Statistical Natural Language Processing

Çağrı Çöltekin
/tʃaːrˈɯ tʃœltecˈɪn/
ccoltekin@sfs.uni-tuebingen.de

University of Tübingen
Seminar für Sprachwissenschaft

Summer Semester 2019

Why study (statistical) NLP

- (Most of) you are studying in a 'computational linguistics' program
- Many practical applications
- Investigating basic questions in linguistics and cognitive science (and more)

Application examples

Just a few examples

For profit (engineering):

- Machine translation
- Question answering
- Information retrieval
- Dialog systems
- Summarization
- Text classification
- Text mining/analytics
- Sentiment analysis
- Speech recognition and synthesis
- Automatic grading
- Forensic linguistics

For fun (research):

- Modeling language processing learning
- Investigating language change through time and space
- (Aiding) language documentation through text processing
- (Automatic) corpus annotation for linguistic research
- Stylometry, author identification

Layers of linguistic analysis

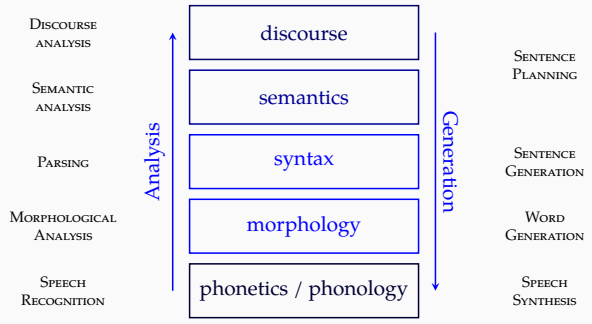

Annotation layers: example

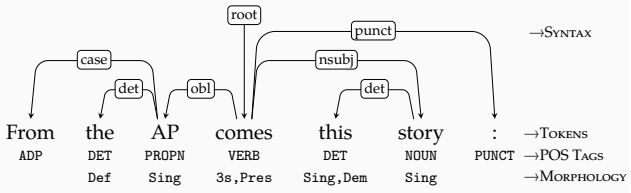

Typical NLP pipeline

- Text processing / normalization
- Word/sentence tokenization
- POS tagging
- Morphological analysis
- Syntactic parsing
- Semantic parsing
- Named entity recognition
- Coreference resolution

Do we need a pipeline?

- Most "traditional" NLP architectures are based on a pipeline approach:
  - tasks are done individually, results are passed to upper level
- Joint learning (e.g., POS tagging and syntax) often improves the results
- End-to-end learning (without intermediate layers) is another (recent/trending) approach

On the word 'statistical'

> *But it must be recognized that the notion 'probability of a sentence' is an entirely useless one, under any known interpretation of this term.* — *Chomsky (1968)*

- Some linguistic traditions emphasize(d) use of 'symbolic', rule-based methods
- Some NLP systems are based on rule-based systems (esp. from 80's 90's)
- Virtually, all modern NLP systems include some sort of statistical component

What is difficult with NLP?

- Combinatorial problems - computational complexity
- Ambiguity
- Data sparseness

NLP and computational complexity

- How many possible parses a sentence may have?
- How many ways can you align two (parallel) sentences?
- How to calculate probability of sentence based on the probabilities of words in it?
- Many similar questions we deal with have an exponential search space
- Naive approaches often are computationally intractable

Combinatorial problems

A typical linguistic problem: parsing

How many different binary trees can span a sentence of N words?

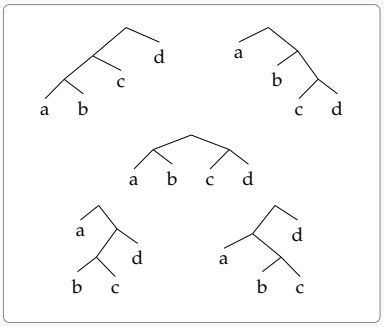

words	trees
2	1
3	2
4	5
5	14
10	4862
20	1 767 263 190
...	...

NLP and ambiguity

fun with newspaper headlines

FARMER BILL DIES IN HOUSE
TEACHER STRIKES IDLE KIDS
SQUAD HELPS DOG BITE VICTIM
BAN ON NUDE DANCING ON GOVERNOR'S DESK
PROSTITUTES APPEAL TO POPE
KIDS MAKE NUTRITIOUS SNACKS
DRUNK GETS NINE MONTHS IN VIOLIN CASE
MINERS REFUSE TO WORK AFTER DEATH

More ambiguities

we do not recognize many of them at first read

- Time flies like an arrow;
  fruit flies like a banana.
- Outside of a dog, a book is a man's best friend;
  inside it's too hard to read.
- One morning I shot an elephant in my pajamas.
  How he got in my pajamas, I don't know.
- Don't eat the pizza with knife and fork;
  the one with anchovies is better.

Even more ambiguities

with pretty pictures

Cartoon Theories of Linguistics, SpecGram Vol CLIII, No 4, 2008. http://specgram.com/CLIII.4/school.gif

Statistical methods and data sparsity

- Statistical methods (machine learning) are the best way we know to deal with ambiguities
- Even for rule-based approaches, a statistical disambiguation component is often needed
- Machine learning methods require (annotated) data
- But ...

Languages are full of rare events

word frequencies in a small corpus

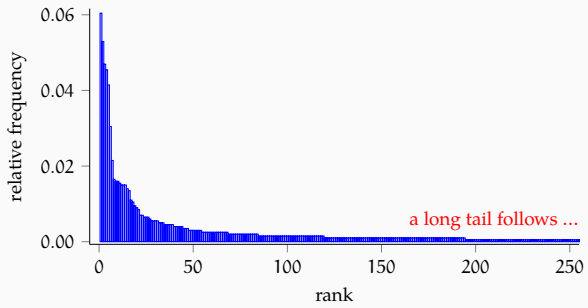

What is difficult in CL?

and how can machine learning help?

