Big Data: Data Wrangling Boot Camp
R Sentiment Analysis

Chuck Cartledge, PhD

24 February 2018

Table of contents (1 of 1)

What are we going to cover?

- Look to the future
- Talk briefly about sentiment analysis
- Address the polyglot of computer languages
- Talk about our sentiment analysis system
- **Data wrangle** tweets using R

Things that we will be doing.

1. **Data wrangle** tweets using R
2. Conduct sentiment analysis on tweets
3. Look at the sentiments in different ways

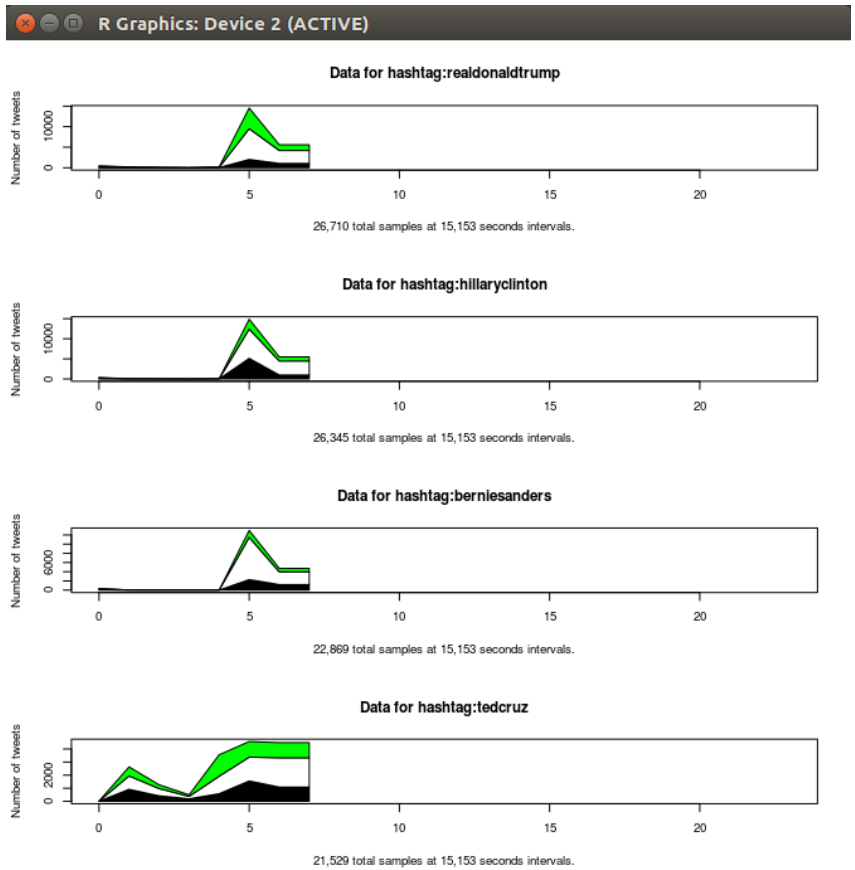

Sentiments over time.

