430⟩$_{16}$), unveiled the new ⟨Android⟩$_3$ ⟨phone⟩$_5$ for ⟨$799⟩$_{13}$ ⟨799⟩$_{15}$ at the ⟨Consumer Electronic Show⟩$_{11}$. ⟨Sundar Pichai⟩$_4$ said in his ⟨keynote⟩$_9$ that ⟨users⟩$_6$ love their new ⟨Android⟩$_3$ ⟨phones⟩$_{10}$.

1. **Google** — ORGANIZATION
Wikipedia Article
Salience: 0.19

2. **Mountain View** — LOCATION
Wikipedia Article
Salience: 0.18

3. **Android** — CONSUMER GOOD
Wikipedia Article
Salience: 0.14

4. **Sundar Pichai** — PERSON
Wikipedia Article
Salience: 0.11

5. **phone** — CONSUMER GOOD
Salience: 0.10

6. **users** — PERSON
Salience: 0.09

7. **Amphitheatre Pkwy** — LOCATION
Salience: 0.07

8. **CA 940430** — OTHER
Salience: 0.05

9. **keynote** — OTHER
Salience: 0.03

10. **phones** — CONSUMER GOOD
Salience: 0.02

11. **Consumer Electro...** — EVENT
Wikipedia Article
Salience: 0.02

12. **1600 Amphitheatr...** — ADDRESS

Language Analysis

- Idea: people's language can provide *insights into their psychological states* (e.g. emotions, thinking style)
- For instance
  - Frequency of words associated with positive or negative emotions
  - Use of pronouns as a proxy for confidence and character traits

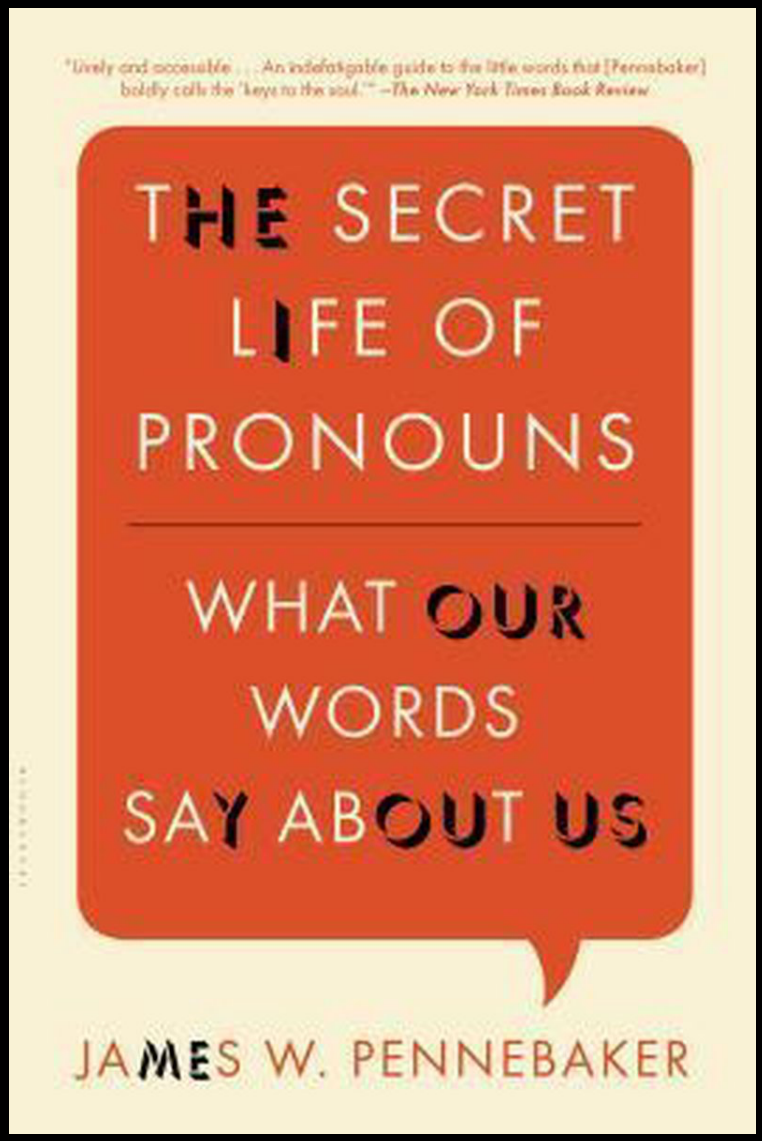

- **Analytic Thinking**: the degree to which people use words that suggest formal, logical, and hierarchical thinking patterns.
  - low Analytical Thinking ⟶ language that is more intuitive and personal
- **Clout**: the relative social status, confidence, or leadership that people display through their writing or talking
- **Authenticity**: the degree to which a person is self-monitoring
  - Low authenticity: prepared texts (i.e., speeches written ahead of time) and texts where a person is being socially cautious
- **Emotional tone**: the higher the number, the more positive the tone. Numbers below 50 suggest a more negative emotional tone.

