

SENTIMENT ANALYSIS

Mausam

(With slides from Jan Wiebe, Kavita Ganesan, Heng Ji, Dan Jurafsky, Chris Manning)

Motivation

“What people think?”

What others think has always been an important piece of information

“Which car should I buy?”

“Which schools should I apply to?”

“Which Professor to work for?”

“Whom should I vote for?”



“Whoala! I have the reviews I need”

Now that I have “too much” information on one topic...I could easily form my opinion and make decisions...

Is this true?

...Not Quite

Searching for reviews may be difficult

Can you search for opinions as conveniently as general Web search?

eg: is it easy to search for “*iPhone vs Google Phone*”?

“Let me look at reviews on one site only...”

Problems?



Biased views

- all reviewers on one site may have the same opinion

Fake reviews/Spam (sites like YellowPages, CitySearch are prone to this)

- people post good reviews about their own product OR services
- some posts are plain spams

Coincidence or Fake?

Reviews for a moving company from YellowPages

- # of merchants reviewed by the each of these reviewers → 1
- Review dates close to one another
- All rated 5 star
- Reviewers seem to know exact names of people working in the company and TOO many positive mentions

THE BEST!!!! 11/30/2007 Posted by [c_karen](#) ★★★★★

NorthStar did an **outstanding job** of packing and moving my things. Quite frankly I was expecting some things to be broken. However, to my surprise not one thing was broken and everything **went as smooth** as could be expected. I had approximately 15,000 lbs. of items to move. I am **very impressed** with NorthStar and I would not hesitate to utilize them again for my next move. All of the young men who assisted in packing and loading were **very hard working and polite**.

Pros: everything was great


GOOD MOVING 10/11/2007 Posted by [loanlee777](#) ★★★★★

About a month ago, on Sep 12, we hired NorthStar Moving to move our belongings from our house in Van Nuys to the Highway Storage place in Santa Clara. We would like to express our sincere **thanks and appreciation for the professional work** that was carried out by NorthStar team of workers. **In particular, we would like to mention the four NorthStar workers: Roy Ashual, Moshiko Haziza, Guillermo Molise and Roberto Mendoza for their very dedicated service.** Besides being **good natured and helpful**, they worked **very well and took good care** of our personal effects. We would definitely refer them and NorthStar Moving to any of our friends who are looking for a **good** moving company.

Great movers 10/08/2007 Posted by [shelly_morgan](#) ★★★★★

I wanted to thank the Northstar Moving group for a fabulous job. We hired Northstar Moving on August 4th to move us out of two storage units and where we were staying to our new home in Los Angeles. I had gone through surgery on the 2nd and was in no condition to move around a lot. The Northstar Moving team was great. I slept in while my husband met them at the first pick-up point. Then they came to the 2nd and that is where I met them. When we arrived at the new house they found something for me to sit on and I **set in one place in the garage telling them which room the items went.** They were **great**. They had **wonderful personalities**; I have never had **so much fun moving** (even if I was in some pain). Northstar thank you again for the **great team** and customer service.

Problem Names



Synonymous
&
Interchangeably Used!

So, what is Subjectivity?


- The **linguistic** expression of somebody's **opinions, sentiments, emotions.....**(private states)
- private state: state that is not open to objective verification (*Quirk, Greenbaum, Leech, Svartvik (1985). A Comprehensive Grammar of the English Language.*)
- **Subjectivity analysis** - is the computational study of **affect, opinions,** and **sentiments** expressed in text
 - blogs
 - editorials
 - reviews (of products, movies, books, etc.)
 - newspaper articles

Subjectivity Analysis on iPhone Reviews

Business' Perspective

- **Apple:** What do consumers think about iPhone?
 - Do they like it?
 - What do they dislike?
 - What are the major complaints?
 - What features should we add?
- **Apple's competitor:**
 - What are iPhone's weaknesses?
 - How can we compete with them?
 - Do people like

Known as Business Intelligence




sy. Maybe I should have
arket research first.”

