

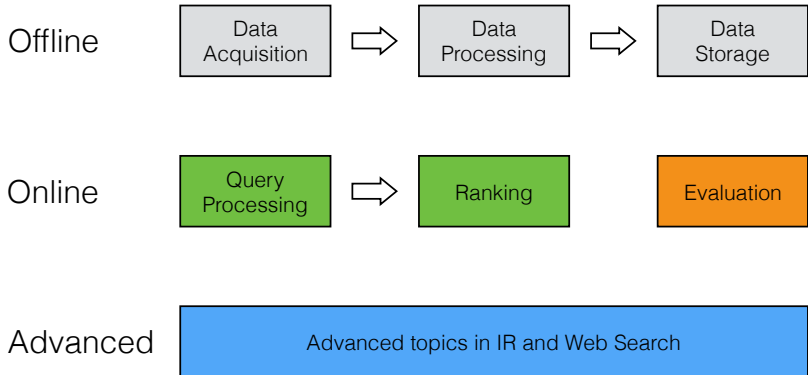
Information Retrieval

Data Processing

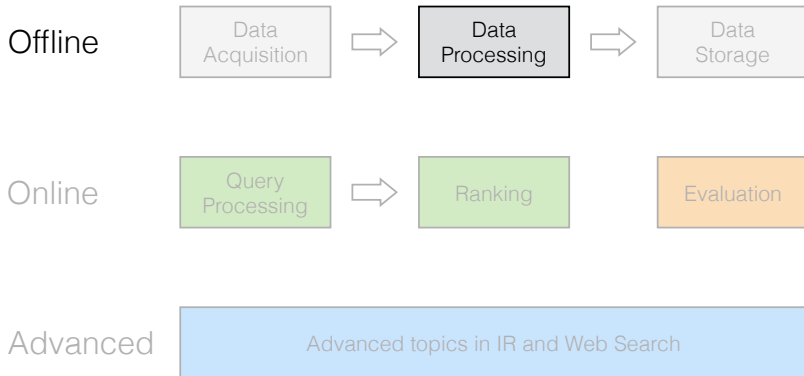
Ilya Markov
i.markov@uva.nl

University of Amsterdam

Course overview



This lecture



Outline

- 1 Data processing pipeline
- 2 Stemming
- 3 Dealing with phrases
- 4 Zipf's and Heaps' laws
- 5 Summary

Outline

- 1 Data processing pipeline
- 2 Stemming
- 3 Dealing with phrases
- 4 Zipf's and Heaps' laws
- 5 Summary

Data processing pipeline

- ① Remove white-spaces and punctuation
- ② Convert terms to lower-case
- ③ Remove stop-words
- ④ Convert terms to their stems
- ⑤ Deal with phrases
- ⑥ Apply language-specific processing rules

Example

- ① To prepare a text for indexing, one needs to split it into tokens, remove stop-words and perform stemming.
- ② to prepare a text for indexing one needs to split it into tokens remove stop words and perform stemming
- ③ prepare indexing needs split
tokens remove stop perform stemming
- ④ prepar index need split
token remov stop perform stem

Stop-word removal

- Frequency-based
 - Set a frequency threshold f
 - Remove words with the frequency higher than f
- Dictionary-based
 - Create a dictionary of stop-words
 - Remove words that occur in this dictionary

Outline

- 1 Data processing pipeline
- 2 Stemming**
- 3 Dealing with phrases
- 4 Zipf's and Heaps' laws
- 5 Summary

Stemming

- ① Algorithmic
- ② Dictionary-based
- ③ Hybrid

Algorithmic stemming (Porter stemmer)

Step 1a:

- Replace *ses* by *ss* (e.g., stresses → stress).
- Delete *s* if the preceding word part contains a vowel not immediately before the *s* (e.g., gaps → gap but gas → gas).
- Replace *ied* or *ies* by *i* if preceded by more than one letter, otherwise by *ie* (e.g., ties → tie, cries → cri).
- If suffix is *us* or *ss* do nothing (e.g., stress → stress).

Step 1b:

- Replace *eed*, *eedly* by *ee* if it is in the part of the word after the first non-vowel following a vowel (e.g., agreed → agree, feed → feed).
- Delete *ed*, *edly*, *ing*, *ingly* if the preceding word part contains a vowel, and then if the word ends in *at*, *bl*, or *iz* add *e* (e.g., fished → fish, pirating → pirate), or if the word ends with a double letter that is not *ll*, *ss* or *zz*, remove the last letter (e.g., falling → fall, dripping → drip), or if the word is short, add *e* (e.g., hoping → hope).
- Whew!