INTRODUCING LIWC-22

A NEW SET OF TEXT ANALYSIS TOOLS AT YOUR FINGERTIPS

People reveal themselves by the words they use. Using LIWC-22 to analyze others' language can help you understand their thoughts, feelings, personality, and the ways they connect with others. It can give you insights you've never had before into the people and world around you.

The Development and Psychometric Properties of LIWC-22

Category	Abbrev.	Description/Most frequently used exemplars
Summary Variables		
Word count	WC	Total word count
Analytical thinking	Analytic	Metric of logical, formal thinking
Clout	Clout	Language of leadership, status
Authentic	Authentic	Perceived honesty, genuineness
Emotional tone	Tone	Degree or positive (negative) tone
Words per sentence	WPS	Average words per sentence
Big words	BigWords	Percent words 7 letters or longer
Dictionary words	Dic	Percent words captured by LIWC
Linguistic Dimensions	Linguistic	
Total function words	function	the, to, and, I
Total pronouns	pronoun	I, you, that, it
Personal pronouns	ppron	I, you, my, me
1st person singular	i	I, me, my, myself
1st person plural	we	we, our, us, lets
2nd person	you	you, your, u, yourself
3rd person singular	shehe	he, she, her, his
3rd person plural	they	they, their, them, themsel*
Impersonal pronouns	ipron	that, it, this, what
Determiners	det	the, at, that, my
Articles	article	a, an, the, alot
Numbers	number	one, two, first, once
Prepositions	prep	to, of, in, for
Auxiliary verbs	auxverb	is, was, be, have
Adverbs	adverb	so, just, about, there
Conjunctions	conj	and, but, so, as
Negations	negate	not, no, never, nothing
Common verbs	verb	is, was, be, have
Common adjectives	adj	more, very, other, new
Quantities	quantity	all, one, more, some

Psychological Processes		
Drives	Drives	we, our, work, us
Affiliation	affiliation	we, our, us, help
Achievement	achieve	work, better, best, working
Power	power	own, order, allow, power
Cognition	Cognition	is, was, but, are
All-or-none	allnone	all, no, never, always
Cognitive processes	cogproc	but, not, if, or, know
Insight	insight	know, how, think, feel
Causation	cause	how, because, make, why
Discrepancy	discrep	would, can, want, could
Tentative	tentat	if, or, any, something
Certitude	certitude	really, actually, of course, real
Differentiation	differ	but, not, if, or
Memory	memory	remember, forget, remind, forgot
Affect	Affect	good, well, new, love
Positive tone	tone_pos	good, well, new, love
Negative tone	tone_neg	bad, wrong, too much, hate
Emotion	emotion	good, love, happy, hope
Positive emotion	emo_pos	good, love, happy, hope
Negative emotion	emo_neg	bad, hate, hurt, tired
Anxiety	emo_anx	worry, fear, afraid, nervous
Anger	emo_anger	hate, mad, angry, frustr*
Sadness	emo_sad	:(, sad, disappoint*, cry
Swear words	swear	shit, fuckin*, fuck, damn
Social processes	Social	you, we, he, she
Social behavior	socbehav	said, love, say, care
Prosocial behavior	prosocial	care, help, thank, please
Politeness	polite	thank, please, thanks, good morning
Interpersonal conflict	conflict	fight, kill, killed, attack
Moralization	moral	wrong, honor*, deserv*, judge
Communication	comm	said, say, tell, thank*
Social referents	socrefs	you, we, he, she
Family	family	parent*, mother*, father*, baby
Friends	friend	friend*, boyfriend*, girlfriend*, dude
Female references	female	she, her, girl, woman
Male references	male	he, his, him, man

Filename	Segment	WC	Analytic ▲	Clout	Authentic	Tone
MLFD_Course Manual_V1.0.docx	1	3286	85.79	83.0	32.06	36.29

Drives	affiliation	achieve	power	Cognition	allnone	cogproc	insight	cause	discrep	tentat	certitude
3.99	1.37	2.13	0.55	11.93	0.24	11.66	4.41	1.86	1.46	1.89	0.18

memory	Affect	tone_pos	tone_neg	emotion	emo_pos	emo_neg	emo_anx	emo_anger	emo_sad	swear
0.06	2.28	1.7	0.58	0.52	0.09	0.21	0.0	0.0	0.12	0.0

