# 05 - Indexing and Searching Very Large Texts IA161 Natural Language Processing in Practice 

M. Jakubíček<br>NLP Centre, FI MU, Brno

November 21, 2023
(1) Indexing
(2) Searching

## Searching big text corpora

Corpus:

- positional attributes - word form, lemma, PoS tag, ...
- structures and structure attributes - documents (e.g. with author, id, year, ...), paragraph, sentence
- searching: Manatee/Bonito/Sketch Engine
- http://corpora.fi.muni.cz
- https://app.sketchengine.eu
- SQL unsuitable (independent rows)


## Searching big text corpora

- data too big to be stored in memory
- data too big to be searched sequentially
$\Rightarrow$ preprocessing needed (indexing, alias corpus compilation)
- key decisions are:
- trade off between compile-time (preprocessing) and run-time
- trade off between in memory and off-memory processing


## Zipf's law I



## Zipf's law II

- may be simplified to inductive definition:


## Zipf's law (simplified)

frequency of the $n$-th element $f_{n} \approx \frac{1}{n} \cdot f_{1}$

- $\Rightarrow$ frequency is inversely proportional to the rank according to frequency
- $\Rightarrow$ one needs really large corpora to capture all the variety of many language phenomena


## Zipf's law III

| Word | $\downarrow$ Frequency ? |  | Word | $\downarrow$ Frequency ? |  |
| :---: | :---: | :---: | :---: | :---: | :---: |
| 1 the | 174,935,080 | $\cdots$ | 11 i | 23,989,001 | $\cdots$ |
| 2 of | 88,596,331 | $\cdots$ | 12 on | 20,237,809 | $\cdots$ |
| 3 and | 80,072,865 | -•• | 13 with | 19,230,246 | - |
| 4 to | 77,354,235 | $\cdots$ | 14 as | 19,076,719 | -. |
| 5 a | 59,410,937 | $\cdots$ | 15 be | 18,269,437 | ... |
| 6 in | 54,044,533 | $\cdots$ | 16 was | 16,505,649 | -.. |
| 7 that | 34,942,237 | $\cdots$ | 17 this | 16,475,525 | *. |
| 8 is | 34,190,792 | $\cdots$ | 18 you | 16,268,767 | -.. |
| 9 for | 27,849,928 | $\cdots$ | 19 are | 15,838,329 | - |
| 10 it | 24,609,587 | -•• | 20 by | 14,917,197 | -.. |


| Word | F Frequency? |  |
| :--- | ---: | :--- |
| 21 not | $14,421,888$ | $\ldots$ |
| 22 or | $13,599,707$ | $\ldots$ |
| 23 have | $13,540,277$ | $\ldots$ |
| 24 at | $13,282,835$ | $\ldots$ |
| 25 he | $12,821,501$ | $\ldots$ |
| 26 from | $12,285,435$ | $\ldots$ |
| 27 but | $11,049,177$ | $\ldots$ |
| 28 we | $10,997,497$ | $\ldots$ |
| 29 they | $10,388,785$ | $\ldots$ |
| 30 an | $10,182,791$ | $\ldots$ |

enTenTen2008, 3.2G tokens

## Zipf's law IV

About 1 billion words is enough to have enough evidence for single word units. But not for multiwords:

| word | Brown (1M) | BNC (100M) | enTenTen08 <br> $(2.7 G)$ | enTenTen15 <br> $(15.7 G)$ |
| :---: | :---: | :---: | :--- | :--- |
| carbonation | 0 | 5 | 429 | 2,817 |
| weird phrase | 0 | 0 | 14 | 34 |

## Building corpora

(1) content definition (what will it be used for? how do I get texts?)
(2) obtaining data (e.g. crawling)
(3) data cleaning (spam, boilerplate, duplicates)
(9) tokenization
(3) sentence segmentation
(6) further annotation (PoS tagging)
(1) corpus indexing and analysis

## Building corpora

(1) content definition (what will it be used for? how do I get texts?)
(2) obtaining data (e.g. crawling)
(3) data cleaning (spam, boilerplate, duplicates)
(9) tokenization
(3) sentence segmentation
(0) further annotation (PoS tagging)
(1) corpus indexing and analysis