Bing Shopping

HP Officejet 6500A E710N Multifunction Printer

[Product summary](#) [Find best price](#) **Customer reviews** [Specifications](#) [Related items](#)



\$121.53 - \$242.39 (14 stores)

Compare

Average rating ★★★★★ (144)



Most mentioned




Show reviews by source



Twitter sentiment versus Gallup Poll of Consumer Confidence


Brendan O'Connor, Ramnath Balasubramanyan, Bryan R. Routledge, and Noah A. Smith. 2010. From Tweets to Polls: Linking Text Sentiment to Public Opinion Time Series. In ICWSM-2010



Twitter sentiment:


Johan Bollen, Huina Mao, Xiaojun Zeng.
2011. [Twitter mood predicts the stock market](#),

Journal of Computational Science 2:1, 1-8. 10.1016/j.jocs.2010.12.007.



Bollen et al. (2011)

- CALM today predicts DJIA 3 days later
- At least one current hedge fund uses this algorithm



Definition

Sentiment Analysis

- Sentiment analysis is the detection of **attitudes**
 - “enduring, affectively colored beliefs, dispositions towards objects or persons”
 - 1. **Holder (source)** of attitude
 - 2. **Target (aspect)** of attitude
 - 3. **Type** of attitude
 - From a set of types
 - *Like, love, hate, value, desire, etc.*
 - Or (more commonly) simple weighted **polarity**:
 - *positive, negative, neutral, together with strength*
 - 4. **Text** containing the attitude
 - Sentence or entire document

Sentiment Analysis

- Simplest task:
 - Is the attitude of this text positive or negative?
- More complex:
 - Rank the attitude of this text from 1 to 5
- Advanced:
 - Detect the target, source, or complex attitude types

Sentiment Analysis

- Simplest task:
 - Is the attitude of this text positive or negative?
- More complex:
 - Rank the attitude of this text from 1 to 5
- Advanced:
 - Detect the target, source, or complex attitude types

Baseline Algorithms

Sentiment Classification in Movie Reviews

Bo Pang, Lillian Lee, and Shivakumar Vaithyanathan. 2002. Thumbs up? Sentiment Classification using Machine Learning Techniques. EMNLP-2002, 79–86.

Bo Pang and Lillian Lee. 2004. A Sentimental Education: Sentiment Analysis Using Subjectivity Summarization Based on Minimum Cuts. ACL, 271-278

- Polarity detection:
 - Is an IMDB movie review positive or negative?
- Data: *Polarity Data 2.0*:
 - <http://www.cs.cornell.edu/people/pabo/movie-review-data>

IMDB data in the Pang and Lee database



when `_star wars_` came out some twenty years ago , the image of traveling throughout the stars has become a commonplace image . [...]

when han solo goes light speed , the stars change to bright lines , going towards the viewer in lines that converge at an invisible point .

cool .

`_october sky_` offers a much simpler image— that of a single white dot , traveling horizontally across the night sky . [. . .]



“ snake eyes ” is the most aggravating kind of movie : the kind that shows so much potential then becomes unbelievably disappointing . it’s not just because this is a brian depalma film , and since he’s a great director and one who’s films are always greeted with at least some fanfare .

and it’s not even because this was a film starring nicolas cage and since he gives a brauvara performance , this film is hardly worth his talents .

Baseline Algorithm (adapted from Pang and Lee)

- Tokenization
- Feature Extraction
- Classification using different classifiers
 - Naïve Bayes
 - MaxEnt
 - SVM

Sentiment Tokenization Issues

- Deal with HTML and XML markup
- Twitter mark-up (names, hash tags)
- Capitalization (preserve for words in all caps)
- Phone numbers, dates
- Emoticons

Extracting Features for Sentiment Classification

- How to handle negation
 - I **didn't** like this movie
 - vs
 - I really like this movie
- Which words to use?
 - Only adjectives
 - All words
 - All words turns out to work better, at least on this data

Negation

Das, Sanjiv and Mike Chen. 2001. Yahoo! for Amazon: Extracting market sentiment from stock message boards. In Proceedings of the Asia Pacific Finance Association Annual Conference (APFA).

Bo Pang, Lillian Lee, and Shivakumar Vaithyanathan. 2002. Thumbs up? Sentiment Classification using Machine Learning Techniques. EMNLP-2002, 79—86.

