

Part 3: Sentiment Analysis



© The University of Sheffield, 1995-2015

This work is licenced under the [Creative Commons Attribution-NonCommercial-ShareAlike Licence](#)

What is Sentiment Analysis?

- Also known as Opinion Mining
- 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
- Web 2.0 nowadays provides a great medium for people to share things.
- This provides a great source of unstructured information (especially opinions) that may be useful to others (e.g. companies and their rivals, other consumers...)

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



Opinion Mining is Big Business

- Someone who wants to buy a camera
 - Looks for comments and reviews
- Someone who just bought a camera
 - Comments on it
 - Writes about their experience
- Camera Manufacturer
 - Gets feedback from customer
 - Improve their products
 - Adjust Marketing Strategies



Cafe Pie



15 Reviews

#24 of 1,084 Restaurants in Sheffield

#26 of 1,169 Places to Eat in Sheffield

£


Café, British, English

Overview

Reviews (15)

Q&A


Location



TripAdvisor Reviewer Highlights

[Read all 15 reviews](#)

Visitor rating



“Mrs ”

Amazing pies and great friendly staff my kids loved the kids snack packs too excellent idea! I'll definitely be back soon to try a full English breaky!!



Reviewed 6 May 2015
Spring1986



Alex P
1 review

“Proper homemade tasty English grub”

★★★★★ Reviewed 3 May 2015  via mobile

Finally went to the pie cafe for lunch yesterday. Had pie mash and red cabbabe, really tasty, pastry was perfect, plenty of filling and red cabbage had loads of flavour. Loved the diner style decor that still stuck to its Yorkshire routes. I'll definately be back for more. Yum yum!

Was this review helpful? [Yes](#)



RyuOnline

Reviewer

★ 4 reviews

 4 restaurant reviews

“The Famous Cafe Pie in Walkley ”

★★★★★ Reviewed 30 April 2015

Superb cafe with great selection of fresh sandwiches & meals. Sunday roasts with pie or meat are to die for. This is not your average greasy vera sandwich shop. They take great pride in all the food they produce from bacon bap to full themed dishes. & with the name of the shop called "Cafe Pie" it's a must to...

[More](#) ▼

It's not just about product reviews

- Much opinion mining research has been focused around reviews of films, books, electronics etc.
- But there are many other uses
 - companies want to know what people think
 - finding out political and social opinions and moods
 - investigating how public mood influences the stock market
 - investigating and preserving community memories
 - drawing inferences from social analytics

But there are lots of tools that “analyse” social media already....

- Streamcrab <http://www.streamcrab.com/>
- Semantria <http://semantria.com>
- Social Mention <http://socialmention.com/>
- Sentiment140: <http://www.sentiment140.com/>
- TipTop: <http://feeltiptop.com/>

Why not use existing sentiment apps?

- 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 Lady Gaga's dress, you can often only search on “Lady Gaga” to get hits
- They're suitable for a quick sanity check of social media, but not really for business needs
- And the opinion finding they do isn't always very good...

Why are these sites unsuccessful?

- They don't work well at more than a very basic level
- They mainly use dictionary lookup for positive and negative words
- Or they use ML, which only works for text that's similar in style
- They classify the tweets as positive or negative, but not with respect to the keyword you're searching for
- First, the keyword search just retrieves any tweet mentioning it, but not necessarily about it as a topic
- Second, there is no correlation between the keyword and the sentiment: the sentiment refers to the tweet as a whole
- Sometimes this is fine, but it can also go horribly wrong

“Positive” tweets about fracking

- *Help me stop fracking. Sign the petition to David Cameron for a #frack-free UK now!*
- *I'll take it as a sign that the gods applaud my new anti-fracking country love song.*
- *#Cameron wants to change the law to allow #fracking under homes without permission. Tell him NO!!!!*

Be careful!



Sentiment analysis isn't just about looking at the sentiment words

- “It's a great movie if you have the taste and sensibilities of a 5-year-old boy.”
- “It's terrible Candidate X did so well in the debate last night.”
- “I'd have liked the film a lot more if it had been a bit shorter.”

Situation is everything. If you and I are best friends, then my graceful swearing at you is different from if it's at my boss.

Death confuses opinion mining tools

- Opinion mining tools are good for a general overview, but not for some situations



Nobody liked Leonard Nimoy



Or did they?



Peter Cook @Slublog · May 22

Oh, hell. Watching the 2009 "Star Trek" just made me remember that **Leonard Nimoy** is **dead**. And now I'm sad.



aahelejandro @Dank_Herbert · May 12

Has **Leonard Nimoy** been **dead** long enough that it's okay to say he wrote the worst poetry on the planet



Romantic Lost Prince @danaFringe · Mar 8

The great actor, **Leonard Nimoy** is **dead**, February 27 2015. I'm so very sad and I very much miss him! pinterest.com/pin/4321345267...




