



---

# **Automatic Detection of Political Opinions in Tweets**

**Diana Maynard and Adam Funk**  
**University of Sheffield, UK**

# What is Opinion Mining?

- OM is a recent discipline that studies the extraction of opinions using IR, AI and/or NLP techniques.
- More informally, it's about extracting the opinions or sentiments given in a piece of text
- Also referred to as Sentiment Analysis (though technically this is a more specific task)
- Social media provides a great medium for people to share opinions
- This provides a useful source of unstructured information that may be useful to others (e.g. companies and their rivals, other consumers...)
- But the problem lies in getting this useful information out.

# It's about finding out what people think...




# Venus Williams causes controversy...



# Opinion mining exposes these insights

### Get Tips by Timeline ?



■ Positive  
■ Negative  
■ Neutral

Move slider to search

Day(s) Just Current Latest  
earlier prior

Mostly from 4 days -Now

### Discover Tips by Topic ?

[pie](#) [venue](#) [dress is](#) [court](#)  
[williams](#) [crust](#) [tennis](#)

## 😊 Positive Tips ?



I see a lot of people on here bitching about **Venus Williams' dress**. I really like it. She looks really good in it.

#WTA #Tennis



**Venus Williams' dress** looks as if it might be worn by Hannibal Lecter to keep him from eating people. Only in a cheery color.

## 😞 Negative Tips ?



Watching the Australian Open and am quite disturbed by **Venus Williams' dress** which has completely unnecessary holes all around



Srsly, this **Venus Williams dress** is as bad as her backhand, which is all elbows.

## 😐 Neutral Tips ?



Someone took **Venus Williams' dress** and police have cordoned off the area. #ausopen



Soooo I just saw **Venus Williams' dress** and now I want a Belgian Waffle from @MaxBrenner.

# Online social media sentiment apps

- There are lots of these apps available:
  - Twitter sentiment <http://twittersentiment.appspot.com/>
  - Twends: <http://twendz.waggeneratedstrom.com/>
  - Twittratr: <http://twitrratr.com/>
  - SocialMention: <http://socialmention.com/>
- Easy to search for opinions about famous people, brands and so on
- Hard to search for more abstract concepts, perform a non-keyword based string search
- e.g. to find opinions about Venus Williams' dress, you can only search on "Venus Williams" to get hits

# Opinion mining and social media

- Social media provides a wealth of information about a user's behaviour and interests:
  - *explicit*: John likes tennis, swimming and classical music
  - *implicit*: people who like skydiving tend to be big risk-takers
  - *associative*: people who buy Nike products also tend to buy Apple products
- While information about individuals isn't useful on its own, finding defined clusters of interests and opinions is

If many people talk on social media sites about fears in airline security, life insurance companies might consider opportunities to sell a new service
- This kind of predictive analysis is all about understanding your potential audience at a much deeper level - this can lead to improved advertising techniques such as personalised ads to different groups



# Analysing and preserving opinions

---

- Useful to collect, store and later retrieve public opinions about events and their changes or developments over time
- One of the difficulties lies in distinguishing what is important
- Opinion mining tools can help here
- Not only can online social networks provide a snapshot of such situations, but they can actually trigger a chain of reactions and events
- Ultimately these events might lead to societal, political or administrative changes



# The Royal Wedding leads to Pilates

- One of the biggest Royal Wedding stories on Social Media sites was Pippa Middleton's "assets"
- Her bottom now has its own [twitter account](#), [facebook page](#) and [website](#).
- Pilates classes have become incredibly popular since the Royal Wedding, solely as a result of all the social media



Pippa Middleton has revealed the secret to her perfect figure - Pilates classes. <http://dlvr.it/S9Cy8>

[CutePippaFace](#): 57 minutes ago.

[Reply](#) [View Tweet](#)

# Accuracy of twitter sentiment apps

---

- Mine the social media sentiment apps and you'll find a huge difference of opinions about Pippa Middleton:
- **TweetFeel**: 25% positive, 75% negative
- **Twendz**: no results
- **TipTop**: 42% positive, 11% negative
- **Twitter Sentiment**: 62% positive, 38% negative
- Accuracy is therefore very questionable

# Language analysis is not always easy



“Rubbish hotel in Madrid”



## It's not just about bottoms and dresses ...

---

- Film, theatre, books, fashion etc
  - impacts on the whole industry
  - predictions about changing society, trends etc.
- Monitoring political views
- Feedback/opinions about multimedia productions, e.g. documentaries, broadcasts etc.
- Feedback about events, e.g. conferences
- Scientific and technological monitoring, competitor surveillance etc.
- Monitoring public opinion
- Creating community memories

## Tracking opinions over time

---

- Opinions can be extracted with a time stamp and/or a geo-location
- We can then analyse changes to opinions about the same entity/event over time, and other statistics
- We can also measure the impact of an entity or event on the overall sentiment about an entity or another event, over the course of time
- In politics, crucial to know how political events impact on people's opinions towards a particular party, minister, law etc.