Sentiment by sending device

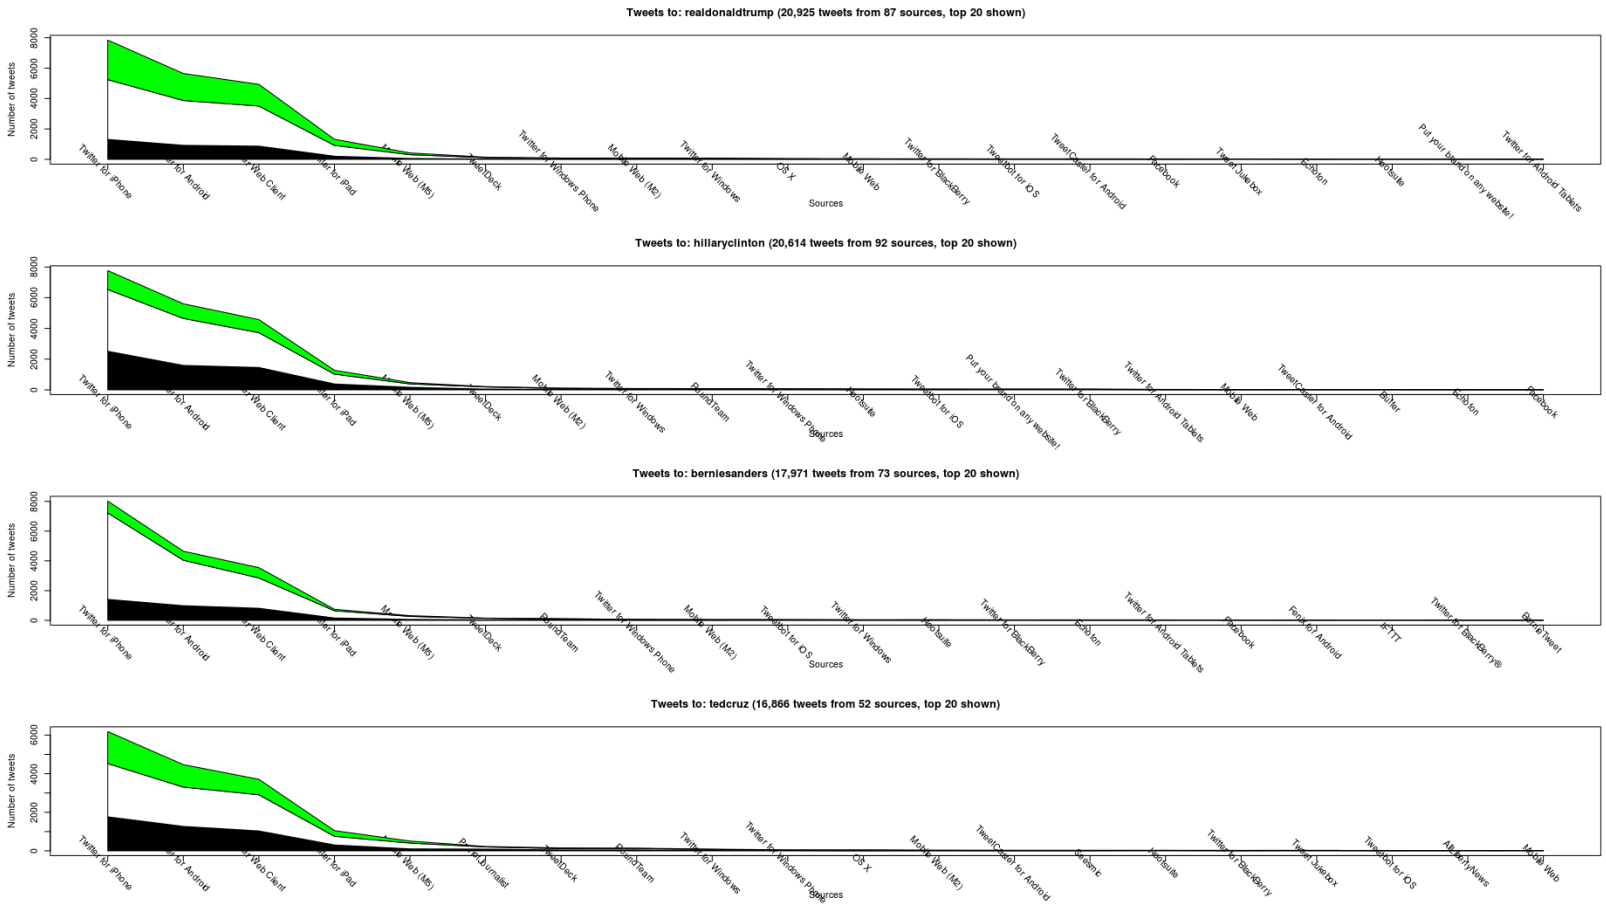

Sentiment by geographic location

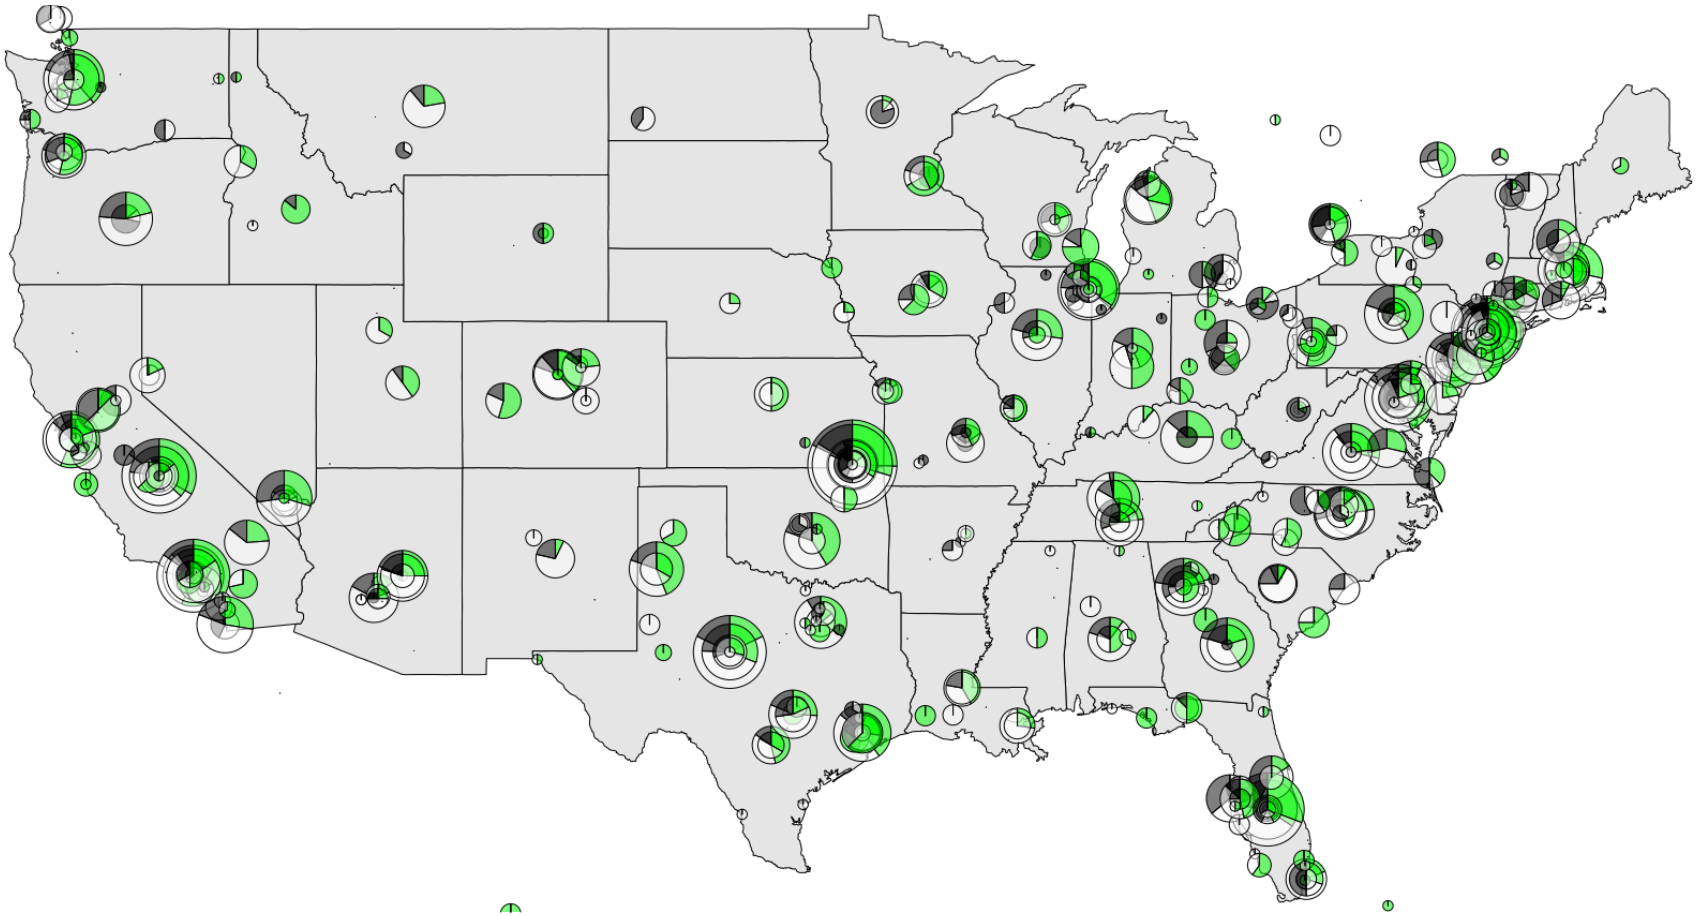

How we'll get to the images

We'll walk before we run.

- Start with a replay file
- **Data wrangle** using the library file
- Go live and download live tweets

Data flows through the backend, into the database, out the frontend.

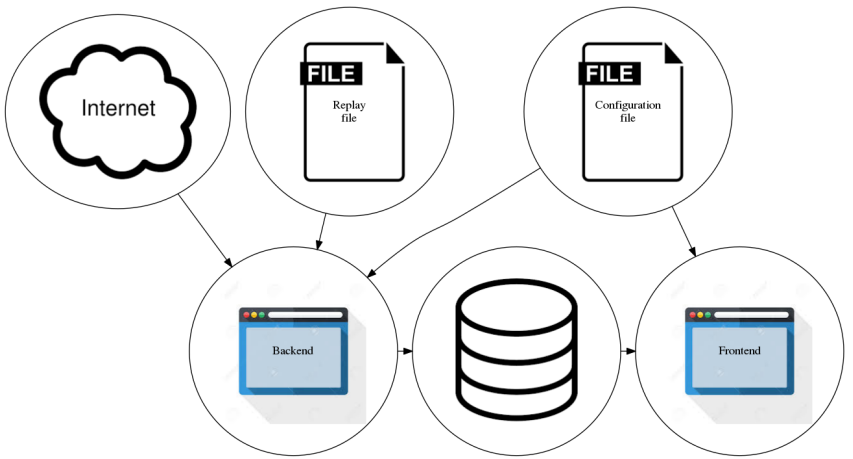

The software design document (attached) contains lots of details.

Same image.