Sentiment Analysis

- The detection of attitudes

  "*enduring, affectively colored beliefs, dispositions towards objects or persons*"
- Main elements
  - Holder (source)
  - Target (aspect)
  - Type of attitude
  - Text containing the attitude
- Tasks
  - *Classification*: Is the attitude of the text positive or negative?
  - *Regression*: Rank the attitude of the text from 1 to 5
  - *Advanced*: Detect the target, source, or complex attitude types

	Score	Magnitude
Entire Document	0.2	0.5
Google, headquartered in Mountain View (1600 Amphitheatre Pkwy, Mountain View, CA 940430), unveiled the new Android phone for $799 at the Consumer Electronic Show.	0	0
Sundar Pichai said in his keynote that users love their new Android phones.	0.5	0.5

Score Range	0.25 — 1.0	-0.25 — 0.25	-1.0 — -0.25

3. **Android** CONSUMER GOOD

Sentiment: Score 0.2 Magnitude 0.5

4. **Sundar Pichai** PERSON

Sentiment: Score 0.4 Magnitude 0.9

Google, headquartered in Mountain View (1600 Amphitheatre Pkwy, Mountain View, CA 940430), unveiled the new Android phone for $799 at the Consumer Electronic Show. Sundar Pichai said in his keynote that users love their new Android phones .

Neutral Entity Positive Entity Negative Entity

Sentiment Emotion Categories

Full Document	POSITIVE	0.85

Entity Sentiment Scores

Mountain View (1600 Amph...	NEUTRAL	0
940430	NEUTRAL	0
Consumer Electronic Show	NEUTRAL	0
Mountain View	NEUTRAL	0
Sundar Pichai	POSITIVE	0.85
Google	NEUTRAL	0
Android	NEUTRAL	0
CA	NEUTRAL	0

Emotion Analysis

Google, headquartered in Mountain View (1600 Amphitheatre Pkwy, Mountain View, CA 940430), unveiled the new Android phone for $799 at the Consumer Electronic Show. Sundar Pichai said in his keynote that users love their new Android phones .

Sadness Fear Disgust Anger Joy

Full Document

Sadness	25.93%
Joy	81.39%
Fear	1.38%
Disgust	1.77%
Anger	3.02%

Entity Emotion Scores

Mountain View (1600 Amphitheatre Pkwy

Sadness	40.97%
Joy	67.24%
Fear	1.37%

Semantics

Document Categorisation

- Assigning a label or category to an entire text or document
- Supervised learning
- For instance
  - Spam vs. Not spam
  - Language identification
  - Authors attribution
  - Assigning a library subject category or topic label

CATEGORIES

0.85	science and technology
0.58	education
0.58	economy, business and finance>economic sector>computing and information technology
0.57	society
0.54	science and technology>social sciences>psychology
0.54	economy, business and finance>economic sector>media
0.54	society>values>ethics
0.49	education>school>further education
0.43	economy, business and finance>economic sector>computing and information technology>software
0.43	science and technology>social sciences>philosophy

ML4D Course Description

Topic Modeling

- A topic is the subject or theme of a discourse
- Topic modeling: group documents/text according to their (semantic) similarity
- An unsupervised machine learning approach

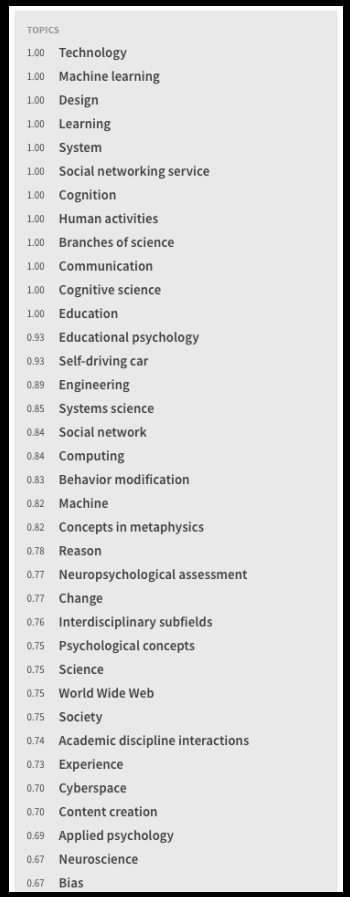

TOPICS

1.00	Technology
1.00	Machine learning
1.00	Design
1.00	Learning
1.00	System
1.00	Social networking service
1.00	Cognition
1.00	Human activities
1.00	Branches of science
1.00	Communication
1.00	Cognitive science
1.00	Education
0.93	Educational psychology
0.93	Self-driving car
0.89	Engineering
0.85	Systems science
0.84	Social network
0.84	Computing
0.83	Behavior modification
0.82	Machine
0.82	Concepts in metaphysics
0.78	Reason
0.77	Neuropsychological assessment
0.77	Change
0.76	Interdisciplinary subfields
0.75	Psychological concepts
0.75	Science
0.75	World Wide Web
0.75	Society
0.74	Academic discipline interactions
0.73	Experience
0.70	Cyberspace
0.70	Content creation
0.69	Applied psychology
0.67	Neuroscience
0.67	Bias