Add NOT_ to every word between negation and following punctuation:

didn't like this movie , but I



didn't NOT_like NOT_this NOT_movie but I

Accounting for Negation

- Let us consider the following positive sentence:
 - Example: *Luckily, the **smelly poo** did not leave **awfully** **nasty** **stains** on my **favorite** shoes!*
- Rest of Sentence (RoS):
 - Following: *Luckily, the **smelly poo** did not leave awfully nasty stains on my favorite shoes!*
 - Around: *Luckily, the smelly poo did not leave awfully nasty stains on my favorite shoes!*
- First Sentiment-Carrying Word (FSW):
 - Following: *Luckily, the **smelly poo** did not leave awfully **nasty** **stains** on my **favorite** shoes!*
 - Around: *Luckily, the **smelly** poo did not leave awfully **nasty** **stains** on my **favorite** shoes!*

Accounting for Negation

- Let us consider the following positive sentence:
 - Example: *Luckily, the **smelly poo** did not leave **awfully** **nasty** **stains** on my **favorite** shoes!*
- Next Non-Adverb (NNA):
 - Following: *Luckily, the **smelly poo** did not leave **awfully** nasty **stains** on my **favorite** shoes!*
- Fixed Window Length (FWL):
 - Following (3): *Luckily, the **smelly poo** did not leave awfully nasty **stains** on my **favorite** shoes!*
 - Around (3): *Luckily, the smelly poo did not leave awfully **stains** nasty on my **favorite** shoes!*

KEYWORDS SELECTION FROM TEXT

- Pang et. al. (2002)
 - Binary Classification of unigrams
 - Positive
 - Negative
 - Unigram method reached 80% accuracy.

N-GRAM BASED CLASSIFICATION

- Learn N-Grams (frequencies) from pre-annotated training data.
- Use this model to classify new incoming sample.

PART-OF-SPEECH BASED PATTERNS

- Extract POS patterns from training data.
- Usually used for subjective vs objective classification.
- Adjectives and Adverbs contain sentiments
- Example patterns
 - *-JJ-NN : trigram pattern
 - JJ-NNP : bigram pattern
 - *-JJ : bigram pattern
- Used as features in NB/Logistic Regression

Problems:

What makes reviews hard to classify?

- Subtlety:
 - Perfume review in *Perfumes: the Guide*:
 - “If you are reading this because it is your darling fragrance, please wear it at home exclusively, and tape the windows shut.”
 - Dorothy Parker on Katherine Hepburn
 - “She runs the gamut of emotions from A to B”

CHALLENGES

- Ambiguous words
 - This music cd is literal waste of time.
(negative)
 - Please throw your waste material here.
(neutral)
- Sarcasm detection and handling
 - “All the features you want - too bad they don’t work. :-P”
- (Almost) No resources and tools for low/scarce resource languages like Indian languages.

User written: grammar, spellings...

Hi,




I have Haier phone.. It was good when i was buing this phone.. But I invented A lot of bad features by this phone those are It's cost is low but Software is not good and Battery is very bad.., Ther are no signals at out side of the city...,, People can't understand this type of software...,, There aren't features in this phone, Design is better not good...,, Sound also bad..So I'm not intrest this side They are giving heare phones it is good. They are these are also good.They are giving also good because other phones low wait.




**Lack of punctuation marks,
Grammatical errors**

Wait.. err.. Come again

From: www.mouthshut.com

Alternating Sentiment

I suggest that instead of fillings songs in tunes you should
 tunes (not made of songs) only. The ph has good
popularity in old age people. Third  d tried much for its
data cable but i find it nowhere. It should be supplied with
set with some extra cost.

