

CS 6120/CS 4120: Natural Language Processing

Instructor: Prof. Lu Wang

College of Computer and Information Science

Northeastern University

Webpage: www.ccs.neu.edu/home/luwang

Logistics

- Minor updates on assignment 2 (available on both course website and blackboard).
 - Questions 2's rules are updated.
- Prelim takes place on November 6, 2:50-4:30pm.
 - Open-book
 - You can take computer but no Internet access, no messaging.
 - Textbooks and handouts are also fine.
 - Practice questions for prelim are posted on blackboard.

Sentiment Analysis

- Sentiment analysis tasks
- Features for building machine learning models
- Sentiment lexicons

Positive or negative movie review?



- unbelievably disappointing



- Full of zany characters and richly applied satire, and some great plot twists



- this is the greatest screwball comedy ever filmed




- It was pathetic. The worst part about it was the boxing scenes.



Apple iPhone 8 - 256 GB - Gold - T-Mobile - GSM

\$850 online ★★★★★ 2,602 product reviews

Reviews




★★★★★ **iPhone 8 vs iPhone 6s Plus, iPhone SE, & iPhone X.** – December 24, 2017

– Review provided by Best Buy
December 24, 2017

I have recent experience using both the iPhone SE and iPhone 6s Plus. The Plus model was too big since I use a case with a belt clip to carry the phone, and the SE's screen was a bit too small. I am going to compare my review of the iPhone 8 (purchased unlocked at full price and used with Verizon prepaid) mostly to the iPhone 6s Plus, but one has to understand that the SE especially at prepaid price is an excellent, outstanding phone too with almost all of the same features as the Plus!) Pros: The iPhone 8 is an upgrade in a few ways. Apple includes a compare feature on its website so I won't go into all of the details, but I will try to address the ones that are upgrades to the iPhone 6s Plus. True Tone display does make the screen easier to read because the lighting isn't ... [more »](#)


Twitter sentiment versus Gallup Poll of Consumer Confidence

Brendan O'Connor, Ramnath Balasubramanian, 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.



Sentiment analysis has many other names

- Opinion extraction
- Opinion mining
- Sentiment mining
- Subjectivity analysis

Why sentiment analysis?

- *Movie*: is this review positive or negative?
- *Products*: what do people think about the new iPhone?
- *Public sentiment*: how is consumer confidence? Is despair increasing?
- *Politics*: what do people think about this candidate or issue?
- *Prediction*: predict election outcomes or market trends from sentiment

Scherer Typology of Affective States

- **Emotion:** brief organically synchronized ... evaluation of a major event
 - *angry, sad, joyful, fearful, ashamed, proud, elated*
- **Mood:** diffuse non-caused low-intensity long-duration change in subjective feeling
 - *cheerful, gloomy, irritable, listless, depressed, buoyant*
- **Interpersonal stances:** affective stance toward another person in a specific interaction
 - *friendly, flirtatious, distant, cold, warm, supportive, contemptuous*
- **Attitudes:** enduring, affectively colored beliefs, dispositions towards objects or persons
 - *liking, loving, hating, valuing, desiring*
- **Personality traits:** stable personality dispositions and typical behavior tendencies
 - *nervous, anxious, reckless, morose, hostile, jealous*

Scherer Typology of Affective States

- Emotion and Mood
 - Annoyance in talking to dialog systems
 - Uncertainty of students in tutoring
 - Detecting trauma or depression
- Interpersonal Stance
 - Romantic interest, flirtation, friendliness
 - Alignment/accommodation/entrainment
- Attitudes = Sentiment (positive or negative)
 - Movie or Products or Politics: is a text positive or negative?
 - “Twitter mood predicts the stock market.”
- Personality Traits
 - Open, Consciencious, Extroverted, Anxious