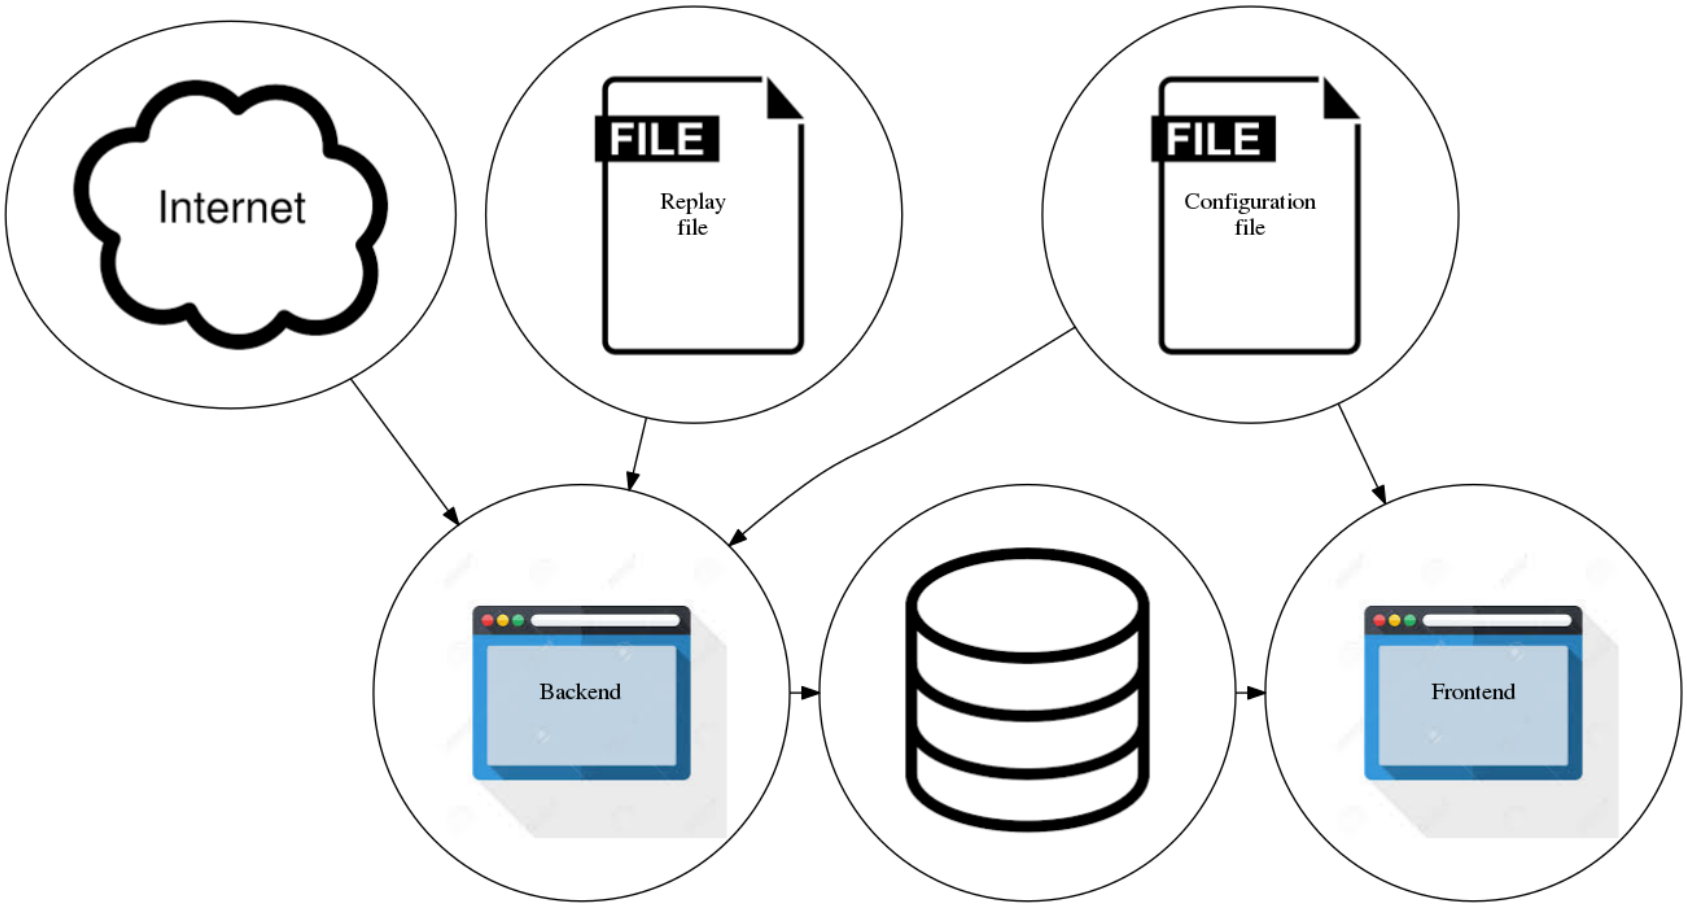

The software design document (attached) contains lots of details.

What some of the files do:

- Replay file – previously recorded tweets
- Configuration file – directives used by both backend and frontend script files
- Shared library file – routines that are common to backend and frontend script files
- Backend file – script file that populates the database from the replay file, or the Internet
- Frontend file – script file that extracts data from the database and presents results

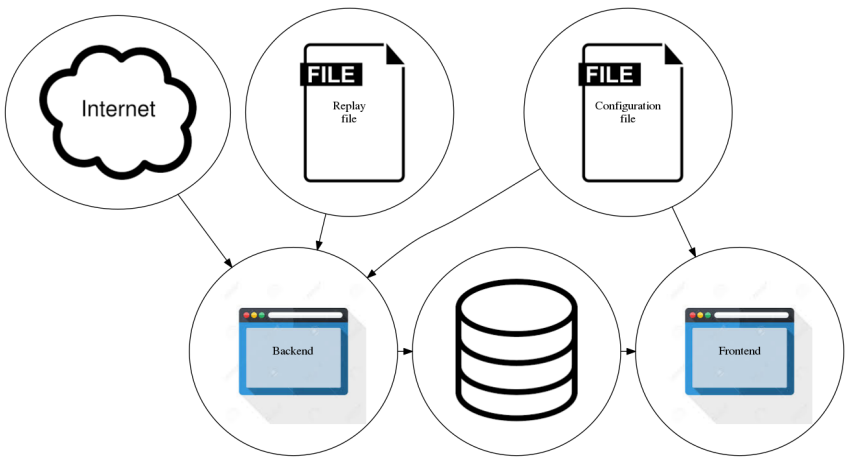

A working definition

"Sentiment analysis (also known as opinion mining) refers to the use of natural language processing, text analysis and computational linguistics to identify and extract subjective information in source materials. Sentiment analysis is widely applied to reviews and social media for a variety of applications, ranging from marketing to customer service."

W. Staff [3]

More formal definitions

"The field of opinion mining and sentiment analysis is well-suited to various types of intelligence applications. Indeed, business intelligence seems to be one of the main factors behind corporate interest in the field."

Pang and Lee [2]

"Sentiment analysis, also called opinion mining, is the field of study that analyzes peoples opinions, sentiments, evaluations, appraisals, attitudes, and emotions towards entities such as products, services, organizations, individuals, issues, events, topics, and their attributes."

Liu [1]

Our approach to sentiment analysis

We will:

1. Search the "twitterverse" for tweets using specific hashtags
2. Tokenize each tweet
3. **Data wrangle** each token
4. Remove all stop words from the tokens
5. Count number of positive and negative tokens
6. Compute the positive, negative, or neutral sentiment for the tokens
7. Display the results

Our approach is language agnostic.

A "Universe" of words

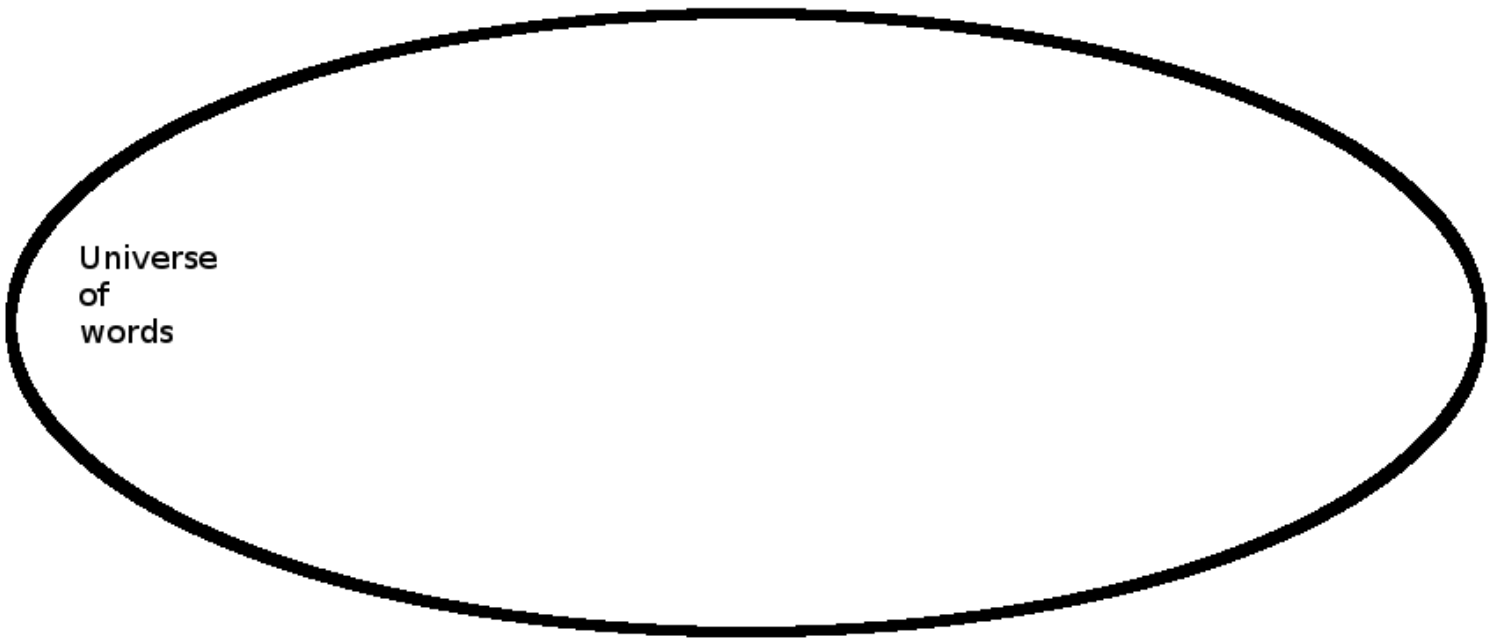

Some words are "Positive"