Go features of this phone are its cheapest price and
durability . It should  ve some features more than nokia
1200. it is easily available in market an pair is also
available

Subject Centrality

- I have this personal experience of using this cell phone. I bought it one and half years back. It had modern features that a normal cell phone has, and the look is excellent. I was very impressed by the design. I bought it for Rs. 8000. It was a gift for someone. It worked fine for first one month, and then started the series of multiple faults it has. First the speaker didnt work, I took it to the service centre (which is like a govt. office with no work). It took 15 days to repair the handset, moreover they ~~charged me Rs. 500. Then after 15~~ days again the mike didnt work, then again same set of time was consumed for the repairs and it continued. Later the camera didnt work, the speakes were rubbish, it used to hang. It started restarting automatically. And the govt. office had staff which I doubt have any knoledge of cell phones??

These multiple faults continued for as long as one year, when the warranty period ended. In this period of time ~~I spent a considerable amount on the petrol, a lot of time (as~~ the service centre is a govt. office). And at last the phone is still working, but now it works as a paper weight. The company who produces such items must be sacked. I understand that it might be fault with one prticular handset, but the company itself never bothered for replacement and I have never seen such miserable cust service. For a comman man like me, Rs. 8000 is a big amount. And I spent almost the same amount to get it work, if any ~~has a good suggestion and can gude me how to sue such~~ companies, please guide.

For this the quality team is faulty, the cust service is really miserable and the worst condition of any organisation I have ever seen is ~~is with the service centre for Fly and Sony Erricson, (it's near~~ Sancheti hospital, Pune). I dont have any thing else to say.

Thwarted Expectations and Ordering Effects

- “This film should be brilliant. It sounds like a great plot, the actors are first grade, and the supporting cast is good as well, and Stallone is attempting to deliver a good performance. However, it can’t hold up.”
- Well as usual Keanu Reeves is nothing special, but surprisingly, the very talented Laurence Fishbourne is not so good either, I was surprised.

Thwarted Expectations and Ordering Effects

- “This film should be **brilliant**. It sounds like a **great** plot, the actors are **first grade**, and the supporting cast is **good** as well, and Stallone is attempting to deliver a good performance. However, it **can't hold up**.”
- Well as usual Keanu Reeves is nothing special, but surprisingly, the **very talented** Laurence Fishbourne is **not so good** either, I was surprised.

Sentiment Lexicons

The General Inquirer

Philip J. Stone, Dexter C Dunphy, Marshall S. Smith, Daniel M. Ogilvie. 1966. The General Inquirer: A Computer Approach to Content Analysis. MIT Press

- Home page: <http://www.wjh.harvard.edu/~inquirer>
- List of Categories:
<http://www.wjh.harvard.edu/~inquirer/homecat.htm>
- Spreadsheet: <http://www.wjh.harvard.edu/~inquirer/inquirerbasic.xls>
- Categories:
 - Positive (1915 words) and Negative (2291 words)
 - Strong vs Weak, Active vs Passive, Overstated versus Understated
 - Pleasure, Pain, Virtue, Vice, Motivation, Cognitive Orientation, etc
- Free for Research Use

LIWC (Linguistic Inquiry and Word Count)

Pennebaker, J.W., Booth, R.J., & Francis, M.E. (2007). Linguistic Inquiry and Word Count: LIWC 2007. Austin, TX

- Home page: <http://www.liwc.net/>
- 2300 words, >70 classes
- **Affective Processes**
 - negative emotion (*bad, weird, hate, problem, tough*)
 - positive emotion (*love, nice, sweet*)
- **Cognitive Processes**
 - Tentative (*maybe, perhaps, guess*), Inhibition (*block, constraint*)
- **Pronouns, Negation** (*no, never*), **Quantifiers** (*few, many*)
- \$30 or \$90 fee

MPQA Subjectivity Cues Lexicon

Theresa Wilson, Janyce Wiebe, and Paul Hoffmann (2005). Recognizing Contextual Polarity in Phrase-Level Sentiment Analysis. Proc. of HLT-EMNLP-2005.

Riloff and Wiebe (2003). Learning extraction patterns for subjective expressions. EMNLP-2003.

- Home page: http://www.cs.pitt.edu/mpqa/subj_lexicon.html
- 6885 words
 - 2718 positive
 - 4912 negative
- Each word annotated for intensity (strong, weak)
- GNU GPL

Bing Liu Opinion Lexicon

Minqing Hu and Bing Liu. Mining and Summarizing Customer Reviews. ACM SIGKDD-2004.