Scherer Typology of Affective States

- **Emotion:** brief organically synchronized ... evaluation of a major event
 - *angry, sad, joyful, fearful, ashamed, proud, elated*
- **Mood:** diffuse non-caused low-intensity long-duration change in subjective feeling
 - *cheerful, gloomy, irritable, listless, depressed, buoyant*
- **Interpersonal stances:** affective stance toward another person in a specific interaction
 - *friendly, flirtatious, distant, cold, warm, supportive, contemptuous*
- **Attitudes: enduring, affectively colored beliefs, dispositions towards objects or persons**
 - *liking, loving, hating, valuing, desiring*
- **Personality traits:** stable personality dispositions and typical behavior tendencies
 - *nervous, anxious, reckless, morose, hostile, jealous*

Sentiment Analysis

- Extraction of opinions and attitudes from text and speech
- When we say “sentiment analysis”
 - We often mean a binary or an ordinal task
 - like X/ dislike X
 - one-star to 5-stars

Sentiment Analysis

- Sentiment analysis is the detection of **attitudes**
“enduring, affectively colored beliefs, dispositions towards objects or persons”

Emily told Charlie that the new movie is disappointing.

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, often 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

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 Analysis

- Sentiment analysis tasks
- • Features for building machine learning models
- Sentiment lexicons

What features to design?



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 .

Negation in Sentiment Analysis

They have not succeeded, and will never succeed, in breaking the will of this valiant people.

Negation in Sentiment Analysis

They have not **succeeded**, and will never succeed, in breaking the will of this valiant people.

Negation in Sentiment Analysis

They have **not succeeded**, and will never succeed, in breaking the will of this valiant people.

Negation in Sentiment Analysis

They **have not succeeded, and will never succeed, in breaking** the will of this valiant people.

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

Reminder: Naïve Bayes

$$c_{NB} = \operatorname{argmax}_{c_j \in C} P(c_j) \prod_{i \in \text{positions}} P(w_i | c_j)$$

$$\hat{P}(w | c) = \frac{\text{count}(w, c) + 1}{\text{count}(c) + |V|}$$

Binarized (Boolean feature)

- Intuition:
 - For sentiment (and for other text classification domains)
 - Word occurrence may matter more than word frequency
 - The occurrence of the word *fantastic* tells us a lot
 - The fact that it occurs 5 times may not tell us much more.
 - Boolean Multinomial Naïve Bayes
 - Clips all the word counts in each document at 1

Boolean Multinomial Naïve Bayes: Learning

- From training corpus, extract Vocabulary
- Calculate $P(c_j)$ terms
 - For each c_j in C do
 - $docs_j \leftarrow$ all docs with class = c_j
- Calculate $P(w_k | c_j)$ terms
 - Remove duplicates in each doc:
 - For each word type w in doc_j
 - Retain only a single instance of w
 - $Text_j \leftarrow$ single doc containing all $docs_j$
 - For each word w_k in *Vocabulary*
 - $n_k \leftarrow$ # of occurrences of w_k in $Text_j$

$$P(c_j) \leftarrow \frac{|docs_j|}{|\text{total \# documents}|}$$

$$P(w_k | c_j) \leftarrow \frac{n_k + \alpha}{n + \alpha |Vocabulary|}$$

Boolean Multinomial Naïve Bayes on a test document d

- First remove all duplicate words from d
- Then compute NB using the same equation:

$$c_{NB} = \operatorname{argmax}_{c_j \in C} P(c_j) \prod_{i \in \text{positions}} P(w_i | c_j)$$

Binarized (Boolean feature) Multinomial Naïve Bayes

B. Pang, L. Lee, and S. Vaithyanathan. 2002. Thumbs up? Sentiment Classification using Machine Learning Techniques. EMNLP-2002, 79—86.

V. Metsis, I. Androutsopoulos, G. Paliouras. 2006. Spam Filtering with Naive Bayes – Which Naive Bayes? CEAS 2006 - Third Conference on Email and Anti-Spam.

K.-M. Schneider. 2004. On word frequency information and negative evidence in Naive Bayes text classification. ICANLP, 474-485.