- Combinatorial problems - computational complexity
  - Often we resort to approximate methods: the answer to 'what is a good approximation?' comes from ML.
- Ambiguity
  - The answer to 'what is the best choice?' comes from ML.
- Data sparseness
  - Even here, ML can help.

What is in this course

- Quick introduction / refreshers on important prerequisites
- The computational linguist's toolbox: basic methods and tools in NLP
- Some applications of NLP

What is in this course

Preliminaries

- Linear algebra, some concepts from calculus
- Probability theory
- Information theory
- Statistical inference
- Some topics from machine learning
  - Regression & classification
  - Sequence learning (HMMs)
  - Neural networks and deep learning
  - Unsupervised learning

What is in this course

NLP Tools and techniques

- Tokenization, normalization, segmentation
- N-gram language models
- Part of speech tagging
- Statistical parsing
- Distributed representations (of words, and other linguistic objects)

What is in this course

Applications

- Text classification
  - sentiment analysis
  - language detection
  - authorship attribution
  - ...

If time allows

- Statistical machine translation
- Named entitiy recognition
- Text summarization
- Dialog systems
- ...

What is not in this course

- Cutting edge, latest methods & applications
- In-depth treatment of particular topics
- Introduction to terms / concepts from linguistics

Logistics

- Lectures: Mon/Fri 12:15 at Hörsaal 0.02
- Practical sessions: Wed 10:15 at Hörsaal 0.02
- Office hours: Mon 14:00-15:00 (room 1.09), or by appointment (email ccoltekin@sfs.uni-tuebingen.de)
- Course web page: `http://sfs.uni-tuebingen.de/~ccoltekin/courses/snlp`
- We will use GitHub classroom in this class (more on this soon)

Reading material

- Daniel Jurafsky and James H. Martin (2009). *Speech and Language Processing: An Introduction to Natural Language Processing, Computational Linguistics, and Speech Recognition*. second. Pearson Prentice Hall. ISBN: 978-0-13-504196-3
  - Draft chapters of the third edition is available at `http://web.stanford.edu/~jurafsky/slp3/`
- Trevor Hastie, Robert Tibshirani, and Jerome Friedman (2009). *The Elements of Statistical Learning: Data Mining, Inference, and Prediction*. Second. Springer series in statistics. Springer-Verlag New York. ISBN: 9780387848587. URL: `http://web.stanford.edu/~hastie/ElemStatLearn/`
- Course notes for some lectures
- Other online references

Grading / evaluation

- 7 graded assignments (6-best counts, 10 % each)
- Final exam (40 %)
- Attendance
  - 5 % (bonus) if you miss only one or two classes
  - you lose one bonus point for each additional class you miss
- Up to 5 % additional bonus points for Easter eggs:
  - first person finding (intentional, trivial) mistakes in the course material gets 1 %

Assignments

- For distribution and submission of assignments, we will use GitHub Classroom
- The amount of `git` usage required is low, but learning/using `git` well is strongly recommended
- You are encouraged work on the assignments in pairs, but **you can work with the same person only once**
- Late assignments up to one week, will be graded up to half points indicated
- The solutions will be discussed in the tutorial session after one week from deadline
- Poll: a match-making system for working in random groups?

Assignment 0

- Your first assignment is already posted on the web page
- By completing assignment 0, you will
  - register for the course
  - have access to the non-public course material
  - exercise with the way later assignments will work
  - provide some data for future exercises
- The repository created for assignment 0 is private, and can only be accessed you and the instructors

Practical sessions

- Tutors: Marko Lozajic & Maxim Korniyenko
- You need to bring your own computer, make sure you have a working Python interpreter
- You are encouraged to ask questions about the exercises during practical sessions
- The solutions will be discussed during tutorial sessions
- Poll: Python tutorial?

Further git/GitHub usage

- Once you complete Assignment 0, you will be a member of the 'organization' `snlp2019`
- You will get access to
  - private course material
  - assignment links
  - news and announcements

  through the repository at
  `https://github.com/snlp2018/snlp2019`
- Make sure to watch this repository
- You are also encouraged to use 'issues' in this repository as a place to discuss course topics, ask questions about the material and assignments

Next

Mon Mathematical preliminaries (some linear algebra and bits from calculus)

Fri Probability theory

References / additional reading material

Bishop, Christopher M. (2006). *Pattern Recognition and Machine Learning*. Springer. ISBN: 978-0387-31073-2.

Chomsky, Noam (1968). "Quine's empirical assumptions". In: *Synthese* 19.1, pp. 53–68. DOI: 10.1007/BF00568049.

Hastie, Trevor, Robert Tibshirani, and Jerome Friedman (2009). *The Elements of Statistical Learning: Data Mining, Inference, and Prediction*. Second. Springer series in statistics. Springer-Verlag New York. ISBN: 9780387848587. URL: http://web.stanford.edu/~hastie/ElemStatLearn/.

Jurafsky, Daniel and James H. Martin (2009). *Speech and Language Processing: An Introduction to Natural Language Processing, Computational Linguistics, and Speech Recognition*. second. Pearson Prentice Hall. ISBN: 978-0-13-504196-3.

Manning, Christopher D. and Hinrich Schütze (1999). *Foundations of Statistical Natural Language Processing*. MIT Press. ISBN: 9780262133609.