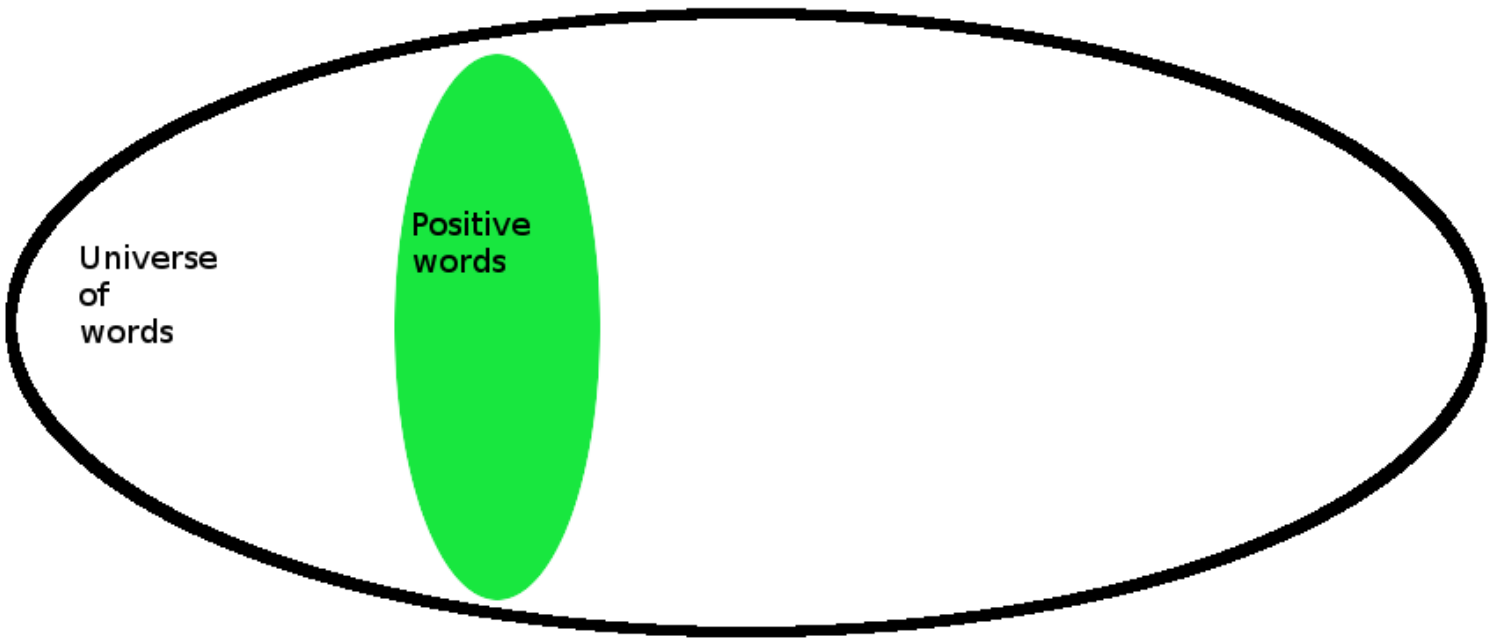

Some words are "Negative"

