Sentiment Analysis

These slides are based on:

Dan Jurafsky and James H. Martin, Speech and Language Processing (3rd ed. draft)
https://web.stanford.edu/~jurafsky/slp3/
(Chapters 7 and 21)

Outline

- Introduction
- Sentiment lexicons
- Expanding lexicons
- Sentiment and ratings of reviews
- Supervised baseline algorithm for polarity classification
- More complex tasks

Introduction

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.

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.

Applications of lexicons

http://hedonometer.org

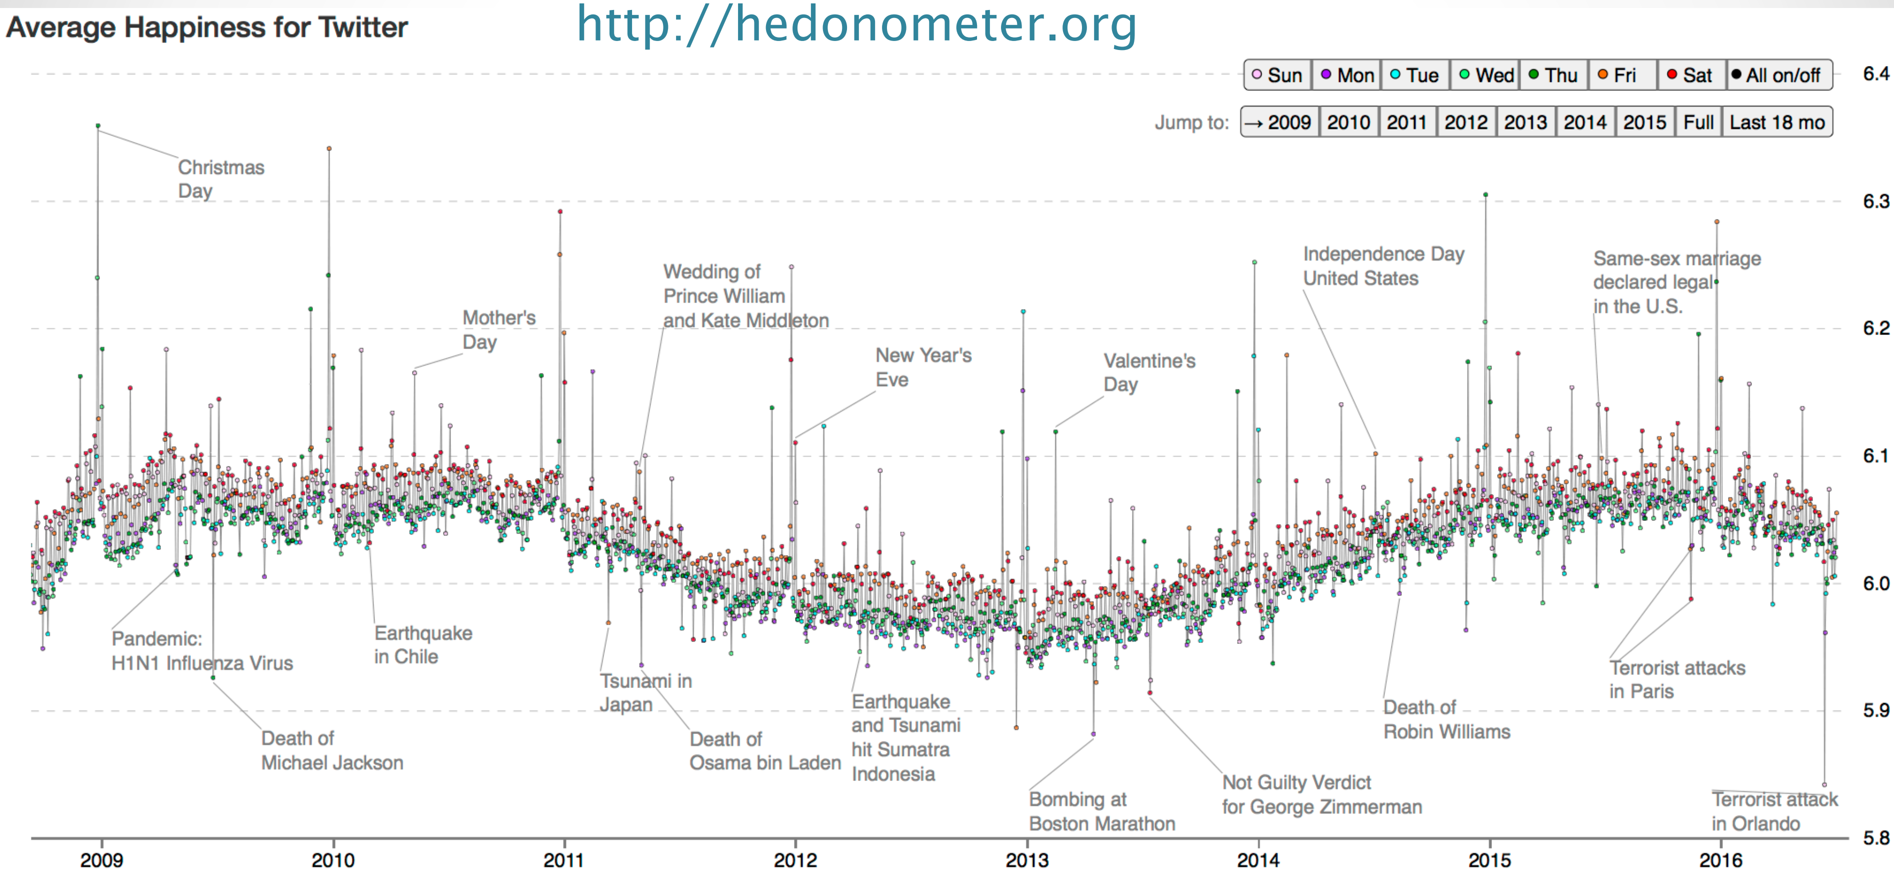

Target Sentiment on Twitter

http://www.sentiment140.com

Sentiment140 Tweet Like 864 G+1 205

teslamotors English Search

Sentiment analysis for teslamotors