- [Bing Liu's Page on Opinion Mining](#)
- <http://www.cs.uic.edu/~liub/FBS/opinion-lexicon-English.rar>
- 6786 words
 - 2006 positive
 - 4783 negative

SentiWordNet

Stefano Baccianella, Andrea Esuli, and Fabrizio Sebastiani. 2010
SENTIWORDNET 3.0: An Enhanced Lexical Resource for Sentiment Analysis
and Opinion Mining. LREC-2010

- Home page: <http://sentiwordnet.isti.cnr.it/>
- All WordNet synsets automatically annotated for degrees of positivity, negativity, and neutrality/objectiveness
- [estimable(J,3)] “may be computed or estimated”
Pos 0 Neg 0 Obj 1
- [estimable(J,1)] “deserving of respect or high regard”
Pos .75 Neg 0 Obj .25

ADVANTAGES AND DISADVANTAGES

- Advantages
 - Fast
 - No Training data necessary
 - Good initial accuracy
- Disadvantages
 - Does not deal with multiple word senses
 - Does not work for multiple word phrases
 - May not deal with domain-specific constructions

Can we use IMDB corpus to learn general sentiment words?

Analyzing the polarity of each word in IMDB

Potts, Christopher. 2011. On the negativity of negation. SALT 20, 636-659.

- How likely is each word to appear in each sentiment class?
- Count("bad") in 1-star, 2-star, 3-star, etc.
- But can't use raw counts:


- Instead, **likelihood**:

$$P(w | c) = \frac{f(w, c)}{\sum_{w \in \mathcal{V}} f(w, c)}$$

- Make them comparable between words


- **Scaled likelihood**:

$$\frac{P(w | c)}{P(w)}$$



Analyzing the polarity of each word in IMDB

Potts, Christopher. 2011. On the negativity of negation. SALT 20, 636-659.




Other sentiment feature: Logical negation

Potts, Christopher. 2011. On the negativity of negation. SALT 20, 636-659.


- Is logical negation (*no*, *not*) associated with negative sentiment?
- Potts experiment:
 - Count negation (*not*, *n't*, *no*, *never*) in online reviews
 - Regress against the review rating

Potts 2011 Results: More negation in negative sentiment

IMDB (4,073,228 tokens)



Five-star reviews (846,444 tokens)



Corpus-based Methods

Semi-supervised learning of lexicons

- Use a small amount of information
 - A few labeled examples
 - A few hand-built patterns
- To **bootstrap** a lexicon

Hatzivassiloglou and McKeown intuition for identifying word polarity

Vasileios Hatzivassiloglou and Kathleen R. McKeown. 1997. Predicting the Semantic Orientation of Adjectives. ACL, 174–181

- Adjectives conjoined by “*and*” have same polarity
 - Fair **and** legitimate, corrupt **and** brutal
 - *fair **and** brutal, *corrupt **and** legitimate
- Adjectives conjoined by “*but*” do not
 - fair **but** brutal

Hatzivassiloglou & McKeown 1997

Step 1

- **Label seed set** of 1336 adjectives (all >20 in 21 million word WSJ corpus)
 - 657 positive
 - adequate central clever famous intelligent remarkable reputed sensitive slender thriving...
 - 679 negative
 - contagious drunken ignorant lanky listless primitive strident troublesome unresolved unsuspecting...

Hatzivassiloglou & McKeown 1997

Step 2

- Expand seed set to conjoined adjectives



"was nice and"

[Nice location in Porto and the front desk staff was nice and helpful...](http://www.tripadvisor.com/ShowUserReviews-g189180-d206904-r12068...)

www.tripadvisor.com/ShowUserReviews-g189180-d206904-r12068...

Mercure Porto Centro: Nice location in Porto and the front desk staff **was nice and helpful** - See traveler reviews, 77 candid photos, and great deals for Porto, ...

nice, helpful

[If a girl was nice and classy, but had some vibrant purple dye in ...](http://answers.yahoo.com)

answers.yahoo.com › Home › All Categories › Beauty & Style › Hair

4 answers - Sep 21


Question: Your personal opinion or what you think other people's opinions might ...