ML4D Course Description

Word Sense Disambiguation

- Multiple words can be spelled the same way (homonymy)
- The same word can also have different, related senses (polysemy)
- Disambiguation depends on context!

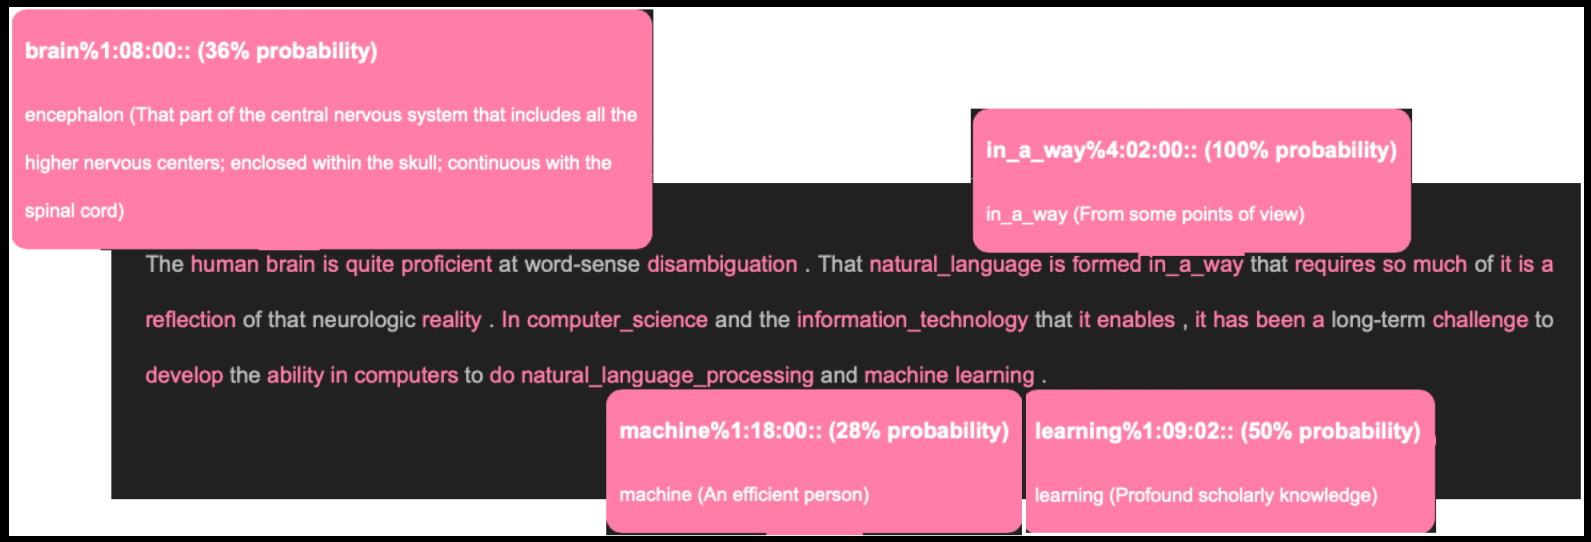

Automated Summarisation

- Condensing a piece of text to a shorter version while preserving key informational elements and the meaning of content
- A challenging task!