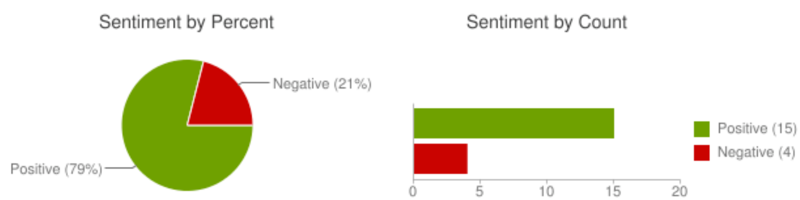

brijwaasi: RT @IndiaTodayTech: Sleeping driver, terrified granny among **@TeslaMotors** Autopilot users on YouTube https://t.co/PX4dFIYUAQ https://t.co/7g ?
Posted: 1 hour ago

stevehammershow: https://t.co/9PVAStujXB #SelfDrivingCar not ready for prime time yet. Perhaps #Mars might be a better introduction? @google **@TeslaMotors**
Posted: 1 hour ago

SanjayAroraIN: @timesofindia @TOIAuto **@TeslaMotors** #Google has had driverless cars on road longer than #Tesla but they have not deployed untested tech
Posted: 1 hour ago

silverfighter: **@TeslaMotors** statement to the auto-pilot crash https://t.co/1o1iLmgAH5 Even with the last paragraph-a human died / not just a customer issue
Posted: 1 hour ago