Usum @usumtwit · May 9

Leonard Nimoy Dead: 9 Great Moments of His Career Besides Spock
ift.tt/1P8UyqH #News, leonard, leonard nimoy, nimoy, spock



Opinion Mining for Stock Market Prediction

- It might be only fiction, but using opinion mining for stock market prediction has been already a reality for some years
- Research shows that opinion mining outperforms event-based classification for trend prediction [Bollen2011]
- Many investment companies offer products based on (shallow) opinion mining



Derwent Capital Markets

- Derwent Capital Markets launched a £25m fund that makes its investments by evaluating whether people are generally happy, sad, anxious or tired, because they believe it will predict whether the market will move up or down.
- Bollen told the Sunday Times: "We recorded the sentiment of the online community, but we couldn't prove if it was correct. So we looked at the Dow Jones to see if there was a correlation. We believed that if the markets fell, then the mood of people on Twitter would fall."
- "But we realised it was the other way round — that a drop in the mood or sentiment of the online community would precede a fall in the market."

Derwent Capital Markets

CAYMAN



" Using global sentiment analysis to trade the financial markets "

Social Media Sentiment Trading - Private Managed Accounts



Due to massive demand we have decided to apply our social media sentiment analysis technology to managed trading accounts allowing private investors the opportunity to invest upwards of £10,000 GBP.

If you are a sophisticated investor or high net worth person and not a US resident and would like to open an account then please click the link below...

But don't believe all you read...





- It's not really possible to predict the stock market in this way
- Otherwise we'd be millionaires by now!
- In Bollen's case. the advertised results were biased by selection (they picked the winners after the race and tried to show correlation)
- The accuracy claim is too general to be useful (you can't predict individual stock prices, only the general trend)
- There's no real agreement about what's useful and what isn't
- <http://sellthenews.tumblr.com/post/21067996377/noitdoesnot>

Who Wants to be a Millionaire?

Child Care   0:24


Also known as the “cry it out” method, the Ferber method is a technique for teaching children to do what?

| | | | | | |
|----------|----------------------|-----|----------|-------------------------|-----|
| A | Eat their vegetables | 38% | B | Learn new words | 3% |
| C | Walk on their own | 15% | D | Sleep through the night | 44% |

Ask the audience or phone a friend?

Which do you think is better?



What's the capital of Spain?

A: Barcelona

B: Madrid

C: Valencia

D: Seville

What's the height of Mt Kilimanjaro?

A: 19,341 ft

B: 23,341 ft

C: 15,341 ft

D: 21,341 ft

Go for the majority or trust an expert?

- It depends what kind of question you're asking
- In *Who Wants to Be a Millionaire*, people tend to ask the audience fairly early on, because once the questions get hard, they can't rely on the audience getting it right

What's the height of Mt Kilimanjaro?

- A: 19,341 ft
- B: 23,341 ft
- C: 15,341 ft
- D: 21,341 ft

What's the capital of Spain?

- A: Barcelona
- B: Madrid
- C: Valencia
- D: Seville

Why bother with opinion mining?

- It depends what kind of information you want
- Don't use opinion mining tools to help you win money on quiz shows
- Recent research has shown that one knowledgeable analyst is better than gathering general public sentiment from lots of analysts and taking the majority opinion
- But only for some kinds of tasks

Whose opinion should you trust?

- Opinion mining gets difficult when the users are exposed to opinions from more than one analyst
- Intuitively, one would probably trust the opinion supported by the majority.
- But some research shows that the user is better off trusting the most credible analyst.
- Then the question becomes: who is the most credible analyst?
- Notions of trust, authority and influence are all related to opinion mining

All opinions are not equal

- Opinion Mining needs to take into account how much influence any single opinion is worth
- This could depend on a variety of factors, such as how much trust we have in a person's opinion, and even what sort of person they are
- Need to account for:
 - experts vs non-experts
 - Spammers
 - frequent vs infrequent posters
 - “experts” in one area may not be expert in another
 - how frequently do other people agree?