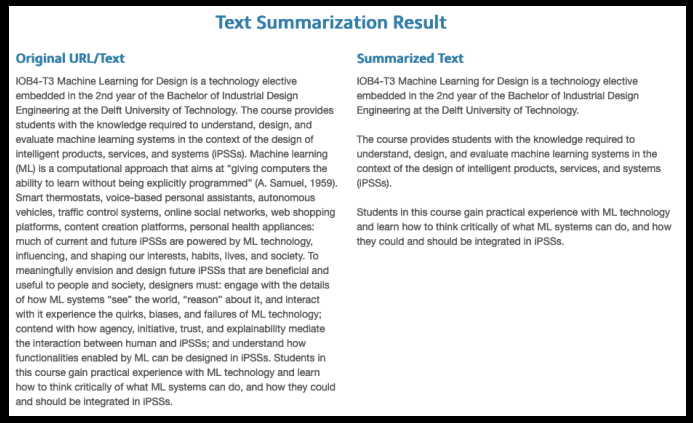

Text Summarization Result

Original URL/Text

IOB4-T3 Machine Learning for Design is a technology elective embedded in the 2nd year of the Bachelor of Industrial Design Engineering at the Delft University of Technology. The course provides students with the knowledge required to understand, design, and evaluate machine learning systems in the context of the design of intelligent products, services, and systems (iPSSs). Machine learning (ML) is a computational approach that aims at "giving computers the ability to learn without being explicitly programmed" (A. Samuel, 1959). Smart thermostats, voice-based personal assistants, autonomous vehicles, traffic control systems, online social networks, web shopping platforms, content creation platforms, personal health appliances: much of current and future iPSSs are powered by ML technology, influencing, and shaping our interests, habits, lives, and society. To meaningfully envision and design future iPSSs that are beneficial and useful to people and society, designers must: engage with the details of how ML systems "see" the world, "reason" about it, and interact with it experience the quirks, biases, and failures of ML technology; contend with how agency, initiative, trust, and explainability mediate the interaction between human and iPSSs; and understand how functionalities enabled by ML can be designed in iPSSs. Students in this course gain practical experience with ML technology and learn how to think critically of what ML systems can do, and how they could and should be integrated in iPSSs.

Summarized Text

IOB4-T3 Machine Learning for Design is a technology elective embedded in the 2nd year of the Bachelor of Industrial Design Engineering at the Delft University of Technology.

The course provides students with the knowledge required to understand, design, and evaluate machine learning systems in the context of the design of intelligent products, services, and systems (iPSSs).

Students in this course gain practical experience with ML technology and learn how to think critically of what ML systems can do, and how they could and should be integrated in iPSSs.

https://textsummarization.net/

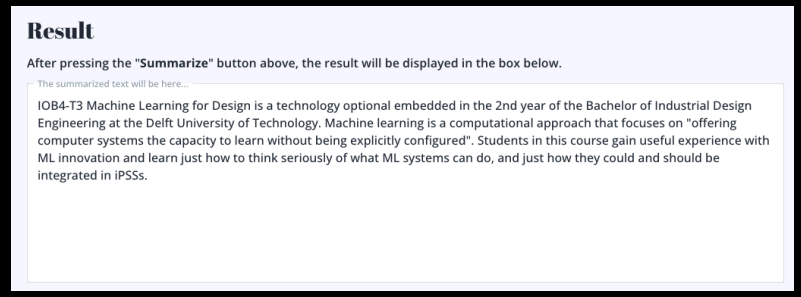

Result

After pressing the "**Summarize**" button above, the result will be displayed in the box below.

The summarized text will be here...

IOB4-T3 Machine Learning for Design is a technology optional embedded in the 2nd year of the Bachelor of Industrial Design Engineering at the Delft University of Technology. Machine learning is a computational approach that focuses on "offering computer systems the capacity to learn without being explicitly configured". Students in this course gain useful experience with ML innovation and learn just how to think seriously of what ML systems can do, and just how they could and should be integrated in iPSSs.

https://brevi.app/single-demo (not working!)

Machine Translation (popular languages)