## Corpus indexing

- text corpus is a database
- standard (=relational) database management systems are not suitable at all
- text corpus does not have relational nature
- special database management systems needed
$\Rightarrow$ Manatee


## REGEX <br> $\bigcirc \bigcap$ eat

- 

exactly one unspecified character
w.n win won wen wun wan
ca. cat car cap cab can
zero or more occurrences of the

| preceding character |  |
| :--- | :--- |
| a*h | h ah aah aaah aaaah |
| c.*ing | words starting c- and <br> ending -ing: cooking <br> error: missing character |
| cool | before asterisk <br> word starting c- |

## ?

preceding character is optional colou?r color colour bet?
be bet

not, brackets are compulsory
[^m]et pet get bet let
(but not met)
[^mpglet set let (but not met pet $g e t$ )
()
grouping, prioritizing
meter|re meter or re met(er|re) meter or metre
list, range - one character from content of the brackets

$$
\begin{aligned}
& \text { [mpgb]et } \begin{array}{l}
\text { met pet get bet } \\
\mathrm{m}[2-5] \\
{[0-9]^{*}} \\
{[\mathrm{~m} 2 \mathrm{~m} 3 \mathrm{~m} 4 \mathrm{~m} 5}
\end{array} \\
& {[\mathrm{a}-\mathrm{z}]} \\
& {[\mathrm{A}-\mathrm{Z}]^{*}} \\
& \begin{array}{ll}
\text { any number of digits } \\
\text { ony number of } \\
\text { any nercase letters: } \\
\text { uppercer } \\
\text { UNESCO, UK, WIEI }
\end{array} \\
& \\
& \begin{array}{ll}
\text { any number of } \\
\text { letters (but not } \\
\text { numbers) }
\end{array}
\end{aligned}
$$

OR - the characters to the left or those on the right

$$
\begin{array}{ll}
\text { get|met } & \text { get or met } \\
\text { met(er|re) } & \text { meter or metre }
\end{array}
$$


one or more occurences of the preceding character (compare *)
a+h
ah aah aaah aaah
hallot hallo halloo hallooo
$\}$
repetition
gr $\{2,4\}$ gre grre grrer $[A-Z](3) \quad 3$-letter acronyms (bla) $\{2,3\}$ blabla blablabla

escaping, removes or adds special meaning to a character
finds $3-1$ letter words
\(\backslash w . <br>

)| finds three dots |
| :--- |
| finds any letter |

## REGEX <br> character classes

These classes also cover non-English (Unicode) characters, e.g. ñ ř ç 香 $\nsim \mathcal{U}$
shortcut NOT

| [[:alpha:]] | any letter including Unicode | \w |
| :---: | :---: | :---: |
| [[:digit:]] | any digit, equivalent to [0-9] | $\backslash d$ |
| [ [:alnum:]] | any digit or letter including Unicode |  |
| [ [:lower:]] | any lowercase letter including Unicode |  |
| [ [:upper:]] | any uppercase letter including Unicode |  |
| [ [:punct:]] |  |  |
| [ [:space:]] | whitespace character space, new line, tab, carriage retum | $\backslash s$ |

## D 5 ए

k. * words starting with $k$
*k. * starting, containing or ending with $k$ (including just $k$ )
. +ment words ending with -ment (but not just ment)
$\left[{ }^{\wedge} x\right]+x\left[{ }^{\wedge} x\right]+\quad$ words containing $x$ but not starting or ending with it
[[:upper:]][[:lower:]]* words starting with one capital letter
[[:upper:]]* acronyms incl. unicode: EU, SRPŠ, ЖЭК
$[a-z] * \backslash d[a-z]$ * lowercase words containing a digit: face2face
(kilo|centi)? metre kilometre centimetre metre
dog|.*cat|mouse dog OR cat OR pussycat OR tomcat etc. OR mouse

Extended regex manual on http://ske.li/regex
For complete information, google "regular expressions cheat sheet or tutorial"
SKETCH
ENGiNE

## CQL cheat sheet

```
sample [lemma="go"]
syntax [lemma="work" & tag !="V.*"]
    [tag="N.*"][ ]{1,5}[tag="V.*"] within <s/>
```