JD Rennie, L Shih, J Teevan. 2003. Tackling the poor assumptions of naive bayes text classifiers. ICML 2003

- Binary seems to work better than full word counts
- Other possibility: $\log(\text{freq}(w))$

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”

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 Analysis

- Sentiment analysis tasks
- Features for building machine learning models
- • Sentiment lexicons

- Adjectives

- positive: honest important mature large patient

- He is the only honest man in Washington.

- Her writing is unbelievably mature and is only likely to get better.

- To humour me my patient father agrees yet again to my choice of film

- negative: harmful hypocritical inefficient insecure

- It was a macabre and hypocritical circus.

- Why are they being so inefficient ?

- Verbs
 - positive: **praise, love**
 - negative: **blame, criticize**
- Nouns
 - positive: **pleasure, enjoyment**
 - negative: **pain, criticism**

Phrases

- Phrases containing adjectives and adverbs
 - positive: high intelligence, low cost
 - negative: little variation, many troubles

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:
 - Positiv (1915 words) and Negativ (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*)
- Not free though!

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 from 8221 lemmas
 - 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

Disagreements between polarity lexicons

Christopher Potts, [Sentiment Tutorial](#), 2011

	Opinion Lexicon	General Inquirer	SentiWordNet	LIWC
MPQA	33/5402 (0.6%)	49/2867 (2%)	1127/4214 (27%)	12/363 (3%)
Opinion Lexicon		32/2411 (1%)	1004/3994 (25%)	9/403 (2%)
General Inquirer			520/2306 (23%)	1/204 (0.5%)
SentiWordNet				174/694 (25%)
LIWC				

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**:
$$P(w | c) = \frac{f(w, c)}{\sum_{w \in C} f(w, c)}$$