---

# Case study: Rule-based Opinion Mining from Political Tweets

# Processing political tweets

- GATE-based application to associate people with their political leanings, based on UK 2010 pre-election tweets
- First stage is to find triples <Person, Opinion, Political Party>  
e.g. Bob Smith is pro\_Labour
- Usually, we will only get a single sentiment per tweet
- Where we get conflicting sentiments per tweet, we do not attempt to produce a result
- Later, we can collect all mentions of “Bob Smith” that refer to the same person, and collate the information  
e.g. Bob may be equally in favour of several different parties, not just Labour, but hates the Conservatives above all else



# Creating a corpus

---

- First step is to create a corpus of tweets
- Used the Twitter Streaming API to collect all the tweets over the pre-election period according to various criteria (use of certain hash tags, mention of various political parties etc.)
- Collected tweets in json format and then converted these to xml using JSON-Lib library
- This gives us lots of additional twitter metadata, such as the date and time of the tweet, the number of followers of the person tweeting, the location and other information about the person tweeting, and so on
- This information is useful for disambiguation and for collating the information later

# Tweets with metadata

```
Thu Mar 25 20:06:32 +0000 2010 false 11050953883 <a
href="http://www.trinketsoftware.com/Twikini" rel="nofollow">Twikini</a>Had
pleasure of formally proposing Stuart King as Labour Candidate for Putney.
Yes he can..... false false Fri Jan 23 15:21:58 +0000 2009 Member of
Parliament for Tooting 0 4224 1590 false 19397942 en London, UK Sadiq
Khan MP f0feff
http://a3.twimg.com/profile_background_images/4356861/twitter.jpg false
http://a1.twimg.com/profile_images/427349972/playgroundcropped_normal.JPG
G 0084B4 BDDCAD DDFCC 333333 false SadiqKhan 1390 London
http://www.sadiqkhan.org.uk 0 false
```

Original markups set

- in\_reply\_to\_user\_id
- lang
- location
- name
- notifications
- o
- place
- profile\_background\_color
- profile\_background\_image\_url
- profile\_background\_tile
- profile\_image\_url
- profile\_link\_color
- profile\_sidebar\_border\_color
- profile\_sidebar\_fill\_color
- profile\_text\_color
- protected
- screen\_name
- source
- statuses\_count
- text
- time\_zone

# Metadata

```

Thu Mar 25 20:06:32 +0000 2010 false 11050953883 <a
href="http://www.trinketsoftware.com/Twikini" rel="nofollow">Twikini</a>Had
pleasure of formally proposing Stuart King as Labour Candidate for Putney.
Yes he can..... false false Fri Jan 23 15:21:58 +0000 2009 Member of
Parliament for Tooting 0 4224 1590 false 19397942 en London, UK Sadiq
Khan MP f0feff
http://a3.twimg.com/profile_background_images/4356861/twitter.jpg false
http://a1.twimg.com/profile_images/427349972/playgroundcropped_normal.JP
G 0084B4 BDDCAD DDFFCC 333333 false SadiqKhan 1390 London
http://www.sadiqkhan.org.uk 0 false

```

Date

Tweet

Number of followers

Location

Profile info

Name



## Corpus Size

---

- Raw corpus contained around 5 million tweets
- Many were duplicates due to the way in which the tweets were collected
- Added a de-duplication step during the conversion of json to xml
- This reduced corpus size by 20% to around 4 million
- This still retains the retweets, however (as we may want to do some analysis on these)



# GATE application

---

- Linguistic pre-processing using standard ANNIE components (tokenisation, POS tagging etc)
- No point in attempting parsing
- Apply ANNIE for standard named entities
- Additional targeted gazetteer lookup and JAPE-based (manually developed) grammars
- Grammars first find other entities (political parties etc), and actions such as voting, supporting etc, negatives, questions etc.
- More JAPE grammars combine the previous annotations to form an opinion
- Many of the grammar rules are quite generic so they can be reused in other domains