Trust Recommenders

- Two types of trust: relationship (local) trust and reputation (global) trust.
- **Relationship trust:** if you and I both rate the same things, and our opinions on them match closely, we have high relationship trust. This can be extended to a social networking group --> web of trust.
- **Reputation trust:** if you've recommended the same thing as other people, and usually your recommendation is close to what the majority of people think, then you're considered to be more of an expert and have high reputation trust.
- We can extend relationship trust to form clusters of interests and likes/dislikes
- We can narrow reputation trust to opinions about similar topics

Related (sub)topics: general

- **Opinion extraction:** extract the piece of text which represents the opinion
 - I just bought a new camera yesterday. It was a bit expensive, but the battery life is very good.
- **Sentiment classification/orientation:** extract the polarity of the opinion (e.g. positive, negative, neutral, or classify on a numerical scale)
 - negative: expensive
 - positive: good battery life
- **Opinion summarisation:** summarise the overall opinion about something
 - price:negative, battery life: positive --> overall 7/10

Feature-opinion association

- **Feature-opinion association:** given a text with target features and opinions extracted, decide which opinions comment on which features.
 - “The battery life is good but not so keen on the picture quality”
- **Target identification:** which thing is the opinion referring to?
- **Source identification:** who is holding the opinion?
- There may be attachment and co-reference issues
 - “The camera comes with a free case but I don't like the colour much.”
 - Does this refer to the colour of the case or the camera?

Opinion spamming



Suppose we run a contest where people retweet our ad repeatedly,
and the winner's whoever loses the most followers.

Spam opinion detection (fake reviews)

- Sometimes people get paid to post “spam” opinions supporting a product, organisation or even government
- An article in the New York Times discussed one such company who gave big discounts to post a 5-star review about the product on Amazon
- http://www.nytimes.com/2012/01/27/technology/for-2-a-star-a-retailer-gets-5-star-reviews.html?_r=3&ref=business
- Could be either positive or negative opinions
- Generally, negative opinions are more damaging than positive ones

How to detect fake opinions?


- Review content: lexical features, content and style inconsistencies from the same user, or similarities between different users
- Complex relationships between reviews, reviewers and products
- Publicly available information about posters (time posted, posting frequency etc)
- Detecting inconsistencies, contradictions, lack of entailment etc. is also relevant here

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-taker
 - *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

Social networks can trigger new events

- Not only can online social networks provide a snapshot of current or past situations, but they can actually trigger chains of reactions and events
- Ultimately these events might lead to societal, political or administrative changes
- Since the Royal Wedding, Pilates classes became incredibly popular in the UK solely as a result of social media.
- Why?
- Pippa Middleton's bottom is the answer!
- Pictures of her bottom are allegedly worth more than those of her face!