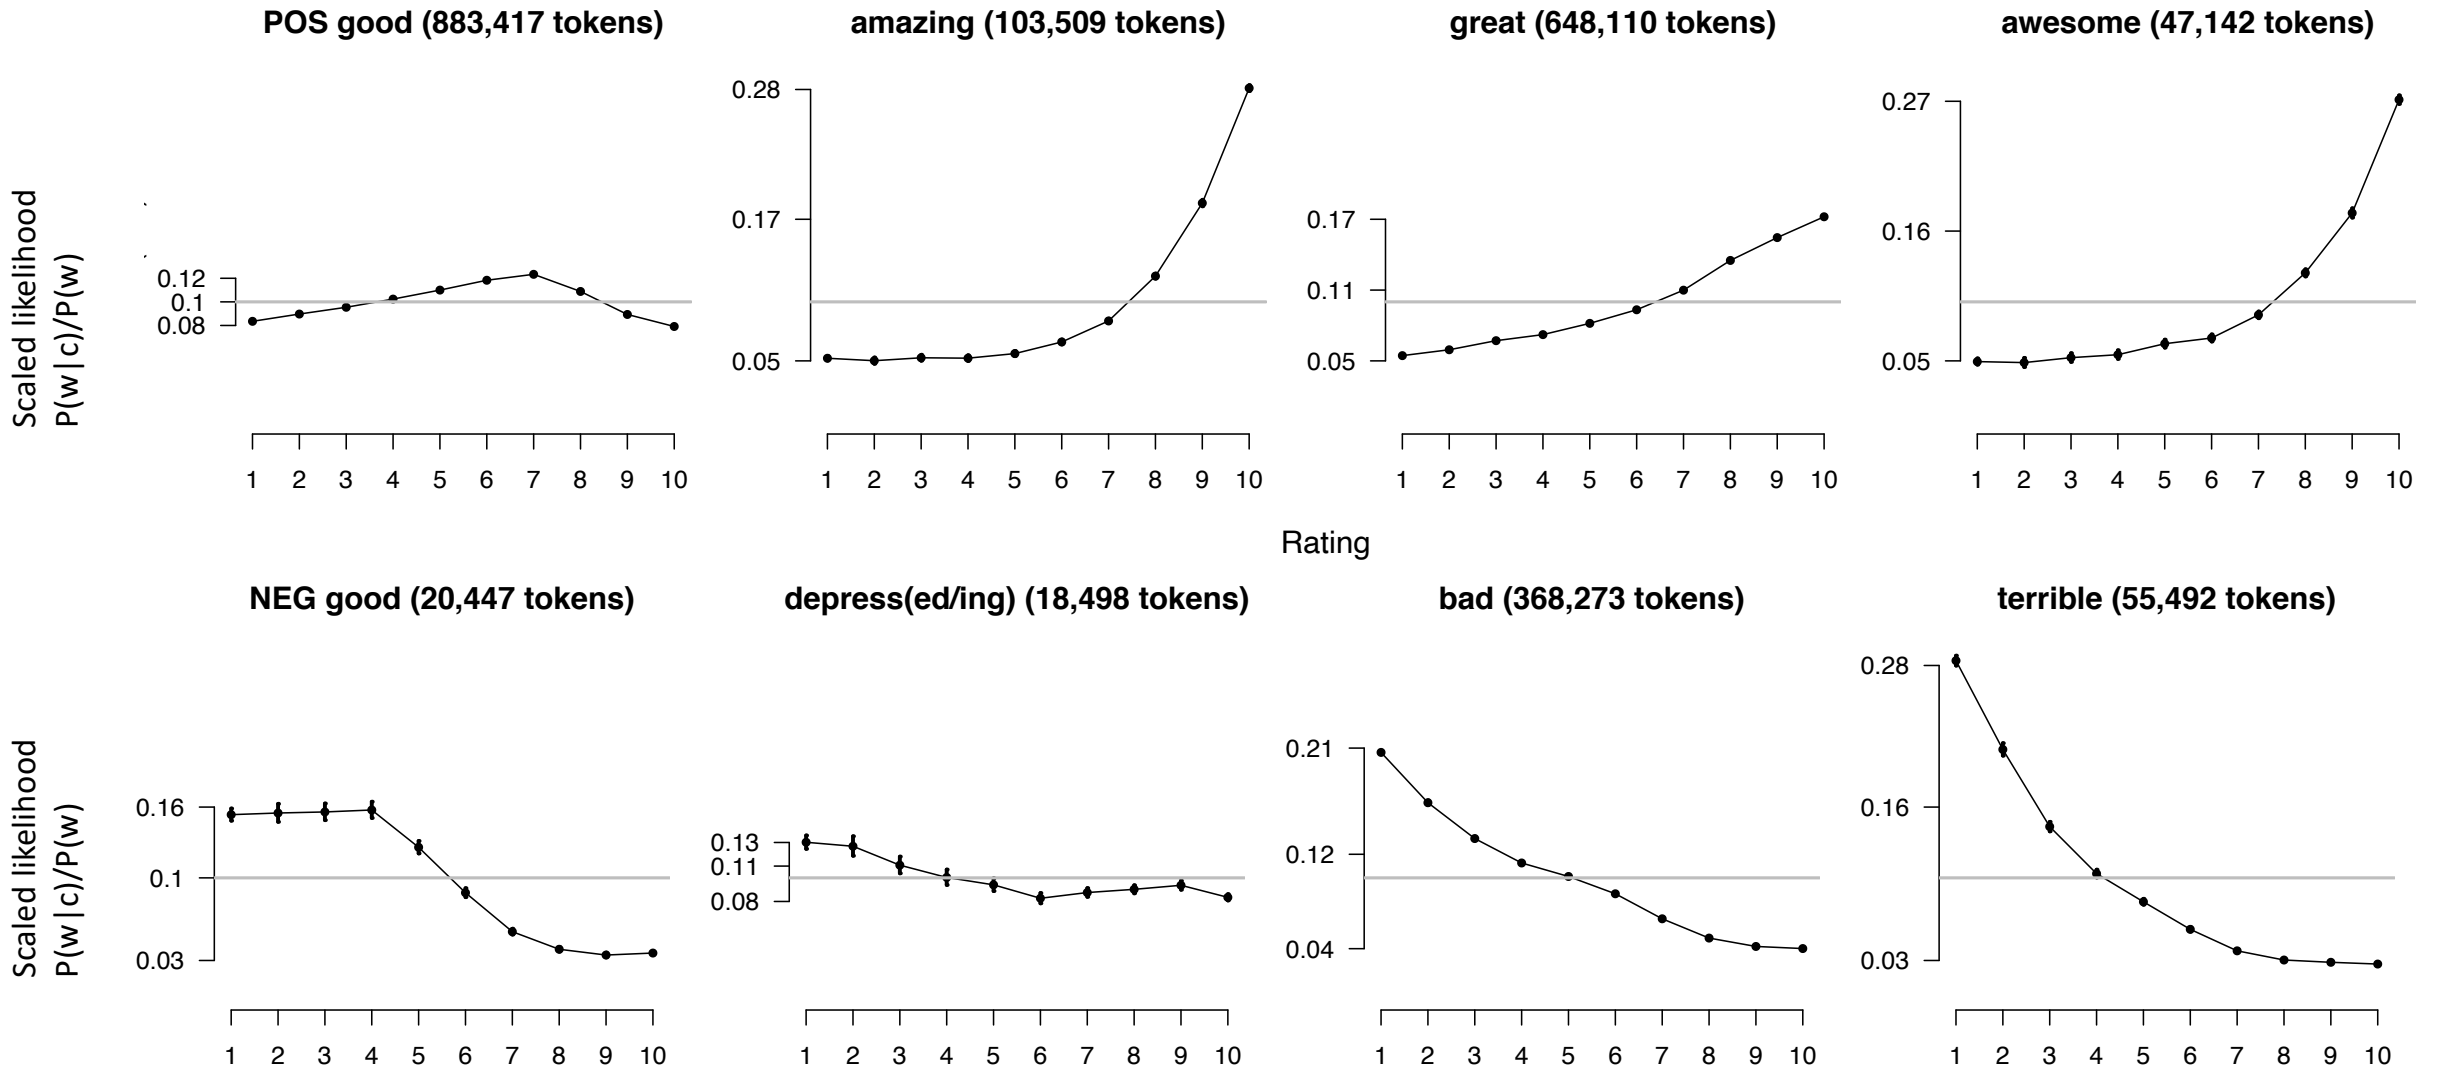
- Make them comparable between words
 - **Scaled likelihood**:

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



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)



Learning Sentiment Lexicons

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 have the same polarity
 - 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** ...](#)

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 ...](#)

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


- Supervised classifier assigns “polarity similarity” to each word pair, resulting in graph:



Hatzivassiloglou & McKeown 1997

Step 4

- Clustering for partitioning the graph into two



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?

Pointwise Mutual Information

- **Pointwise mutual information:**

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

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

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 $\text{count}(\text{word}) / N$
 - -> unigram probability
- $P(\text{word}_1, \text{word}_2)$ by $\text{count}(\text{word}_1 \text{ NEAR } \text{word}_2) / N$
 - -> "NEAR" needs to be defined by window size, e.g. ± 3 words

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

$$\text{Polarity}(\textit{phrase}) = \text{PMI}(\textit{phrase}, \text{"excellent"}) - \text{PMI}(\textit{phrase}, \text{"poor"})$$

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

Using WordNet to learn polarity

S.M. Kim and E. Hovy. 2004. Determining the sentiment of opinions. COLING 2004

M. Hu and B. Liu. Mining and summarizing customer reviews. In Proceedings of KDD, 2004

- WordNet: online thesaurus (covered in later lecture).
- Create positive (“good”) and negative seed-words (“terrible”)
- Find Synonyms and Antonyms
 - Positive Set: Add synonyms of positive words (“well”) and antonyms of negative words
 - Negative Set: Add synonyms of negative words (“awful”) and antonyms of positive words (“evil”)
- Repeat, following chains of synonyms