Top answer: I think she would be cool and confident like katy perry :)

nice, classy


Hatzivassiloglou & McKeown 1997 Step 3

3. A supervised learning algorithm builds a **graph** of adjectives linked by the same or different semantic orientation



Hatzivassiloglou & McKeown 1997 Step 4

4. A clustering algorithm partitions the adjectives into two subsets



Output polarity lexicon

- Positive
 - bold decisive disturbing generous good honest important large mature patient peaceful positive proud sound stimulating straightforward strange talented vigorous witty...
- Negative
 - ambiguous cautious cynical evasive harmful hypocritical inefficient insecure irrational irresponsible minor outspoken pleasant reckless risky selfish tedious unsupported vulnerable wasteful...

Output polarity lexicon

- Positive
 - bold decisive **disturbing** generous good honest important large mature patient peaceful positive proud sound stimulating straightforward **strange** talented vigorous witty...
- Negative
 - ambiguous **cautious** cynical evasive harmful hypocritical inefficient insecure irrational irresponsible minor **outspoken pleasant** reckless risky selfish tedious unsupported vulnerable wasteful...

Turney Algorithm

Turney (2002): Thumbs Up or Thumbs Down? Semantic Orientation Applied to Unsupervised Classification of Reviews

1. Extract a *phrasal lexicon* from reviews
2. Learn polarity of each phrase
3. Rate a review by the average polarity of its phrases

Extract two-word phrases with adjectives

First Word	Second Word	Third Word (not extracted)
JJ	NN or NNS	anything
RB, RBR, RBS	JJ	Not NN nor NNS
JJ	JJ	Not NN or NNS
NN or NNS	JJ	Nor NN nor NNS
RB, RBR, or RBS	VB, VBD, VBN, VBG	anything

How to measure polarity of a phrase?

- Positive phrases co-occur more with “*excellent*”
 - Negative phrases co-occur more with “*poor*”
 - But how to measure co-occurrence?
-
- General principle – distributional similarity

Pointwise Mutual Information

- **Mutual information** between 2 random variables X and Y

$$I(X, Y) = \sum_x \sum_y P(x, y) \log_2 \frac{P(x, y)}{P(x)P(y)}$$

- **Pointwise mutual information:**

- How much more do events x and y co-occur than if they were independent?

$$\text{PMI}(x, y) = \log_2 \frac{P(x, y)}{P(x)P(y)}$$

Pointwise Mutual Information

- **Pointwise mutual information:**

- How much more do events x and y co-occur than if they were independent?

$$\text{PMI}(x, y) = \log_2 \frac{P(x, y)}{P(x)P(y)}$$

- **PMI between two words:**

- How much more do two words co-occur than if they were independent?

$$\text{PMI}(\textit{word}_1, \textit{word}_2) = \log_2 \frac{P(\textit{word}_1, \textit{word}_2)}{P(\textit{word}_1)P(\textit{word}_2)}$$

How to Estimate Pointwise Mutual Information

- Query search engine
 - $P(\text{word})$ estimated by `hits(word) / N`
 - $P(\text{word}_1, \text{word}_2)$ by `hits(word1 NEAR word2) / N`
 -

$$\text{PMI}(\text{word}_1, \text{word}_2) = \log_2 \frac{\frac{1}{N} \text{hits}(\text{word}_1 \text{ NEAR } \text{word}_2)}{\frac{1}{N} \text{hits}(\text{word}_1) \frac{1}{N} \text{hits}(\text{word}_2)}$$

Does phrase appear more with “poor” or “excellent”?