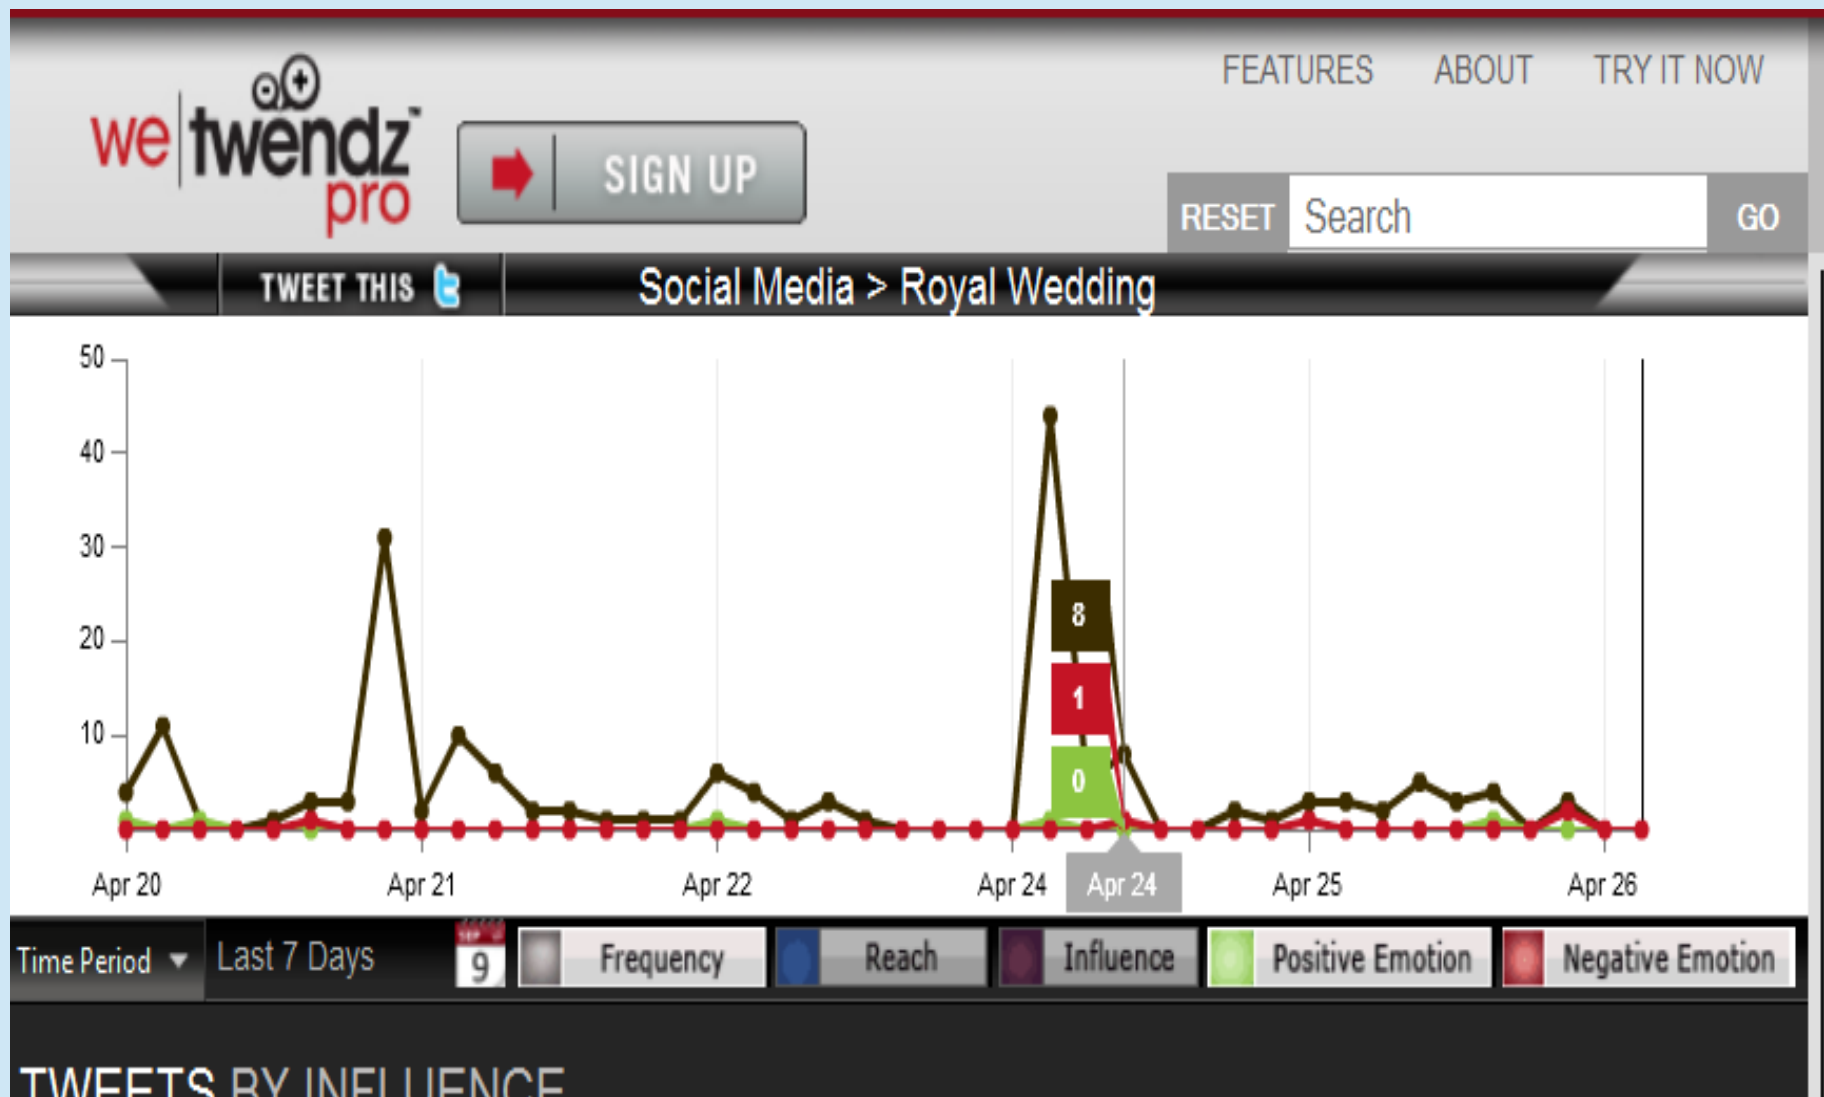
Social media and politics

- Twitter provides real-time feedback on political debates that's much faster than traditional polling.
- Social media chatter can gauge how a candidate's message is being received or even warn of a popularity dive.
- Campaigns that closely monitor the Twittersphere have a better feel of voter sentiment, allowing candidates to fine-tune their message for a particular state: “playing to your audience”.
- But applying complex algorithms to social media is far from perfect for predicting politics, e.g. you can't detect sarcasm reliably.
- Nevertheless, Twitter has played a role in intelligence gathering on uprisings around the world, showing accuracy at gauging political sentiment.
- <http://www.usatoday.com/tech/news/story/2012-03-05/social-super-tuesday-prediction/53374536/1>

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 (e.g. in politics)
- Also possible to incorporate statistical (non-linguistic) techniques to investigate dynamics of opinions, e.g. find statistical correlations between interest in certain topics or entities/events and number/impact/influence of tweets etc.