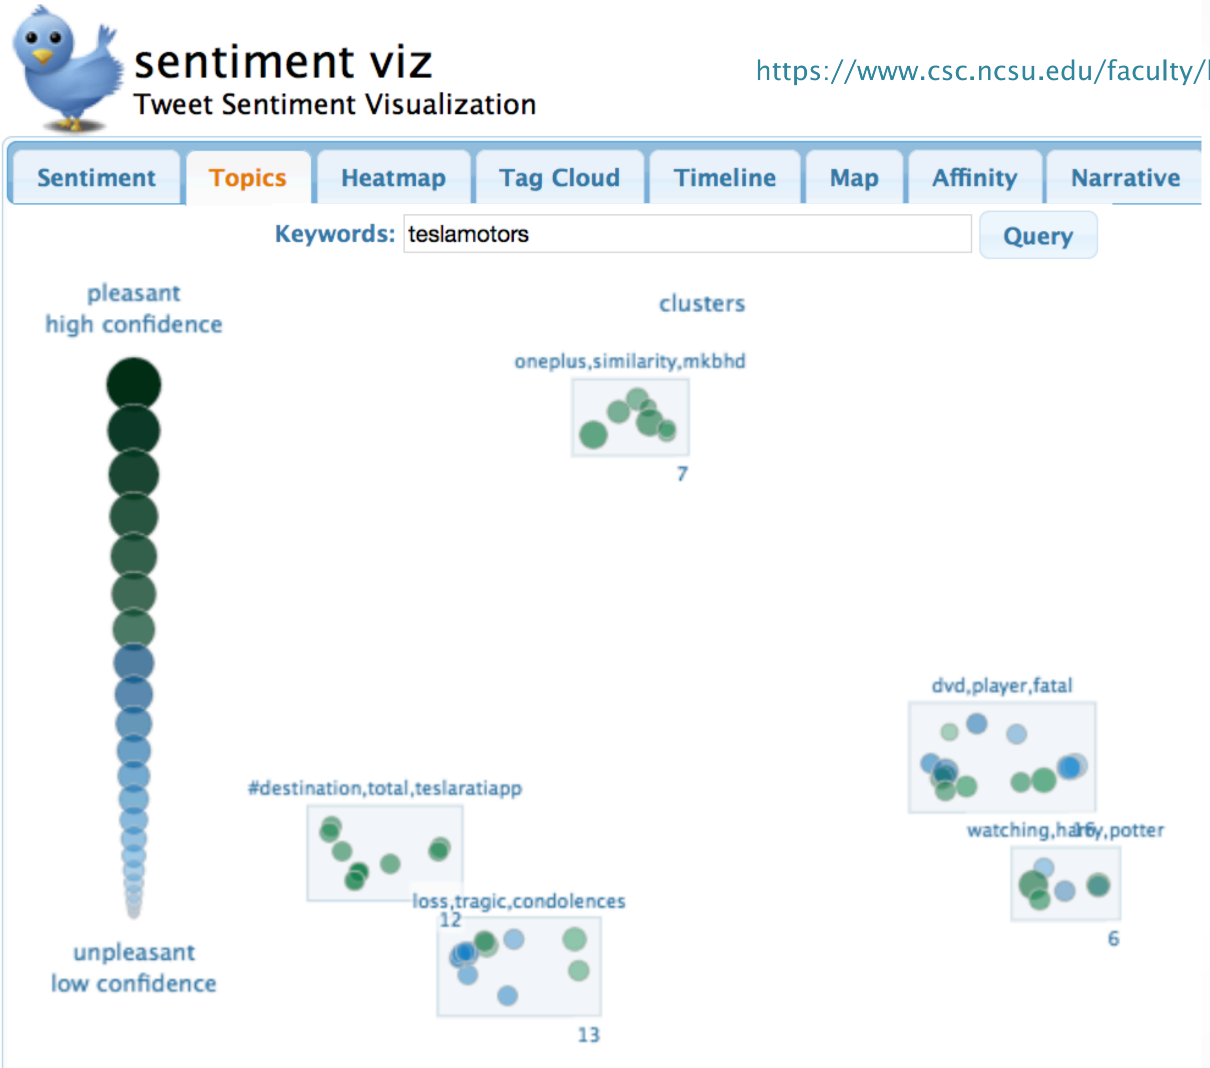

https://www.csc.ncsu.edu/faculty/healey/tweet_viz/tweet_app/

Product search at Google Shopping

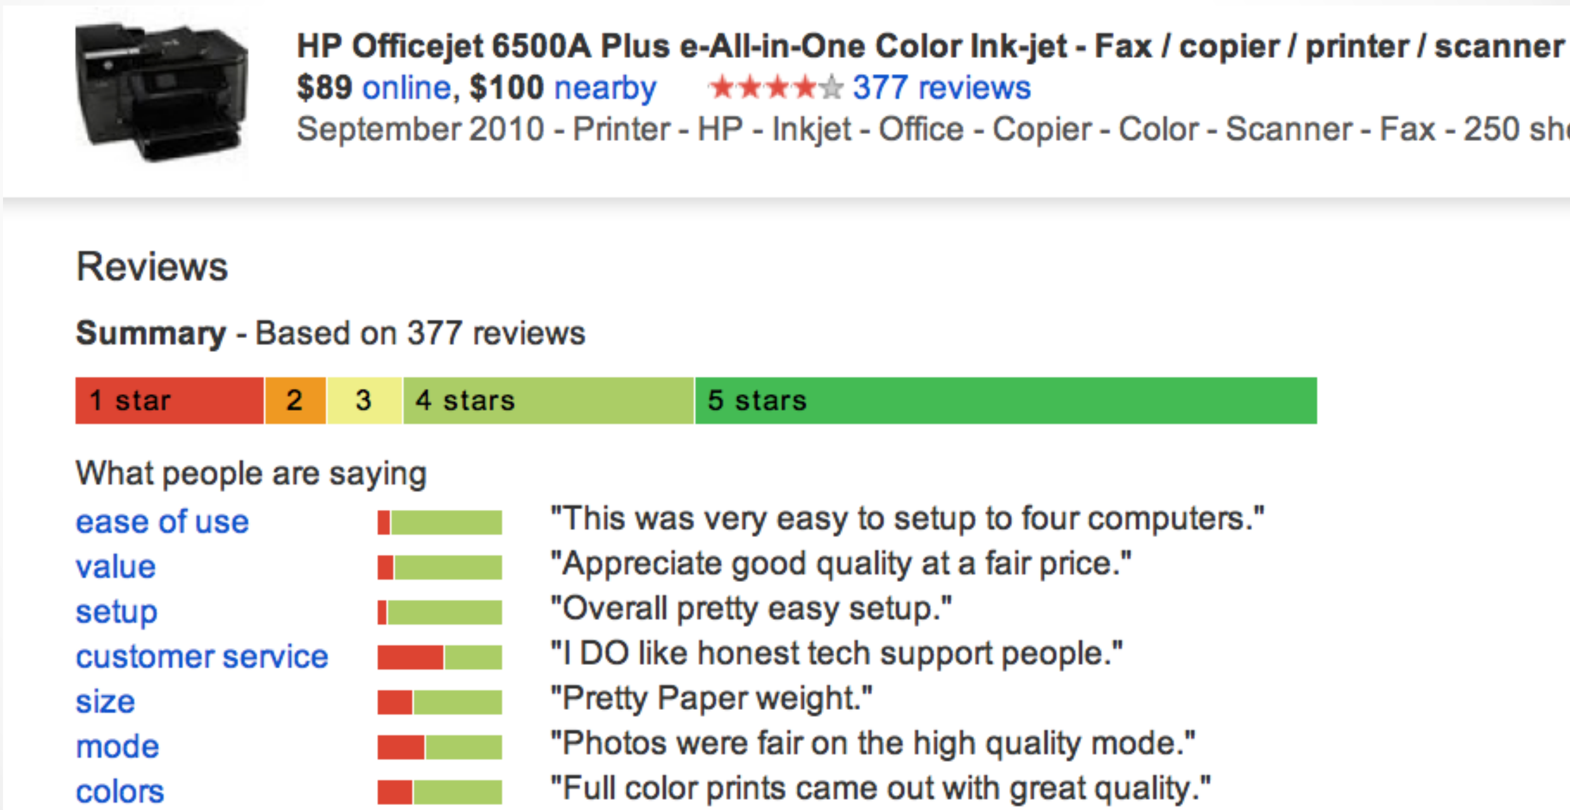

Sentiment analysis has many other names

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

Opinions are pervasive

- *Movie*: is this review positive or negative?
- *Products*: what do people think about the new iPhone or Tesla Motors?
- *Public sentiment*: what do people think about EURO 2016?
- *Politics*: what do people think about this candidate or issue?