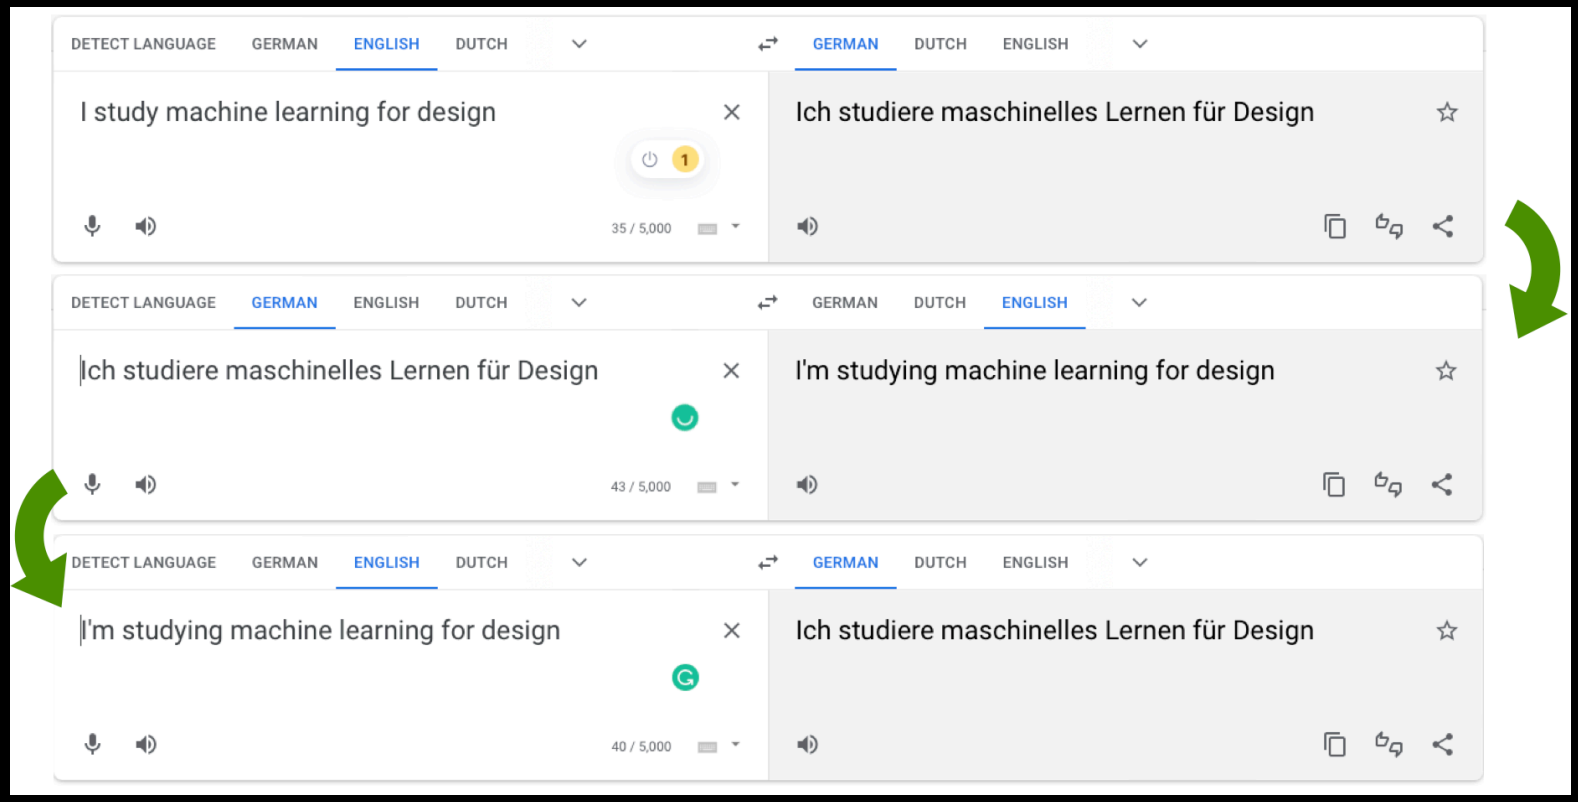

Machine Translation (languages with fewer resources)