\&. joins two or more conditions for the same token

| $\}$ | [word="ha"] (3) finds ha ha ha [tag="N.*"] $\{2,5\}$ finds $2,3,4$ or 5 nouns |
| :---: | :---: |
| $?$ | makes the preceding token optional [lc="new"] [1c="cheap"]? [1c="phone"] finds both new phone and new cheap phone |
| $\dagger$ | ma="accommodate"]\|ma="put"][lc="up"] finds accommodate or put up |
| () | the tokens inside behave as one group [1c="might"](%5B1c=%22as%22%5D%5B1c=%22well%22%5D)? [tag="V.*"] finds both might as well $g o$ and might go |
| * | unlimited (max. 100) repetitions of the preceding token <s> [ ]* [word=" ${ }^{\text {? }}$ "]</s> within <s/> finds sentences finishing with a question mark |
| $<>$ | used for structures such as documents, paragraphs and sentences: <s> beginning " </s> end " <s/> all |
| $\sim$ | searches for chop followed by carrot and its 15 most similar nouns (vegetables) <br> ma="chop"] [](0,3) ~15"carrot-n" |

## Default attribute

Makes queries easier to read. It is applied to each token without square brackets. This query
[lc="might" $]([1 \mathrm{c}=$ "as" $][1 \mathrm{c}=$ "well" $])$ ? [tag="v.*"]
can be simplified like this:

## surnt

word (lowercase) - "might" ("as" "well")? [tag="V.*"]

## within <s/>

ensures that the result is found only if it is inside the same sentence [tag="N.*"] [ ] [tag="V.*"] within <s/>
something shorter within something longer
finds something shorter only if it appears inside something longer, e.g. adjective technical but only if it appears inside a sequence of 3 adjectives [lc="technical"]within [tag="J.*"] (3)

## <s/> containing

finds sentences which contain something else
<s/> containing [tag="N.*"] [ ] [tag="v.*"]
something longer containing something shorter
finds something longer only if it contains something shorter
[tag="N.*"] [ ] 1,3 )[tag="V.*"] containing [lc="often"]
Only the thing before within/containing will be highlighted in red as KWIC. Using the other operator to change the highlighting.

## meet

finds something (staff) only if something else (member) is to the left/right
def.attr. lemma (meet "staff" "member" -1 2)

## Structures

<doc> <p> <s> beginning of a document, paragraph, sentence " </doc> </p> </s> end of a structure " <doc/> <p/> <s/> the whole structure

## Structures in CQL

<doc> [] finds the first token of each document
[lc="local"] within <doc region="UK"/> finds the word local in documents whose region is UK
full CQL manual online: http://ske.li/cql

CQL cheat sheet
SKETCH
ENGiNE
most frequently used tags in the

## English tagset

| N. * | noun |
| :---: | :---: |
| V. * | verb |
| J. * | adjective |
| RB. ? | adverb |
| PP. ? | pronoun |
| CC | conjunction |
| IN | preposition |
| DT | determiner |
| CD | numeral |
| RP | particle |

Click (i) to display the tagset of your corpus.
Boot Camp English

| CC | coordinating conjunction | and |
| :---: | :---: | :---: |
| CD | cardinal number | 1, third |
| CDZ | possesive pronoun | one's |
| DT | determiner | the |
| EX | existential there | there is |
| FW | foreign word | d'hoevre |
| IN | preposition, subord. conjunction | in, of, like |
| IN/t <br> hat | that as subordinator | that |
| JJ | adjective | green |
| JJR | adjective, comparative | greener |
| JJS | adjective, superlative | greenest |
| LS | list marker | 1) |
| MD | modal | could, will |
| NN | noun, singular or mass | table |
| NNS | noun plural | tables |
| NNSZ | possesive noun plural | people's, women's |
| NNZ | possesive noun, singular or mass | year's, world's |
| NP | proper noun, singular | John |
| NPS | proper noun, plural | Vikings |
| NPSZ | possesive proper noun, plural | Boys', <br> Workers' |
| NPZ | possesive noun, singular | Britain's, God's |
| PDT | predeterminer | both the boys |
| PP | personal pronoun | 1, he, it |
| PPZ | possessive pronoun | my, his |
| RB | adverb (however, naturally, here) |  |
| RBR | adverb, comparative | better |
| RBS | adverb, superlative | best |
| RP | particle | give up |