Croft et al., "Search Engines, Information Retrieval in Practice"

Dictionary-based stemming

- Store lists of related words in a dictionary
- Can recognize the relation between “is”, “be”, “was”
- New-words problem

Hybrid stemming (Krovetz stemmer)

- Approach
 - ① Check the word in a dictionary
 - ② If present, either leave it as is
 - ③ If not present, check for suffixes that could be removed
 - ④ After removal, check the dictionary again
- Produces words not stems
- Comparable effectiveness with the Porter stemmer

Stemming example

Original text:

Document will describe marketing strategies carried out by U.S. companies for their agricultural chemicals, report predictions for market share of such chemicals, or report market statistics for agrochemicals, pesticide, herbicide, fungicide, insecticide, fertilizer, predicted sales, market share, stimulate demand, price cut, volume of sales.

Porter stemmer:

document describ market strategi carri compani agricultur chemic report predict market share chemic report market statist agrochem pesticid herbicid fungicid insecticid fertil predict sale market share stimul demand price cut volum sale

Krovetz stemmer:

document describe marketing strategy carry company agriculture chemical report prediction market share chemical report market statistic agrochemic pesticide herbicide fungicide insecticide fertilizer predict sale stimulate demand price cut volume sale

Croft et al., "Search Engines, Information Retrieval in Practice"

Outline

- 1 Data processing pipeline
- 2 Stemming
- 3 Dealing with phrases**
- 4 Zipf's and Heaps' laws
- 5 Summary

Example

To be or not to be...

Dealing with phrases

- ① Detect **noun phrases** using a part-of-speech tagger
 - sequences of nouns
 - adjectives followed by nouns
- ② Detect phrases at the query processing time
 - Use an index with word positions
 - Will be discussed next
- ③ Use frequent **n-grams**, e.g., bigrams and trigrams

Example noun phrases

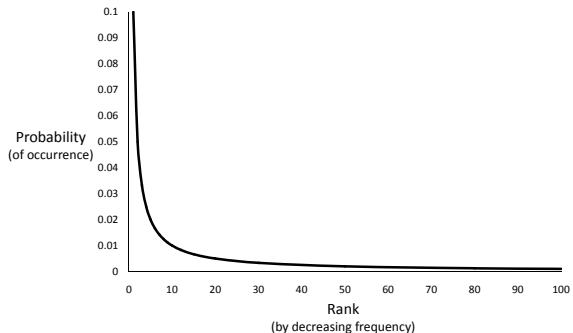
TREC data		Patent data	
<i>Frequency</i>	<i>Phrase</i>	<i>Frequency</i>	<i>Phrase</i>
65824	united states	975362	present invention
61327	article type	191625	u.s. pat
33864	los angeles	147352	preferred embodiment
18062	hong kong	95097	carbon atoms
17788	north korea	87903	group consisting
17308	new york	81809	room temperature
15513	san diego	78458	seq id
15009	orange county	75850	brief description
12869	prime minister	66407	prior art
12799	first time	59828	perspective view
12067	soviet union	58724	first embodiment
10811	russian federation	56715	reaction mixture
9912	united nations	54619	detailed description
8127	southern california	54117	ethyl acetate
7640	south korea	52195	example 1
7620	end recording	52003	block diagram
7524	european union	46299	second embodiment
7436	south africa	41694	accompanying drawings
7362	san francisco	40554	output signal
7086	news conference	37911	first end
6792	city council	35827	second end
6348	middle east	34881	appended claims
6157	peace process	33947	distal end
5955	human rights	32338	cross-sectional view
5837	white house	30193	outer surface

Croft et al., "Search Engines, Information Retrieval in Practice"

Outline

- 1 Data processing pipeline
- 2 Stemming
- 3 Dealing with phrases
- 4 Zipf's and Heaps' laws**
- 5 Summary

Zipf's law

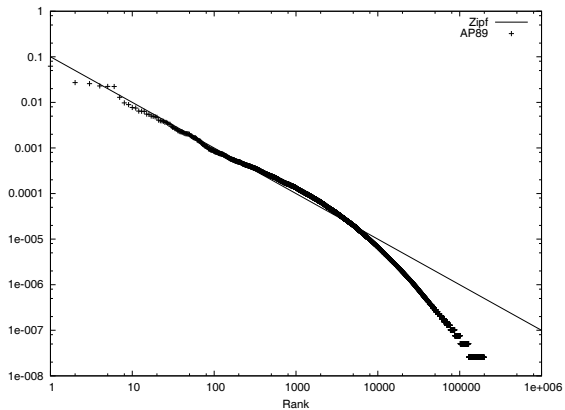


$$\text{rank} \cdot \text{freq} = \text{const}$$

$$\text{rank} \cdot P_r = \text{const}'$$

Croft et al., "Search Engines, Information Retrieval in Practice"