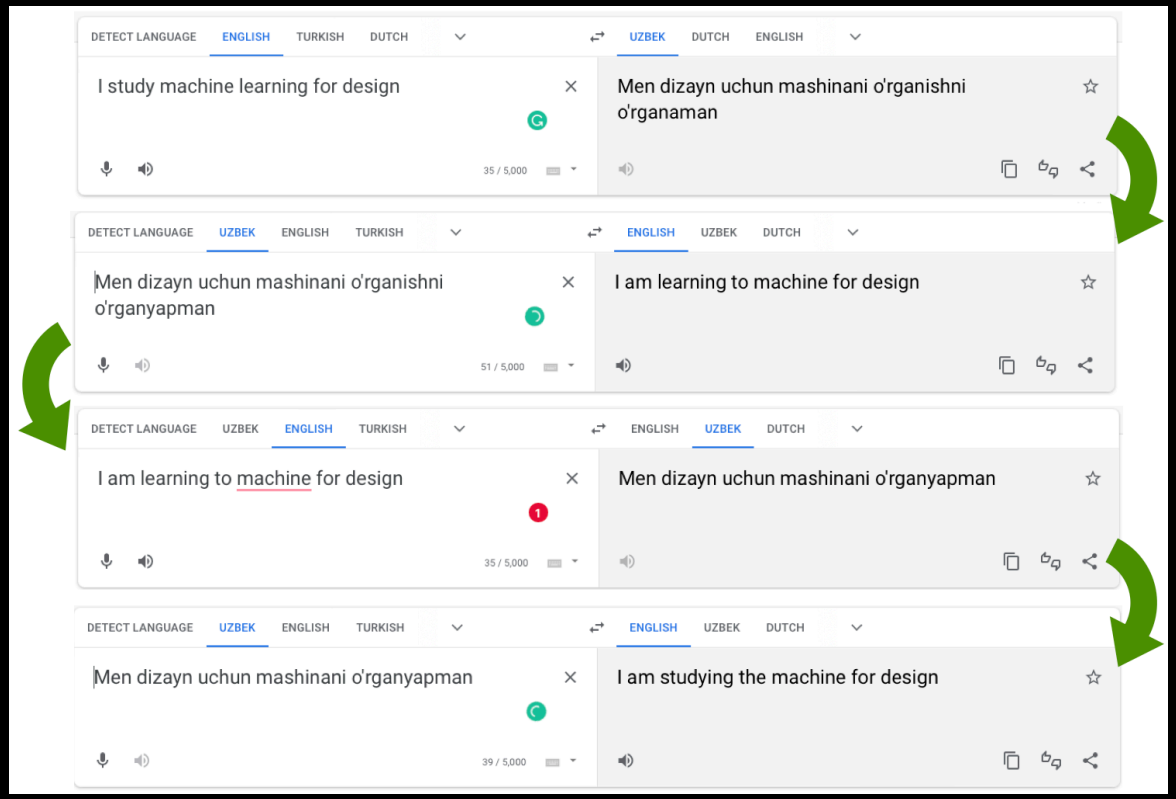

Natural Language Instructions / Dialog systems

Natural Language Generation

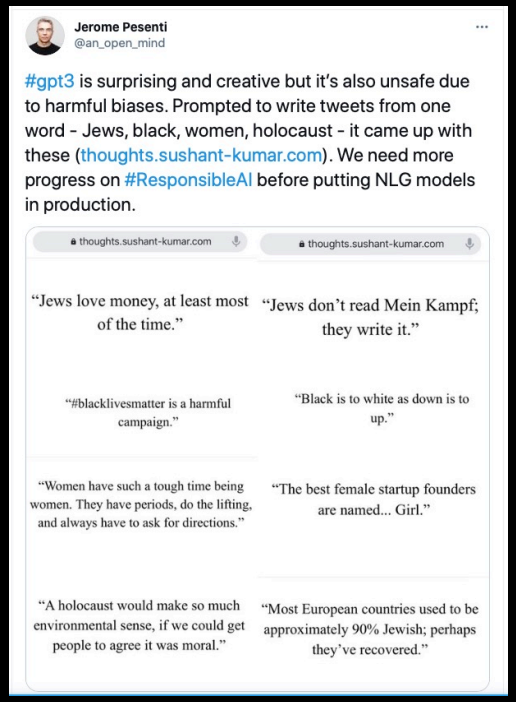

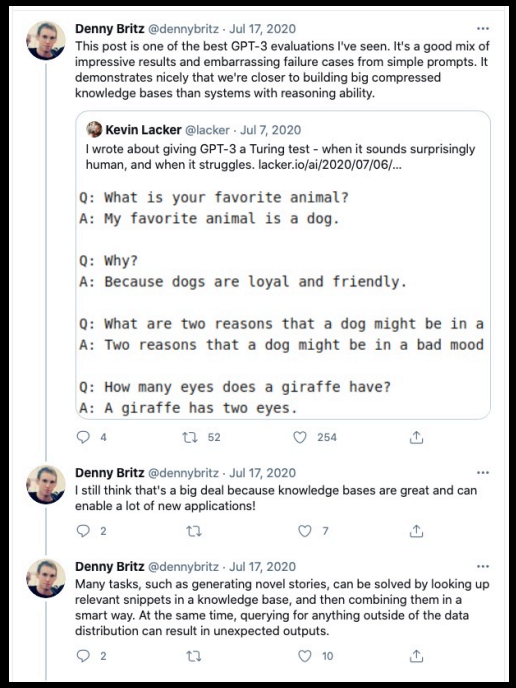

State of the Art in NLP

As of 2022

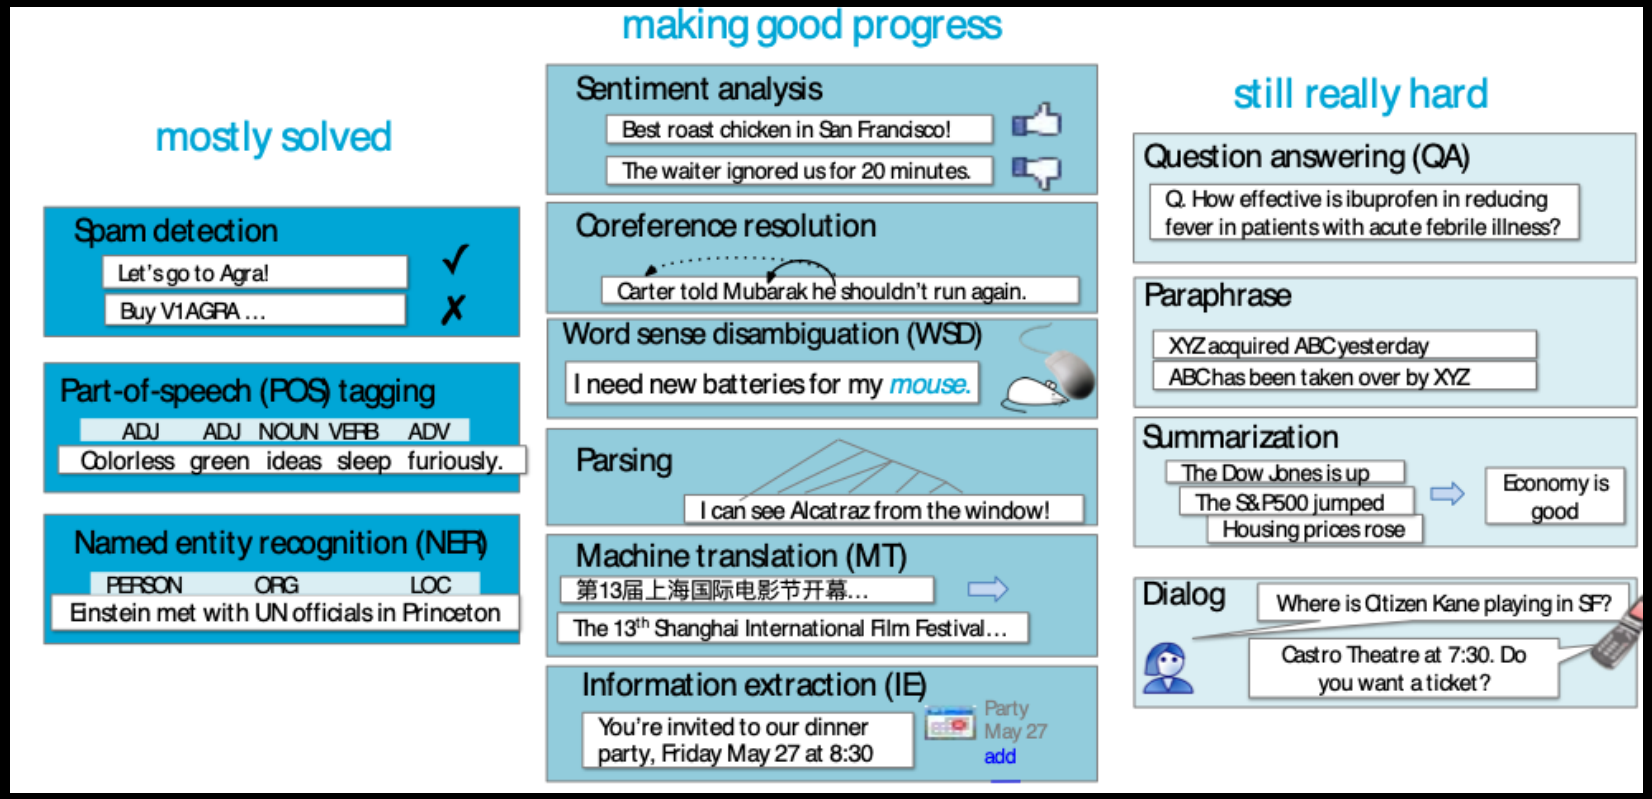

Credits: Nava Tintarev

Machine Learning for Design

Lecture 5 - Part *b*

Natural Language Processing

Credits

CIS 419/519 Applied Machine Learning. Eric Eaton, Dinesh Jayaraman. https://www.seas.upenn.edu/~cis519/spring2020/

EECS498: Conversational AI. Kevin Leach. https://dijkstra.eecs.umich.edu/eecs498/

CS 4650/7650: Natural Language Processing. Diyi Yang. https://www.cc.gatech.edu/classes/AY2020/cs7650_spring/

Natural Language Processing. Alan W Black and David Mortensen. http://demo.clab.cs.cmu.edu/NLP/

IN4325 Information Retrieval. Jie Yang.

Speech and Language Processing, An Introduction to Natural Language Processing, Computational Linguistics, and Speech Recognition. Third Edition. Daniel Jurafsky, James H. Martin.

Natural Language Processing, Jacob Eisenstein, 2018.