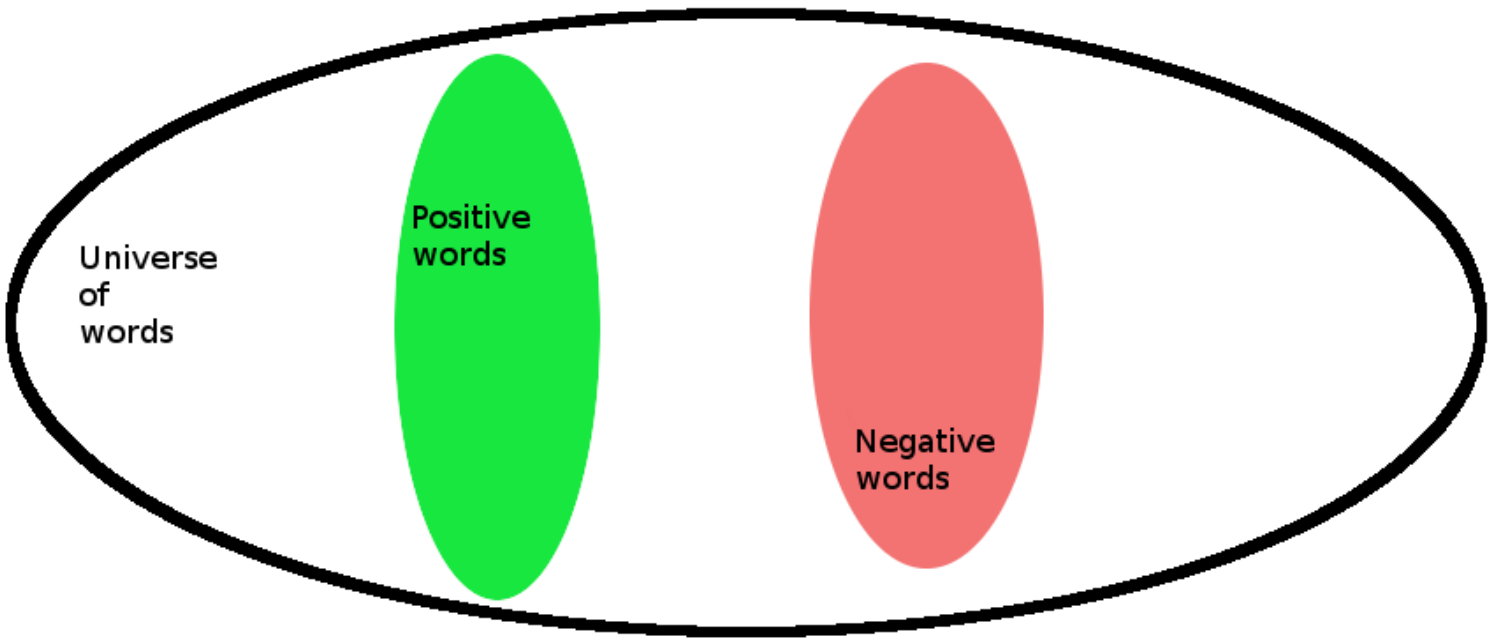

A tweet will/may have positive and negative words

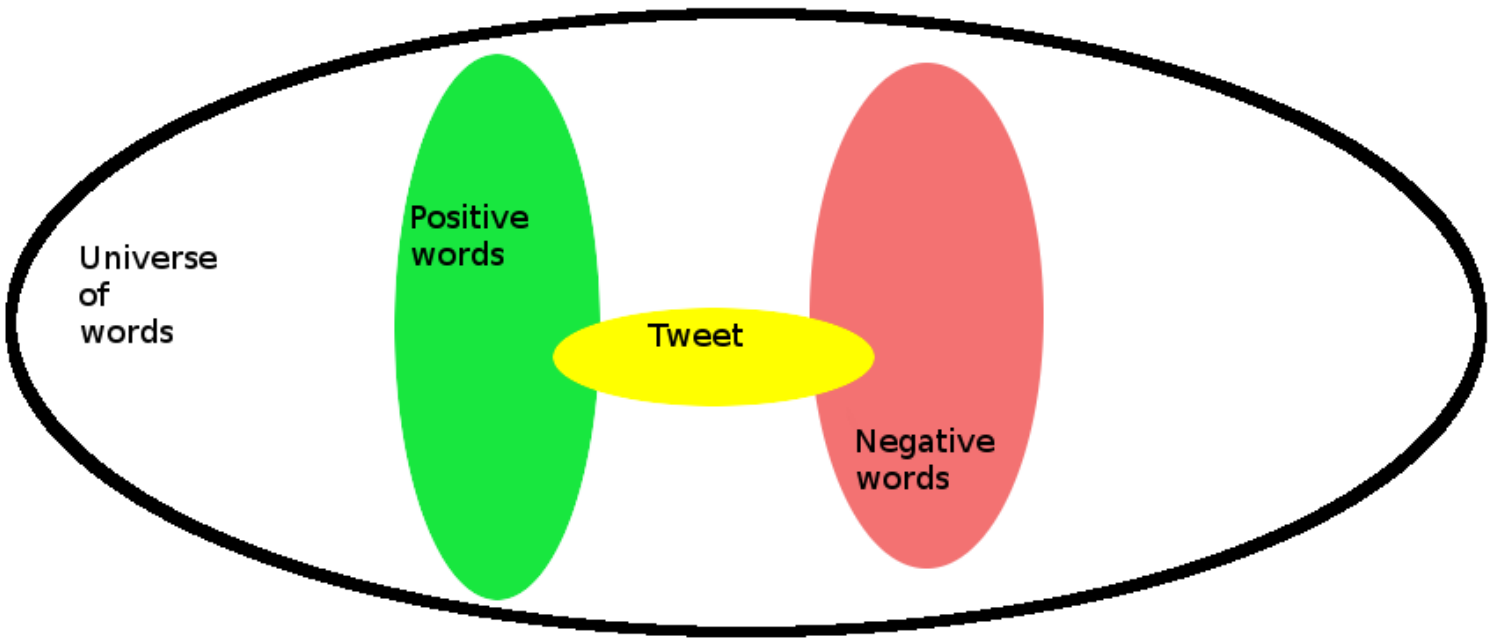

Some words we don't care about

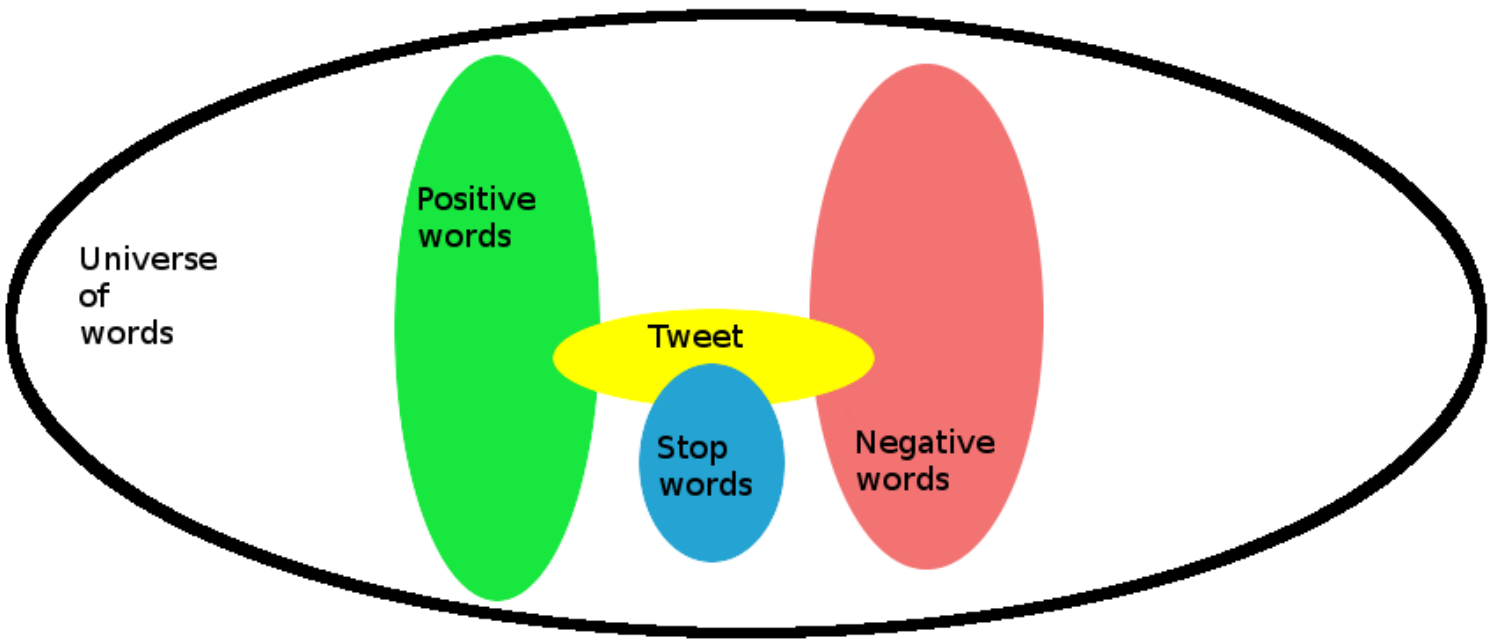

Mechanically this is what we are doing

The steps are:

1. Break the tweet into tokens
2. Remove stop words from the tokens
3. Compute the percentage of remaining tweet tokens that are positive
4. Compute the percentage of remaining tweet tokens that are negative
5. Classify the tokens as positive, negative, or neutral

Mathematically this is what we are doing

The steps are:

$$
\begin{aligned}
tokens &= \{\text{words in tweet}\} \\
tokensLessStop &= tokens - stopWords \\
positivePart &= positiveWords \cap tokensLessStop \\
negativePart &= negativeWords \cap tokensLessStop
\end{aligned}
$$

$$
classification = \begin{cases}
positive, & \text{if} positiveThreshold \leq \frac{positivePart}{tokensLessStop} \\
& \text{AND} \\
& \frac{negativePart}{tokensLessStop} < negativeThreshold \\
negative, & \text{if} negativeThreshold \leq \frac{negativePart}{tokensLessStop} \\
& \text{AND} \\
& \frac{positivePart}{tokensLessStop} < positiveThreshold \\
neutral, & \text{otherwise}
\end{cases}
$$

A software design document

The document contains:

1. Overall system design
2. Algorithms used through out the system
3. Details about the configuration file
4. Details about the database tables

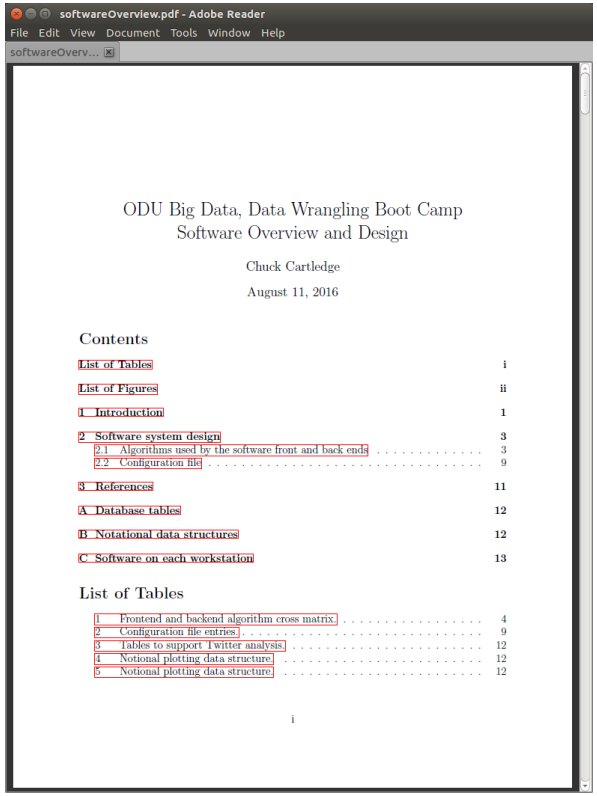
softwareOverview.pdf - Adobe Reader
File Edit View Document Tools Window Help
softwareOverv...

ODU Big Data, Data Wrangling Boot Camp
Software Overview and Design

Chuck Cartledge

August 11, 2016

Contents

List of Tables

i

The file is attached.

What will happen to the tweets when we do nothing.

A histogram of how often a token occurs.

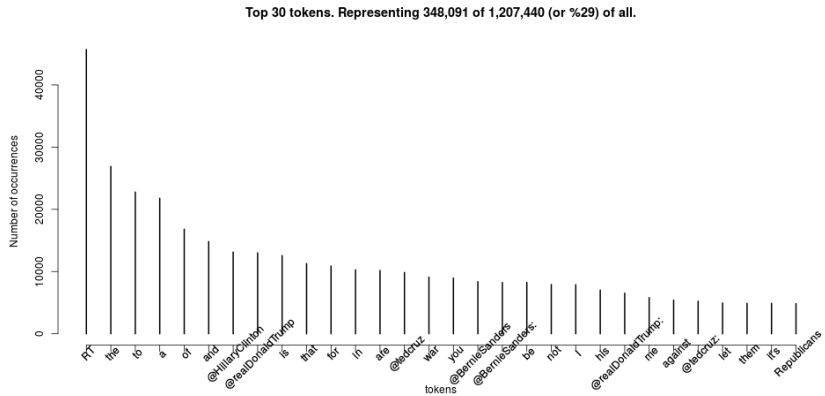

Same image.

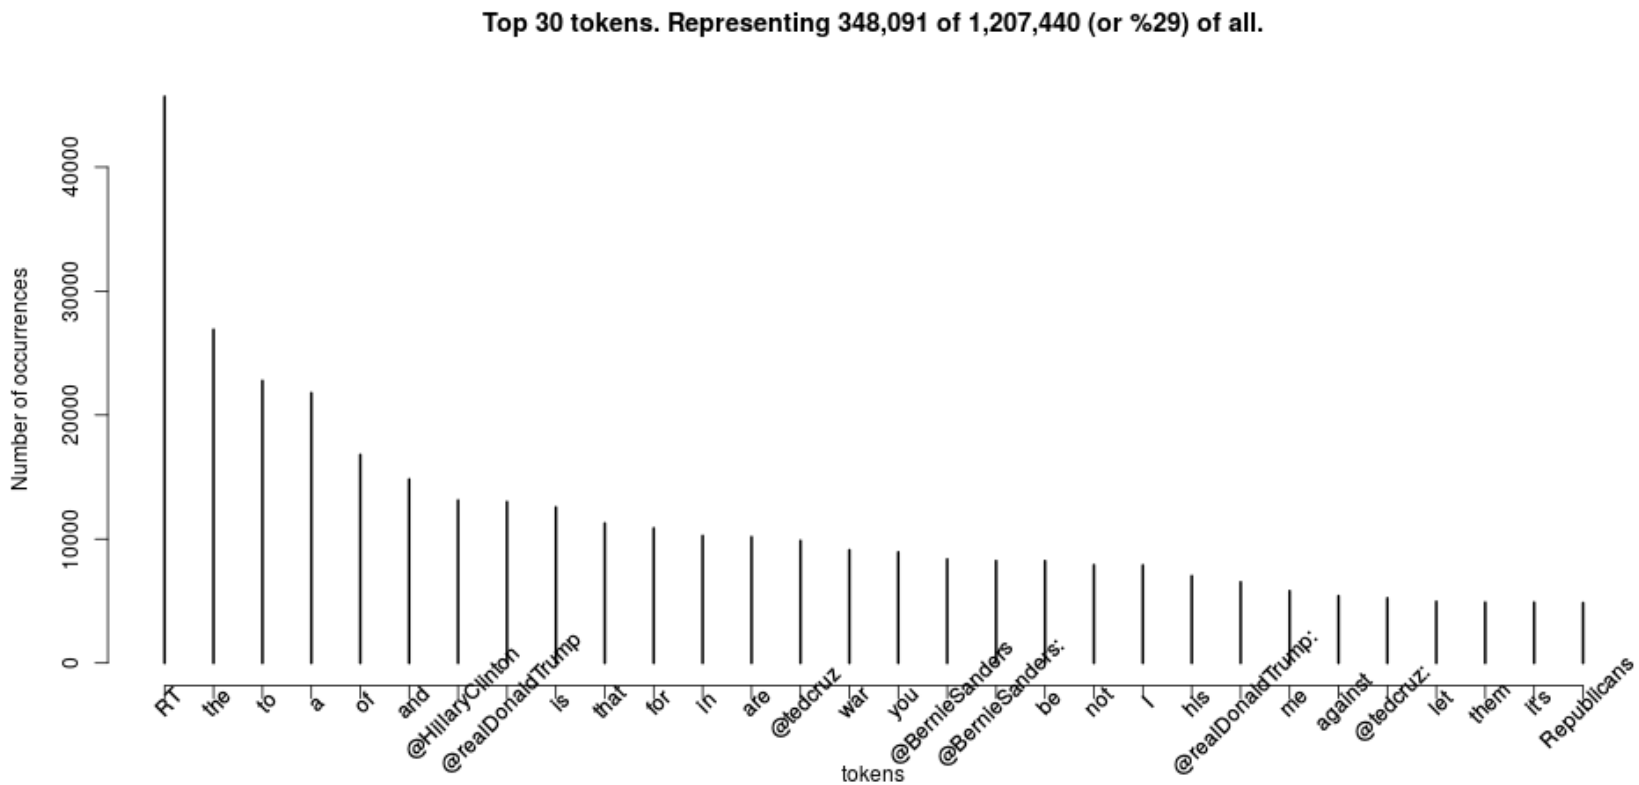
Top 30 tokens. Representing 348,091 of 1,207,440 (or %29) of all.
Number of occurrences
0
10000
20000
30000
40000
RT
the
to
a
of
and
@HillaryClinton
@realDonaldTrump
is
that
for
in
are
@tedcruz
war
you
@BernieSanders
@BernieSanders:
be
not
I
his
@realDonaldTrump:
me
against
@tedcruz:
let
them
it's
Republicans
tokens

