

Lexicon Acquisition

Daniel Zeman

📅 October 16, 2020





EUROPEAN UNION
European Structural and Investment Fund
Operational Programme Research,
Development and Education

Charles University
Faculty of Mathematics and Physics
Institute of Formal and Applied Linguistics



unless otherwise stated

- Some hints only (approach must vary depending on language)
- Identify part of speech and **inflection pattern**
- If affixes restrict possible classes, use them!
 -  Czech: the following suffixes increase likelihood of an infinitive: *-st, -át, -at, -ct, -ci, -ít, -out, -ýt, -ovat, -it, -ět, -et*
 -  English: little inflection but verb forms and derivational suffixes (*-ness, -ity, -able*) can help
- Otherwise, syntax might help
 - E.g. if it's after preposition or article it's likely an adjective or a noun

Lexicon Acquisition

- Create word frequency list
- Identify closed-class words
 - Many of them will be very frequent
 - Textbook and/or bilingual dictionary may help with the rest
 - Parallel corpus + word aligner may supplement the dictionary
- What remains are mostly nouns, adjectives, verbs, and adverbs
 - Try to sort it out by iteratively looking at the word list, identifying repeating affixes etc.
 - If there are no repeating bound morphemes
 - then you may not be able to sort out the parts of speech
 - but maybe the morphology of the language is not so interesting after all



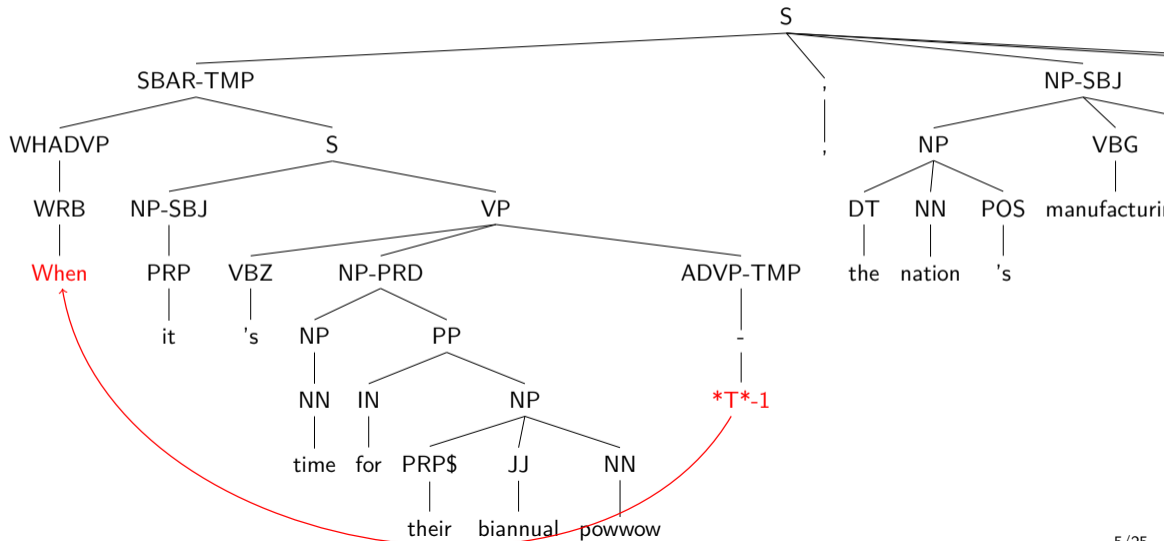
- `http://ufal.mff.cuni.cz/~zeman/vyuka/morfosynt/lab-lexicon/index.html`
- There is also a link from the NPFL094 website



- Example only! Other corpora and languages may require a different approach
- Input: plain-text (taken from Penn Treebank)
 - **Tokenized** (punctuation separated from words)
 - Remove **traces** (non-word terminal nodes in Penn Treebank): all tokens containing “*”?
 - **Lowercase**
 - Later we will want to identify proper nouns
 - Complicated by sentence-initial capitalization



Traces





English Frequency Wordlist

- Penn Treebank 3 / Wall Street Journal:
 - 49,208 sentences
 - 1,253,013 terminal nodes (tokens and traces)
 - 49,817 **word types** (opposed to **word occurrences**) including traces
 - 43,764 lowercased types without traces and some other technical nodes (“error:” etc.)
 - The most frequent types often have these (overlapping) properties:
 - **stopwords**
 - closed-class words
 - short words?