| SENT | Sentence-break, punctuation | !? |
| :---: | :---: | :---: |
| SYM | Symbol | / [ = * |
| TO | infinitive 'to' | to go |
| UH | interjection | Ahh! |
| VB | verb be, base form | be |
| VBD | verb be, past tense | was, were |
| VBG | verb be, gerund/present participle | being |
| VBN | verb be, past participle | been |
| VBP | verb be, sing. present, non-3d | am, are |
| VBZ | verb be, 3 rd person sing. present | is |
| VH | verb have, base form | have |
| VHD | verb have, past tense | had |
| VHG | verb have, gerund/present participle | having |
| VHN | verb have, past participle | had |
| VHP | verb have, sing. present, non-3d | have |
| VHZ | verb have, 3rd person sing. present | has |
| vV | verb, base form | take |
| vVD | verb, past tense | took |
| VVG | verb, gerund/present participle | taking |
| vvN | verb, past participle | taken |
| VVP | verb, present, not 3rd person | take |
| VVZ | verb, 3rd person sing. present | takes |
| WDT | wh-determiner | which |
| WP | wh-pronoun | who, what |
| WPZ | possessive wh-pronoun | whose |
| z | possessive ending | S |

## Vertical text

and other attributes

Well, Theresa May didn't photograph apples with her "Apple" this May.

www.sketchengine.eu

Vertical text
with structures:
sentence and glue

| word | 139 | lc | lemma | lemma.lc | lempos |
| :---: | :---: | :---: | :---: | :---: | :---: |
| <6; |  |  |  |  |  |
| Well | RB | well | well | well | well-a |
| <g/> |  |  |  |  |  |
| , | , | , | , | , | ,-x |
| Theresa | NP | theresa | Theresa | theresa | Theresa-n |
| May | NP | may | May | may | May-n |
| did | MD | did | do | do | do-v |
| <g/> |  |  |  |  |  |
| n't | RB | n't | not | not | not-a |
| photograph | VV | photograph | photograph | photograph | photograph-v |
| apples | NN | apples | apple | apple | apple-n |
| with | IN | with | with | with | with-i |
| her | PPZ | her | her | her | her-d |
| " | * | * | " | " | *-x |
| <g/> |  |  |  |  |  |
| Apple | NP | apple | Apple | apple | Apple-n |
| <g/> |  |  |  |  |  |
| " | * | * | " | " | *-x |
| this | DT | this | this | this | this-x |
| May | NP | may | May | may | May-n |
| , | SENT | . | . | . | --x |
| </s> |  |  |  |  |  |

## Indexing corpora in Manatee

Key data structures for a positional attribute:

- lexicon
- because operations on numbers are just so much faster than on strings
- corpus text
- to iterate over positions
- inverted (reversed) index
- to give fast access to positions for a given value


## How to store integer numbers

- given Zipf's distribution: fixed-length storing very inefficient
- variable-length more complicated but yielding much smaller and quicker indices
- variable-length bit-wise universal Elias' codes: gamma, delta codes
- cf. Huffman coding


## How to store integer numbers

BNC: 112,345,722 tokens

- whole data 4 -byte encoding: 449,382,888 bytes
- whole data delta difference coding: 189 MB
- the: frequency $5,415,707$ (4.8 \%)
- 4-byte integer encoding: 21,662,828 bytes
- delta difference coding: 5,213,473 bytes ( $24 \%$ )
enTenTen20: 43,125,207,462 tokens
- whole data 4 -byte encoding: $172,500,829,848$ bytes
- whole data delta difference coding: 75 GB
- the: frequency $1,915,064,722$ (4.44 \%)
- 4-byte integer encoding: 7,660,258,888 bytes
- delta difference coding: 1,877,715,456 bytes (24.5 \%)


## Indexing corpora in Manatee

Structures and operations:

- operations in between: string (str) - number (id) - position (poss)
- lexicon building: $\Rightarrow$ word-to-id mapping $\Rightarrow$ operations on numbers, not strings $\Rightarrow$ id2str, str2id
- inverted index: id2poss
- corpus text: pos2id
- yields transitively also pos2str, str2poss