What will happen to the tweets when we make everything the same case

Changing case is easy, unless they are emojois.

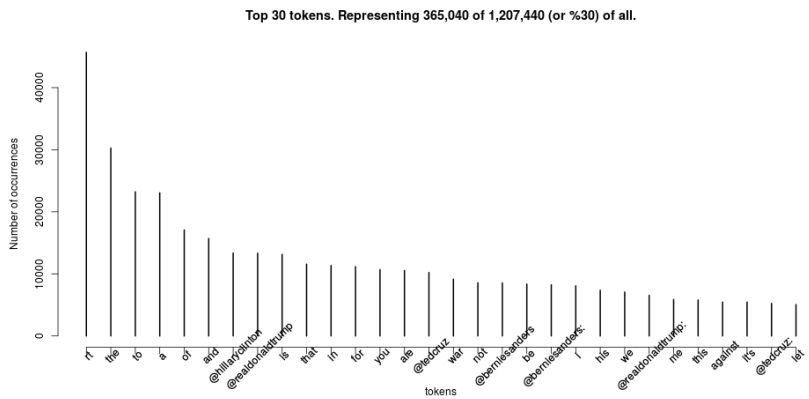

Same image.

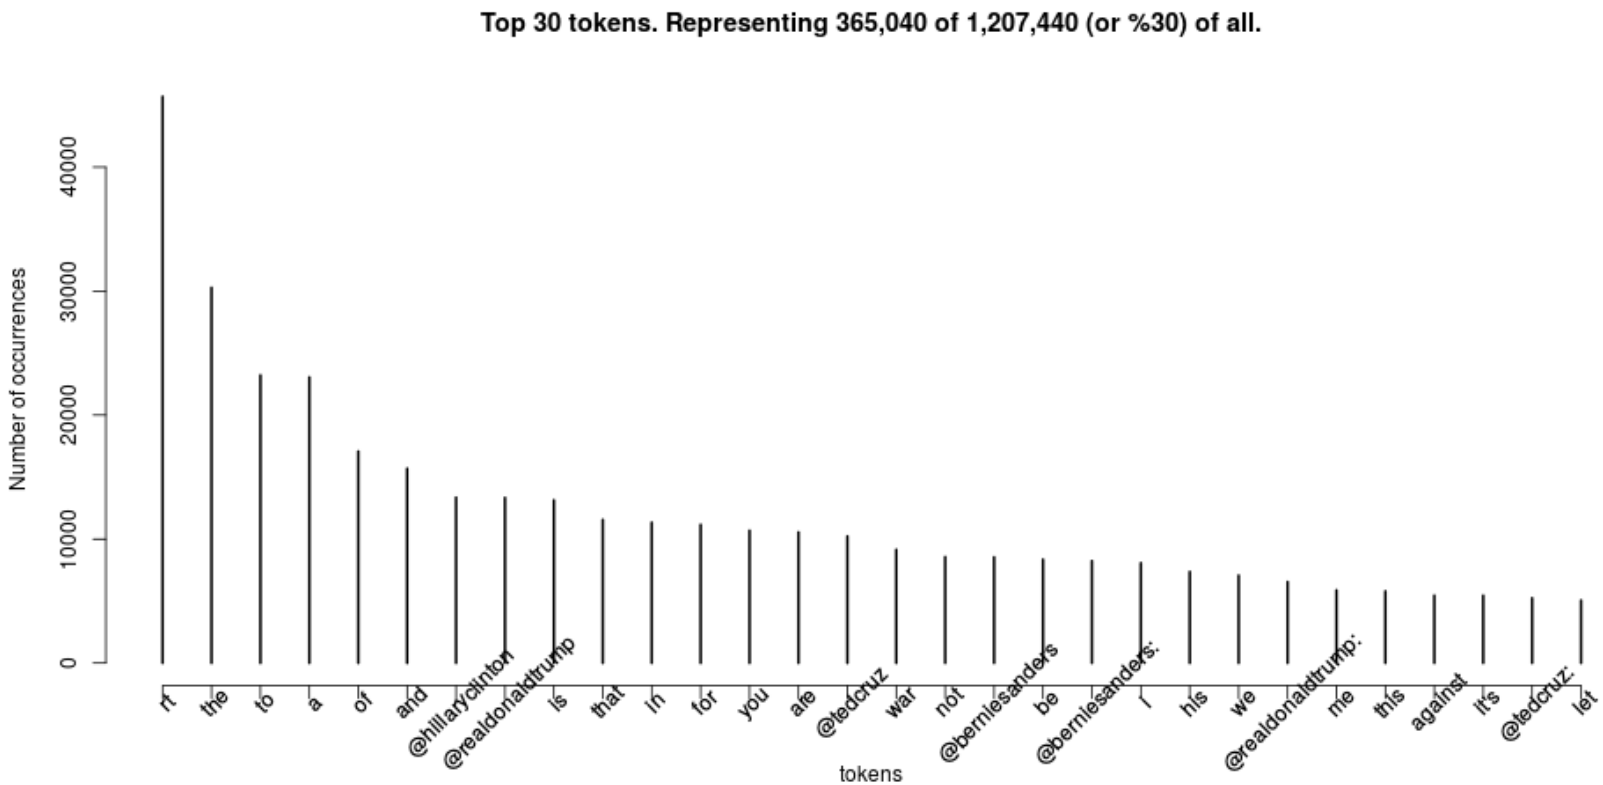
Top 30 tokens. Representing 365,040 of 1,207,440 (or %30) of all.
Number of occurrences
0
10000
20000
30000
40000
rt
the
to
a
of
and
@hillaryclinton
@realdonaldtrump
is
that
in
for
you
are
@tedcruz
war
not
@berniesanders
be
@berniesanders:
i
his
we
@realdonaldtrump:
me
this
against
it's
@tedcruz:
let
tokens

What will happen to the tweets we remove non-ASCII

We need to talk about ASCII and the rest of the character systems.

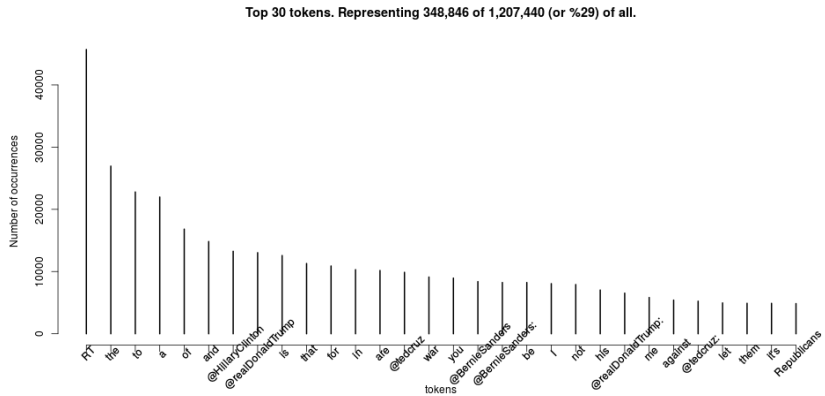

Same image.