Zipf's law vs. real data



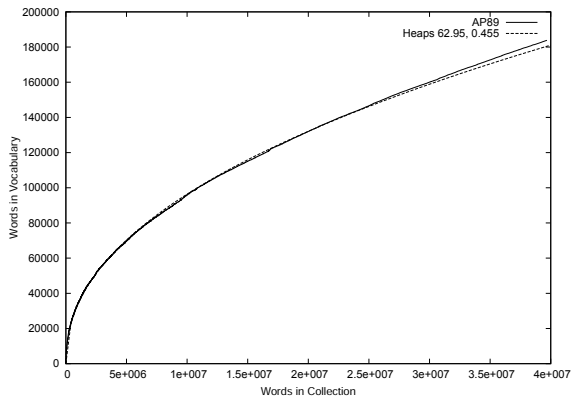
Croft et al., "Search Engines, Information Retrieval in Practice"

Zipf's law example

<i>Word</i>	<i>Freq.</i>	<i>r</i>	<i>P_r(%)</i>	<i>r.P_r</i>	<i>Word</i>	<i>Freq</i>	<i>r</i>	<i>P_r(%)</i>	<i>r.P_r</i>
the	2,420,778	1	6.49	0.065	has	136,007	26	0.37	0.095
of	1,045,733	2	2.80	0.056	are	130,322	27	0.35	0.094
to	968,882	3	2.60	0.078	not	127,493	28	0.34	0.096
a	892,429	4	2.39	0.096	who	116,364	29	0.31	0.090
and	865,644	5	2.32	0.120	they	111,024	30	0.30	0.089
in	847,825	6	2.27	0.140	its	111,021	31	0.30	0.092
said	504,593	7	1.35	0.095	had	103,943	32	0.28	0.089
for	363,865	8	0.98	0.078	will	102,949	33	0.28	0.091
that	347,072	9	0.93	0.084	would	99,503	34	0.27	0.091
was	293,027	10	0.79	0.079	about	92,983	35	0.25	0.087
on	291,947	11	0.78	0.086	i	92,005	36	0.25	0.089
he	250,919	12	0.67	0.081	been	88,786	37	0.24	0.088
is	245,843	13	0.65	0.086	this	87,286	38	0.23	0.089
with	223,846	14	0.60	0.084	their	84,638	39	0.23	0.089
at	210,064	15	0.56	0.085	new	83,449	40	0.22	0.090
by	209,586	16	0.56	0.090	or	81,796	41	0.22	0.090
it	195,621	17	0.52	0.089	which	80,385	42	0.22	0.091
from	189,451	18	0.51	0.091	we	80,245	43	0.22	0.093
as	181,714	19	0.49	0.093	more	76,388	44	0.21	0.090
be	157,300	20	0.42	0.084	after	75,165	45	0.20	0.091
were	153,913	21	0.41	0.087	us	72,045	46	0.19	0.089
an	152,576	22	0.41	0.090	percent	71,956	47	0.19	0.091
have	149,749	23	0.40	0.092	up	71,082	48	0.19	0.092
his	142,285	24	0.38	0.092	one	70,266	49	0.19	0.092
but	140,880	25	0.38	0.094	people	68,988	50	0.19	0.093

Croft et al., "Search Engines, Information Retrieval in Practice"

Heaps' law



$$vocab = const \cdot words^{\beta}$$

Croft et al., "Search Engines, Information Retrieval in Practice"

Outline

- 1 Data processing pipeline
- 2 Stemming
- 3 Dealing with phrases
- 4 Zipf's and Heaps' laws
- 5 Summary

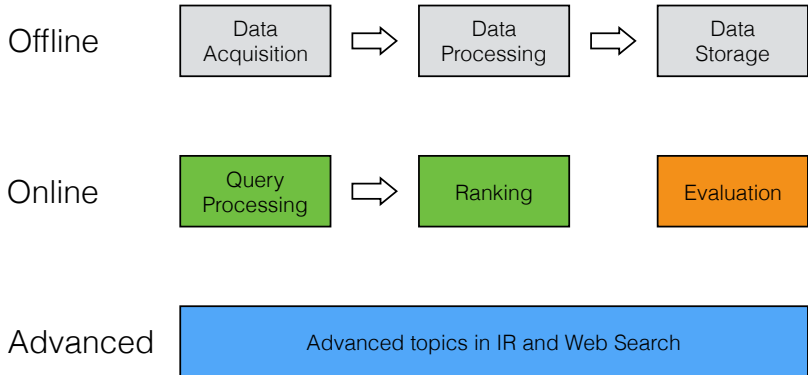
Data processing summary

- Stop-word removal
- Stemming
 - Algorithmic (Porter stemmer)
 - Dictionary-based
 - Hybrid (Krovetz stemmer)
- Dealing with phrases
- Zipf's and Heaps' laws

Materials

- Croft et al., Chapter 4
- Manning et al., Chapter 2.2

Course overview



Next lecture