# Gazetteers

- We create an instance of a flexible gazetteer to match certain useful keywords, in various morphological forms:
  - political parties, e.g. “Conservative”, “LibDem”
  - concepts about winning election, e.g. “win”, “landslide”
  - words for politicians, e.g. “candidate”, “MP”
  - words for voting and supporting a party/ person, e.g. “vote”
  - words indicating negation, e.g. “not”, “never”
- We create another gazetteer containing affect/emotion words from WordNet-Affect, e.g. “beneficial”, “awful”.
  - these have a feature denoting part of speech (category)

# Grammar rules: creating temporary annotations

- Identify questions or doubtful statements as opposed to "factual" statements in tweets: we only care about factual statements
- Create Affect annotations if an "affect" Lookup in the gazetteer is found and if the category matches the POS tag on the Token (this ensures disambiguation of the different possible categories)
  - *"People like her should be shot."* vs *"People like her."*
- We only want to match "affect" adjectives if they're actually being as adjectives to modify some relevant content word



# Example of a grammar rule

Phase: Affect

Input: AffectLookup Token

Options: control = appelt

Check category of both Lookup and Token are adjectives or past participles

Rule: AffectAdjective

(

{AffectLookup.category == adjective, Token.category == VBN}

{AffectLookup.category == adjective, Token.category == JJ}

):tag

-->

```
:tag.Affect = {kind = :tag.AffectLookup.kind,
               category = :tag.AffectLookup.category,
               rule = "AffectAdjective"}
```

copy category and kind values from Lookup to new Affect annotation

# Grammar rules: finding triples

- We first create annotations for Person, Organization, Vote, Party, Negatives etc. based on gazetteer lookup, NEs etc.
- We then create a set of rules (in JAPE) to combine these into pairs or triples:
  - *<Person, Vote, Party> “Tory Phip admits he voted LibDem”.*
  - *<Party, Affect> “When they get a Tory government they'll be sorry.”*
- We create an annotation “Sentiment” which has the following features:
  - **kind**, e.g. “pro\_Labour”, “anti\_LibDem”, etc.
  - **opinion\_holder**, e.g. “John Smith”, “author” etc.

# Identifying the Opinion Holder

- If the opinion holder in the pattern matched is a Person or Organization, we just get the string as the value of `opinion_holder`
- If the opinion holder in the pattern matched is a pronoun, we first find the value of the string of the antecedent and use this as the value of `opinion_holder`, using the pronominal coreference PR and some special JAPE grammars to match the string with the respective proper noun
- If no explicit opinion holder then we use "author" as the value of `opinion_holder`.
- Later we can combine the actual details of the twitterer (from the metadata) instead of just using "author".



# Creating the Application

---

- To process only the actual text of the tweet, we use a special resource in GATE which allows the running of an application over a selected annotation type (in this case, “text” from the Original Markup)
- We still have available all the other metadata if we want to process that too
- We can therefore combine the analysis of the text with analysis of other metadata, within the same application

# Evaluation

- Evaluated Precision on 1000 tweets from corpus
- Manually annotated 150 tweets not identified by the system as opinionated, of which 85% were correctly identified as non-opinionated by the system
- We predict recall on a larger scale from these figures
- Finding a political sentiment correctly (regardless of orientation):  
78% Precision, 47% (predicted) Recall
- For documents known to have a political sentiment, correct opinion polarity 79% Precision
- Overall, 62% Precision, 37% (predicted) Recall
- System has been developed primarily with Precision in mind

## Further work

- Lots of potential for further improvement
- Better processing of hashtags, e.g. #torytombstone, #votefodderforthetories
- Using metadata for training (e.g. political affiliation in profile)
- Better detection of opinionated vs non-opinionated tweets via a separate pre-processing step (primary cause of over/under-generation)
- Improving detection of negation (primary cause of lack of Precision)
- Much world knowledge needed, but even for a human, the task is hard due to irony and missing contextual information: “*Vote Labour. Harry Potter would.*”
- Pre-processing step to include separation of irrelevant material:  
“*I am sooo bored I want to go into labour just for something to do.*”

## More information

---

- Work done in the context of the EU-funded ARCOMEM project
- Dealing with lots of issues about opinion mining from social media
  - with case studies about “Rock am Ring” (a big annual German rock festival) and Greek and Austrian parliaments
- See <http://www.arcomem.eu> for more details