English Frequency Wordlist

,	60484	\$	8817
the	59459	“	8735
.	48144	is	8539
to	29576	”	8506
of	28440	it	7195
a	24781	said	7141
in	21047	on	6646
and	20449	%	6121
's	11556	at	5770
for	10454	by	5705
that	10422	as	5701

Punctuation and Special Characters

m/\pP/

,	60484		
.	48144		
's	11556	third-quarter	333
\$	8817	buy-out	222
"	8735	s&p	164
"	8506	3,000	28
%	6121	3.7	28
mr.	4950		
n't	4006	<i>total types</i>	10607
–	2585	<i>the rest</i>	33157
u.s.	2056		

- Caught, OK
- Not caught (but should have been caught)
- Caught (disputable)
- Caught (tokenization-related)

Numbers

m/\pN/

0	12447
10	668
30	607...
1988	503...
1,000	111...
1/2	105...
1.5	88...
30-year	79...
1980s	53...
ru-486	15...
mid-1980s	12...

b-2	7
19th	7
1989-90	5
80%-owned	4
xr4ti	4
<i>total types</i>	6123
<i>the rest</i>	37641
<i>no punctuation or numbers</i>	32218





Real Words

!m/ [\pP\pN`\$] /

the	59459	on	6646
of	28440	at	5770
to	27448	by	5705
a	24781	as	5701
in	21047	from	5438
and	20449	with	5357
for	10454	million	5335
that	10422	was	4901
is	8539	be	4586
it	7195	its	4571
said	7141	are	4528

- Pronouns / determiners / articles in all cases
 - Personal: *I, me, you, he, him, she, her, it, we, us, they, them*
 - Impersonal: *one* (as in *One has to be careful here.*)
 - Reflexive: *myself, yourself, himself, herself, itself, ourselves, yourselves, themselves, oneself*
 - Possessive: *my, mine, your, yours, his, her, hers, its, our, ours, their, theirs*
 - Demonstrative: *this, these, that, those*
 - Article: *the, a, an*
 - Interrogative / relative: *who, whom, whose, what, which, whoever, whomever, whatever*
 - Indefinite: *some, somebody, someone, something, any, anybody, anyone, anything; many, much, more, most, too, enough, few, little, fewer, less, fewest, least*
 - Total: *every, everybody, everyone, everything, each, all, both*
 - Negative: *no, nobody, nothing, none*

- Numerals
 - Cardinal
 - *zero, one, two, three, four, five, six, seven, eight, nine, ten*
 - *eleven, twelve, thirteen, fourteen, fifteen, sixteen, seventeen, eighteen, nineteen*
 - *twenty, thirty, forty, fifty, sixty, seventy, eighty, ninety*
 - *hundred, thousand, million, billion*
 - Ordinal
 - *first, second, third, fourth, fifth, sixth, seventh, eighth, ninth, tenth* ⇒ morphology *-th*
 - In some languages written as one word, i.e., a nice morphological exercise:
 - 361,972
 -  en: **three** hundred **sixty-one** thousand **nine** hundred and **seventy-two**
 -  de: **drei**hundert**einundsechzig**tausend**neun**hundert**zwei**und**siebzig**

- Auxiliary and modal verbs
 - *be, am, are, is, was, were, been, being, 'm, 's, 're, aren't, isn't, wasn't, weren't, ain't*
 - *have, has, had, having, 've, 's, 'd, haven't, hasn't*
 - *will, would, (willing), 'll, 'd, won't, wouldn't*
 - *can, cannot, could, couldn't*
 - *shall, should, shouldn't*
 - *may, might*
 - *must*
 - *do, does, did, done, doing, don't, doesn't*

- Pronominal adverbs
 - Demonstrative: *here, there, now, then*
 - Interrogative / relative: *where, when, how, why*
 - Indefinite: *somewhere, sometime, sometimes, somehow, anywhere, anytime, anyhow, anyway, anyways*
 - Total: *everywhere, always*
 - Negative: *nowhere, never*

- Prepositions (>60; tagged corpus?)
 - *aboard, about, above, across, after, against, ago, along, alongside, amid, among, amongst, around, as, astride, at, atop, before, behind, below, beneath, beside, besides, between, beyond, by, despite, de, down, during, en, except, for, from, in, inside, into, lest, like, minus, near, next, notwithstanding, of, off, on, onto, opposite, out, outside, over, par, past, per, plus, post, since, through, throughout, 'til, till, to, toward, towards, under, underneath, unlike, until, unto, up, upon, versus, via, vs., with, within, without, worth*
 - `grep 'IN' wsj.mrg | perl -pe 's/^(.*?\(IN(.*?)\).*$/\$1/; $_=lc($_)' | sort -u | less`

- Conjunctions
 - Coordinating: *and, both, but, either, et, less, minus, 'n, 'n', neither, nor, or, plus, so, times, v., versus, vs., yet*
 - Subordinating: *albeit, although, because, 'cause, if, neither, since, so, than, that, though, 'til, till, unless, until, whereas, whether, which, while*
- Particles
 - *yes, no, not, n't, to* (infinitival)
- Found in corpus:
 - 256 closed-class types (out of 307 anticipated) **TODO: update these numbers**
 - 413,914 occurrences (33% of total tokens)



Open-Class Words

- Now there is a nice list of some 32,000 open-class words. What remains is to read them all and sort them out manually 😊
 - Nouns (including proper nouns)
 - Adjectives (including those derived from proper nouns)
 - Verbs (except for auxiliaries and modals)
 - Adverbs
 - (Interjections)

- What else can help us?