## Searching corpora in Manatee

- key idea: operations on sorted forward-only streams of positions
- FastStream - single position stream
- RangeStream - stream of position pairs (structures: from position, to position)
- = Corpus Query Language (Christ and Schulze, 1994)
- positions and positional attributes: [attr="value"]
- structures and structural attributes: <str attr="value">
- example:

$$
\begin{aligned}
& \text { [word=".*ing" \& tag="V.*"] } \\
& \text { <doc id=" } 20[5-9] . * "
\end{aligned}
$$

- established a within <str/> query:
[tag="N.*"]+ within <s/>
and alternative meet/union query:

$$
\begin{aligned}
& \text { (meet }[\text { lemma="take"] }[\text { tag }=" \mathrm{~N} . * "]-5+5) \\
& \text { (union (meet ...) }(\text { meet ...)) }
\end{aligned}
$$

## CQL in Manatee/Bonito

- ehnancements and differences to the original CQL syntax
- within <query> and containing <query>
- meet/union (sub)query
- inequality comparisons
- frequency function


## within/containing queries

- searching for particles:

$$
\begin{aligned}
& \text { [tag="PR.*"] within [tag="V.*"] [tag="ATO"]? } \\
& \text { [tag="AJO"]* [tag="(PR.?lN.*)"] [tag="PR.*"] within } \\
& \text { <s/> }
\end{aligned}
$$

- searching for a Czech idiom "hnout někomu žlučí" ("to get somebody's goat" ):
word-by-word translated as: hnout "move" [V, infinitive] někomu "somebody" [ N , dative] žlučí "bile" [ $N$, instrumental].
<s/> containing [lemma="hnout"] containing [tag=".*c3.*"] containing [word="žlučí"]


## within/containing queries

- structure boundaries: begin: <str>, whole structure: <str/>, end: </str>
- changes: within <str> not allowed anymore, use within <str/>


## meet/union queries

- combined with regular query: <s/>

```
containing (meet [lemma="have"] [tag="P.*"] -5 5) containing (meet [tag="N.*"] [lemma="blue"])
```

- changes: meet/union queries can be used on any position, they can contain labels and no MU keyword is required (and deprecated): (meet 1:[] 2:[]) \& 1.tag = 2.tag


## Inequality comparisons

- former comparisons allowed only equality and its negation: [attr="value"] [attr!="value"]
- inequality comparisons implemented: [attr<="value"] [attr>="value"] [attr!<="value"] [attr!>="value"]
- intended usage:
[tag="AJ.*"] [tag="NN.*"] within <doc year>="2009">
- sophisticated comparison performed on the attribute value: <doc id<="CC20101031B"> matches e.g. BB20101031B, CC20091031B, CC20101030B CC20101031A.


## Fixed string comparisons

- normally the CQL values are regular expressions
- sometimes this is not desirable (batch processing needs escaping of metacharacters)
- new $==$ and $!==$ operator introduced for fixed strings comparison
- no escaping needed except for '"' and '\'
- examples: ".", "\$"," " matches a single dot, dollar sign and tilda, respectively, " $\backslash \mathrm{n}$ " matches a backslash followed by the character n ,


## Frequency function

- a frequency constraint allowed in the global conditions part of CQL: 1:[tag="PP.*"] 2:[tag="NN.*"] \& f(1.word) > 10


## Performance evaluation

Table: Query performance evaluation - corpora legend: ○ BNC (110M tokens),

- BiWeC (version with 9.5G tokens), * Czes (1.2G tokens)

| query | \# of results | time (m:s) |
| :--- | :---: | :---: |
| $\circ$ [lemma="time"] | 179,321 | 0.07 |
| $\circ$ [lemma="t.*"] | $14,660,881$ | 3.12 |
| $\circ$ Ex: particles | $1,219,973$ | 33.36 |
| $\bullet$ Ex: particles | $97,671,485$ | $32: 26.48$ |
| * Ex: idioms | 66 | $1: 6.86$ |
| $\circ$ Ex: meet/union | 3 | 8.47 |
| $\bullet$ Ex: meet/union | 1457 | $7: 13.12$ |

## CQL query evaluation

Example: [tag="ADJ"] [(word="record" | word="process") \& tag="NOUN"] within <doc year="2012"/>



## Conclusions

- special database management systems for processing text corpora needed
- trade-offs between compile-time and run-time, in-memory and off-memory
- CQL
- Manatee


## Assignment