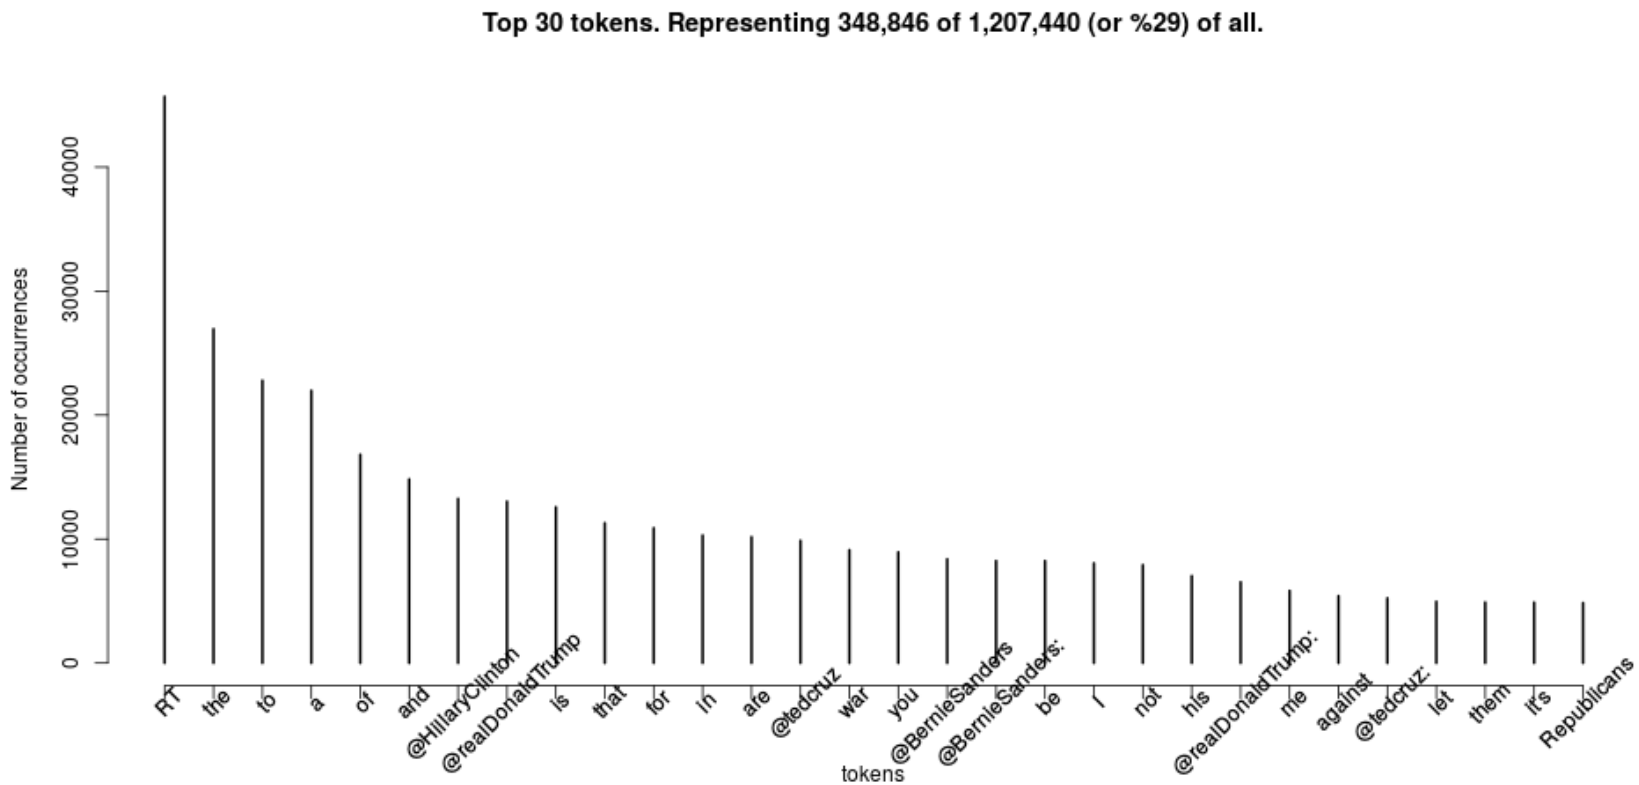
Top 30 tokens. Representing 348,846 of 1,207,440 (or %29) of all.
Number of occurrences
0
10000
20000
30000
40000
RT
the
to
a
of
and
@HillaryClinton
@realDonaldTrump
is
that
for
in
are
@tedcruz
war
you
@BernieSanders
@BernieSanders:
be
I
not
his
@realDonaldTrump:
me
against
@tedcruz:
let
them
it's
Republicans
tokens

What will happen to the tweets when we remove "stop words"

Stop words are words/tokens that have no use in whatever we are doing. Stop words are domain specific.
And we change case, remove non-ASCII, remove punctuation, etc.

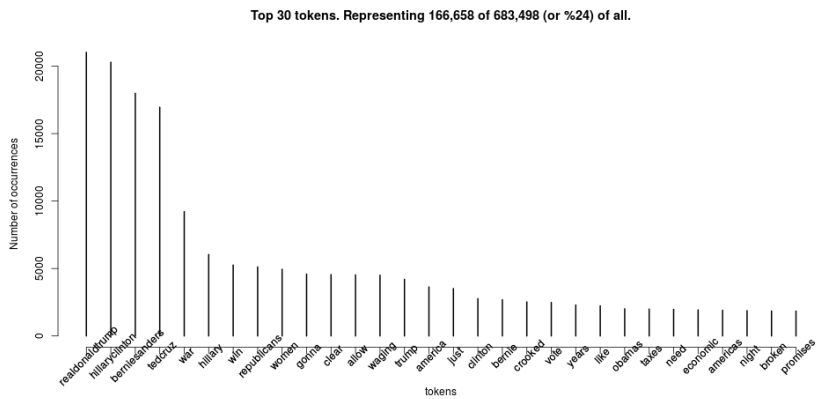

Same image.

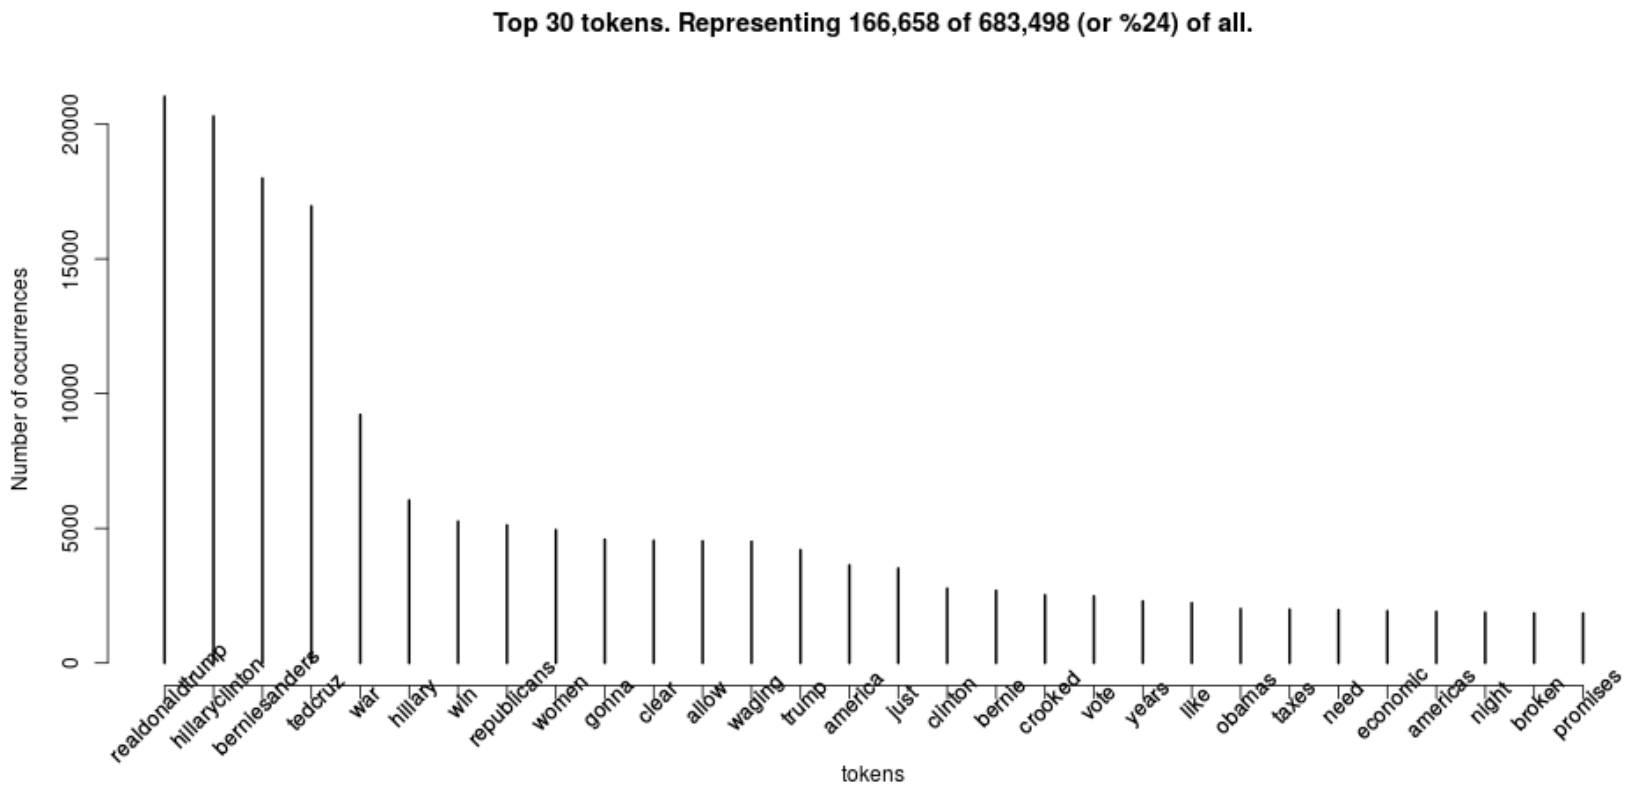
Top 30 tokens. Representing 166,658 of 683,498 (or %24) of all.
Number of occurrences
0
5000
10000
15000
20000
realdonaldtrump
hillaryclinton
berniesanders
tedcruz
war
hillary
win
republicans
women
gonna
clear
allow
waging
trump
america
just
clinton
bernie
crooked
vote
years
like
obamas
taxes
need
economic
americas
night
broken
promises
tokens