Most Frequent Open-Class Words

said	7141	shares	1444
new	3257	president	1431
company	3078	years	1426
year	2753	trading	1415
market	2648	sales	1331
says	2467	only	1188
stock	2002	business	1171
also	1867	such	1164
other	1808	york	1129
share	1798	group	1102
last	1482	time	1032



Most Frequent Open-Class Words

said	7141	shares	1444
new	3257	president	1431
company	3078	years	1426
year	2753	trading	1415
market	2648	sales	1331
says	2467	only	1188
stock	2002	business	1171
also	1867	such	1164
other	1808	york	1129
share	1798	group	1102
last	1482	time	1032

Plural Nouns / 3rd Person Verbs

year	2753	years	1426	4179
company	3078	companies	1020	4098
new	3257	news	424	3680
say	878	says	2467	3345
market	2648	markets	621	3269
stock	2002	stocks	800	2802
other	1808	others	263	2071
price	929	prices	1016	1945
sale	483	sales	1331	1814
last	1482	lasts	8	1490
month	624	months	844	1468
president	1431	presidents	22	1453
business	1171	businesses	267	1438


Total 4448 pairs

Gerunds / Present Participles

market	2648	marketing	211	2859
stock	2002	stocking	2	2004
trade	525	trading	1415	1940
share	1798	sharing	9	1807
last	1482	lasting	9	1491
bank	955	banking	220	1175
time	1032	timing	33	1065
say	878	saying	172	1050
make	739	making	286	1025
price	929	pricing	59	988
sell	603	selling	353	956 ...
even	905	evening	35	940
get	572	getting	201	773 ...

Total 1927 pairs

Tagged Corpus Available?

- Having a tagged corpus does not necessarily mean we have a morphological analyzer, so it still could make sense to construct one
- Now it's trivial to distinguish nouns from verbs, adjectives etc., even if they overlap
- Still, we may need some information not encoded in the tags
- Example: declension class ("pattern") of  Czech nouns:
 - NNF* = **feminine noun** \Rightarrow 4 declension classes:

<i>žena</i> "woman"	-a	-y	-ě	-u	-o	-ě	-ou	-y	-o	-ám	-y	-y	-ách	-ami
<i>růže</i> "rose"	-e	-e	-i	-i	-e	-i	-í	-e	-í	-ím	-e	-e	-ích	-emi
<i>píseň</i> "song"	-o	-ě	-i	-o	-i	-i	-í	-ě	-í	-ím	-ě	-ě	-ích	-ěmi
<i>kost</i> "bone"	-o	-i	-i	-o	-i	-i	-í	-i	-í	-em	-i	-i	-ech	-mi

And So On...

- Using similar heuristics, gradually classify more and more word forms
 - Obviously, not everything can be captured this way
 - Some sets of pairs have multiple interpretations
 - For some words no heuristics exist
 - Or the other member of the pair has not occurred in the corpus
- **Semi-supervised:**
 - You don't know what word form belongs where
 - However, you know how the suffixes look like
- **Unsupervised:**
 - You don't even know the set of affixes
 - However, you know (or assume) that the morphology is concatenative (prefix* stem+ suffix*)
 - Look at the corpus, try to find regularities

Unsupervised Morphemic Segmentation

- Morpho Challenge (shared task) since 2005
- Linguistica (John A. Goldsmith)
(<http://humanities.uchicago.edu/faculty/goldsmith/Linguistica2000/>)
- Morfessor (Mathias Creutz & Krista Lagus)
(<http://www.cis.hut.fi/projects/morpho/>)
- ParaMor (Christian Monson)
(<http://www.cslu.ogi.edu/~monsonc/ParaMor.html>)
- Affisix (Michal Hrušecký, MFF)
- Morseus (Daniel Zeman, MFF)
(<http://ufal.mff.cuni.cz/~zeman/projekty/morseus/>)
- And many others...

Homework

- Pick a language, get data
- Extract lexicon
- Details and data:
<http://ufal.mff.cuni.cz/~zeman/vyuka/morfosynt/lab-lexicon/index.html>
- Deadline:
Thursday November 12, 23:59 CET