Viewing opinion changes over time



Some opinion mining resources

- Sentiment lexicons
- Sentiment-annotated corpora

Bing Liu's English Sentiment Lexicon

- Available from
- <http://www.cs.uic.edu/~liub/FBS/opinion-lexicon-English.rar>
- List of strings
- Positive words: 2006
 - a+, abound, abounds, abundance, abundant
- Negative words: 4783
 - Examples: 2-faced, abnormal, abolish, abominable
- Useful properties: includes mis-spellings, morphological variants, slang

The MPQA Subjectivity Lexicon

- GPL License; download from http://mpqa.cs.pitt.edu/lexicons/subj_lexicon/
 - type - either strongsubj or weaksubj
 - len - length of the clue in words
 - word1 - token or stem of the clue
 - pos1 - part of speech of the clue, may be anypos (any part of speech)
 - stemmed1 - y (yes) or n (no)
 - priorpolarity - positive, negative, both, neutral

| | Strength | Length | Word | Part-of-speech | Stemmed | Polarity |
|----|-----------------|--------|-------------------|----------------|------------|------------------------|
| 1. | type=weaksubj | len=1 | word1=abandoned | pos1=adj | stemmed1=n | priorpolarity=negative |
| 2. | type=weaksubj | len=1 | word1=abandonment | pos1=noun | stemmed1=n | priorpolarity=negative |
| 3. | type=weaksubj | len=1 | word1=abandon | pos1=verb | stemmed1=y | priorpolarity=negative |
| 4. | type=strongsubj | len=1 | word1=abase | pos1=verb | stemmed1=y | priorpolarity=negative |
| 5. | type=strongsubj | len=1 | word1=abasement | pos1=anypos | stemmed1=y | priorpolarity=negative |
| 6. | type=strongsubj | len=1 | word1=abash | pos1=verb | stemmed1=y | priorpolarity=negative |

LIWC

- Linguistic Inquiry and Word Counts (LIWC) – a database of categorized regular expressions
- Costs around \$90
- Some sentiment-relevant categories with example regular expressions below, from <http://sentiment.christopherpotts.net/lexicons.html>

| Category | Examples |
|----------|--|
| Negate | aint, ain't, arent, aren't, cannot, cant, can't, couldnt, ... |
| Swear | arse, arsehole*, arses, ass, asses, asshole*, bastard*, ... |
| Social | acquainta*, admit, admits, admitted, admitting, adult, adults, advice, advis* |
| Affect | abandon*, abuse*, abusi*, accept, accepta*, accepted, accepting, accepts, ache* |
| Anx | afraid, alarm*, anguish*, anxi*, apprehens*, asham*, aversi*, avoid*, awkward* |
| Anger | jealous*, jerk, jerked, jerks, kill*, liar*, lied, lies, lous*, ludicrous*, lying, mad |

Try some different lexicons

- <http://sentiment.christopherpotts.net/lexicon/> Get sentiment scores for single words from a variety of sentiment lexicons
- <http://sentiment.christopherpotts.net/textscores/> Show how a variety of lexicons score novel texts
- Spend a bit of time experimenting with these to see the difference

Problems with Using Sentiment Lexicons

- Sentiment words are context-dependent and ambiguous
 - “a long dress” vs “a long walk” vs “a long battery life”
 - “the camera was cheap” vs “the camera looked cheap”
 - “I like her” vs “People like her should be shot”.
- Solutions involve
 - domain-specific lexicons
 - lexicons including context (see e.g. Scharl's GWAP methods <http://apps.facebook.com/sentiment-quiz>)
 - constraining POS categories
- And we also have to look at things in the context other than sentiment words (negatives, adverbs, sarcasm etc.)

Summary

- In this section, we have introduced the concept of sentiment analysis
 - Why is it important?
 - What are the typical problems associated with it?
 - What problems are incurred when working with social media?
 - How can these be resolved?
- Next, we'll look at how to actually do sentiment analysis in GATE