$$\begin{aligned}
 \text{Polarity}(\textit{phrase}) &= \text{PMI}(\textit{phrase}, \text{"excellent"}) - \text{PMI}(\textit{phrase}, \text{"poor"}) \\
 &= \log_2 \frac{\frac{1}{N} \text{hits}(\textit{phrase} \text{ NEAR } \text{"excellent"})}{\frac{1}{N} \text{hits}(\textit{phrase}) \frac{1}{N} \text{hits}(\text{"excellent"})} - \log_2 \frac{\frac{1}{N} \text{hits}(\textit{phrase} \text{ NEAR } \text{"poor"})}{\frac{1}{N} \text{hits}(\textit{phrase}) \frac{1}{N} \text{hits}(\text{"poor"})} \\
 &= \log_2 \frac{\text{hits}(\textit{phrase} \text{ NEAR } \text{"excellent"})}{\text{hits}(\textit{phrase}) \text{hits}(\text{"excellent"})} \frac{\text{hits}(\textit{phrase}) \text{hits}(\text{"poor"})}{\text{hits}(\textit{phrase} \text{ NEAR } \text{"poor"})} \\
 &= \log_2 \frac{\square \text{hits}(\textit{phrase} \text{ NEAR } \text{"excellent"}) \text{hits}(\text{"poor"}) \square}{\square \text{hits}(\textit{phrase} \text{ NEAR } \text{"poor"}) \text{hits}(\text{"excellent"}) \square}
 \end{aligned}$$

Phrases from a thumbs-up review

Phrase	POS tags	Polarity
online service	JJ NN	2.8
online experience	JJ NN	2.3
direct deposit	JJ NN	1.3
local branch	JJ NN	0.42
...		
low fees	JJ NNS	0.33
true service	JJ NN	-0.73
other bank	JJ NN	-0.85
inconveniently located	JJ NN	-1.5
<i>Average</i>		0.32

Phrases from a thumbs-down review

Phrase	POS tags	Polarity
direct deposits	JJ NNS	5.8
online web	JJ NN	1.9
very handy	RB JJ	1.4
...		
virtual monopoly	JJ NN	-2.0
lesser evil	RBR JJ	-2.3
other problems	JJ NNS	-2.8
low funds	JJ NNS	-6.8
unethical practices	JJ NNS	-8.5
<i>Average</i>		-1.2

Results of Turney algorithm

- 410 reviews from Epinions
 - 170 (41%) negative
 - 240 (59%) positive
- Majority class baseline: 59%
- Turney algorithm: 74%

- Phrases rather than words
- Learns domain-specific information

Summary on Learning Lexicons

- **Advantages:**
 - Can be domain-specific
 - Can be more robust (more words)
- **Intuition**
 - Start with a seed set of words ('good', 'poor')
 - Find other words that have similar polarity:
 - Using "and" and "but"
 - Using words that occur nearby in the same document
 - Using WordNet synonyms and antonyms
 - Use seeds and semi-supervised learning to induce lexicons

Summary on Sentiment

- Generally modeled as classification task
 - Regression if prediction is ordinal
- Features:
 - Negation is important
 - ML: naïve Bayes, Logistic regression, etc
 - Finding subsets of words may help in other tasks
 - Hand-built polarity lexicons
 - Domain adaptation from one corpus
 - Use seeds and semi-supervised learning to induce lexicons
 - Use linguistic insights, e.g., co-occurrence with PMI

Harms in (ML-based) sentiment classifiers

- Kiritchenko and Mohammad (2018) found that most sentiment classifiers assign lower sentiment and more negative emotion to sentences with African American names in them.
- This perpetuates negative stereotypes that associate African Americans with negative emotions

Harms in (ML-based) toxicity classification

- Toxicity detection is the task of detecting hate speech, abuse, harassment, or other kinds of toxic language
- But some toxicity classifiers incorrectly flag as being toxic sentences that are non-toxic but simply mention identities like blind people, women, or gay people.
- This could lead to censorship of discussion about these groups.

What causes these harms?

- Can be caused by:
 - Problems in the training data; machine learning systems are known to amplify the biases in their training data.
 - Problems in the human labels
 - Problems in the resources used (like lexicons)
 - Problems in model architecture (like what the model is trained to optimized)
- Mitigation of these harms is an open research area
- Meanwhile: **model cards**

Model Cards (Mitchell et al., 2019)

- For each algorithm you release, document:
 - training algorithms and parameters
 - training data sources, motivation, and preprocessing
 - evaluation data sources, motivation, and preprocessing
 - intended use and users
 - model performance across different demographic or other groups and environmental situations