Using Rstudio

1. Set the session working directory
2. Load the following files into the editor:
   1. `checkPostgres.R`
   2. `backEnd.R`
   3. `config.txt`
   4. `frontEnd.R`
   5. `library.R`
3. Run the `checkPostgres.R` script to ensure PostGres is configured correctly
4. Run the `backEnd.R` script to populate the database with "canned" data
5. Run the `frontEnd.R` script to analyze the tweets
6. Modify the `config.txt` file, and re-run back and front ends

Q & A time.

Q: How did you get into artificial intelligence?
A: Seemed logical – I didn't have any real intelligence.

What have we covered?

- A preview of today's activities
- An overview of the sentiment analysis

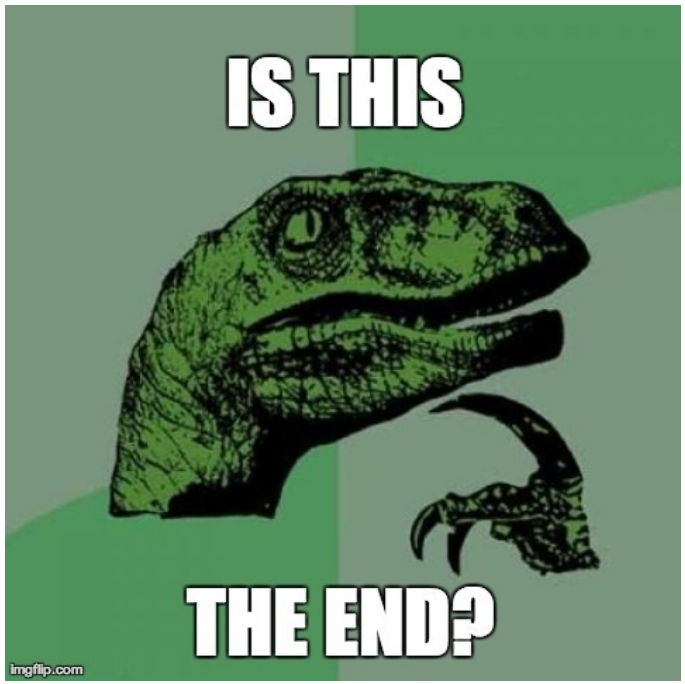

Next: Hands on analysing tweets with R.

References (1 of 1)

[1] Bing Liu, Sentiment analysis and opinion mining, 2012.

[2] Bo Pang and Lillian Lee, Opinion mining and sentiment analysis, 2008.

[3] Wikipedia Staff, Sentiment analysis, `https://en.wikipedia.org/wiki/Sentiment_analysis`, 2016.

Files of interest

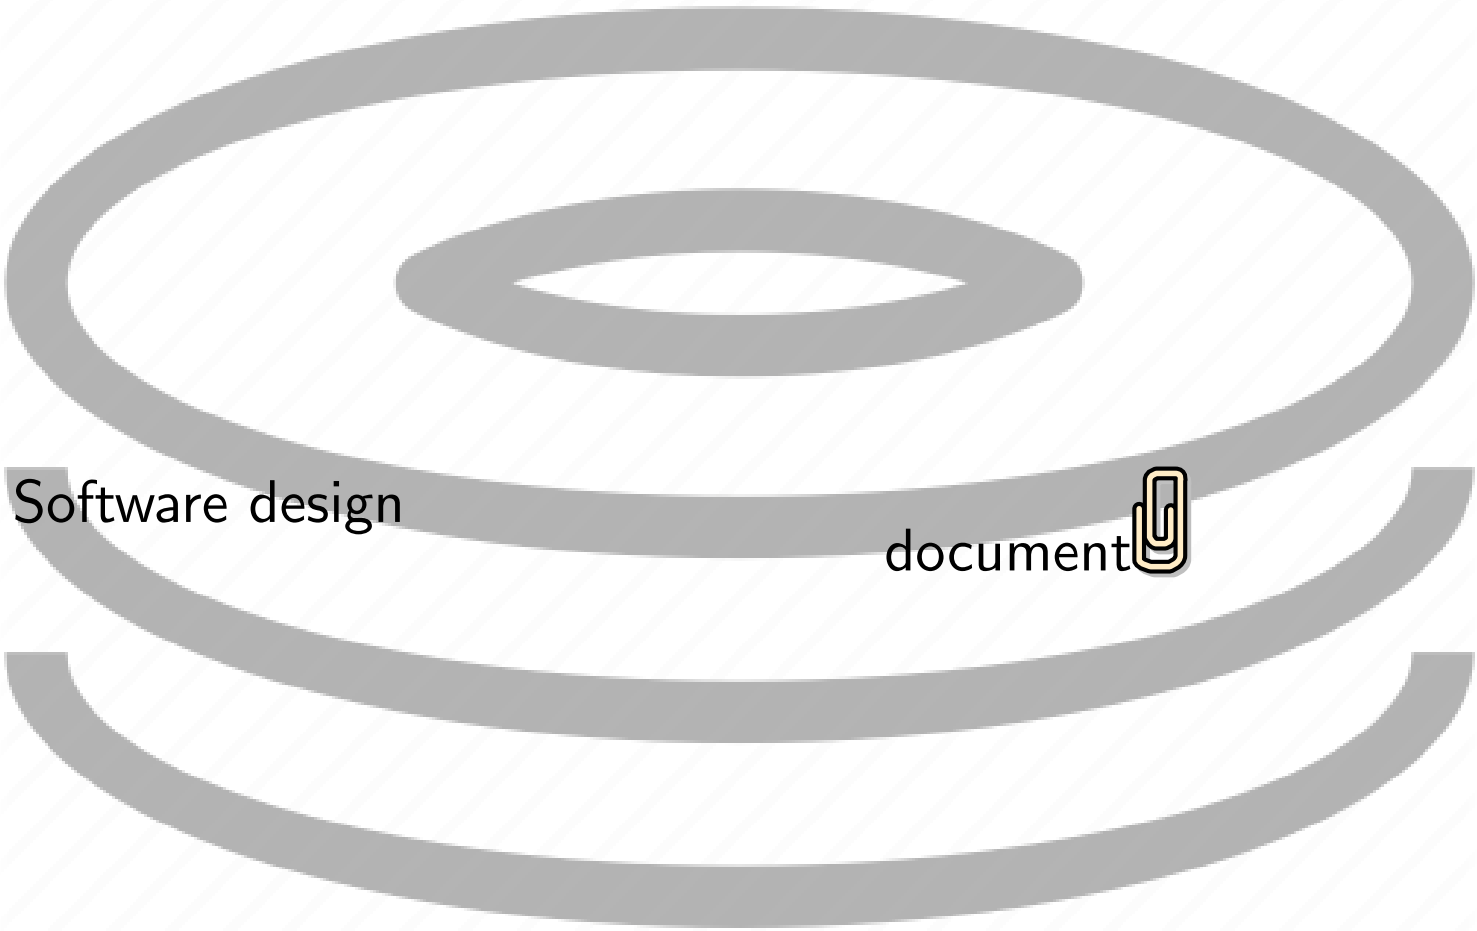

1. Software design document