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**: 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

- 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 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:
    - 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)

- Home page: http://www.liwc.net/
- 2300 words, >70 cathegories
- **Affective Processes**
  - negative emotion (*bad, weird, hate, problem, tough*)
  - positive emotion (*love, nice, sweet*)
  - anxiety, anger, sadness
- Other processes: social, cognitive, perceptual, biological
- Personal concerns, informal language, drives

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

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%)**

Learning sentiment lexicons

Semi-supervised learning of lexicons

- Use a set of seed words
  - A few labeled examples
  - A few hand-built patterns
- Bootstrap a larger lexicon

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
  - NOT fair and brutal, legitimate and corrupt
- Adjectives conjoined by "***but***" do not
  - fair **but** brutal

Step 1

- Label **seed set** of 1336 adjectives
  - 657 positive
    - adequate central clever famous intelligent remarkable reputed sensitive slender thriving…
  - 679 negative
    - contagious drunken ignorant lanky listless primitive strident troublesome unresolved unsuspecting…

Step 2

- Assign "polarity similarity" to each word pair, resulting in graph:

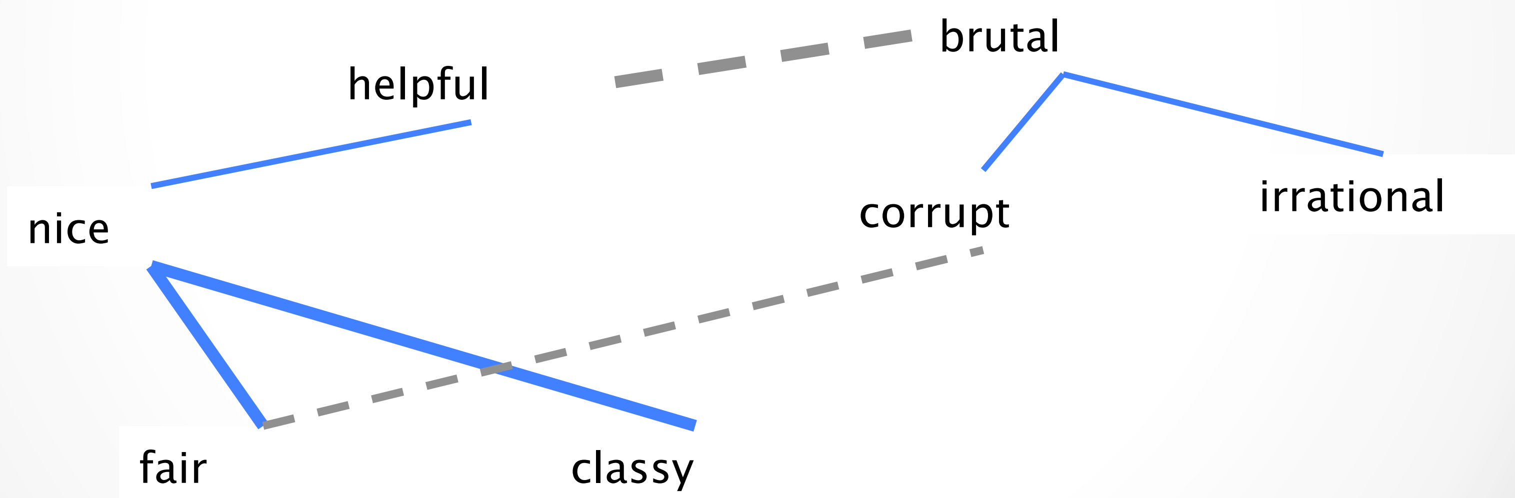

Step 3

- Clustering for partitioning the graph into two

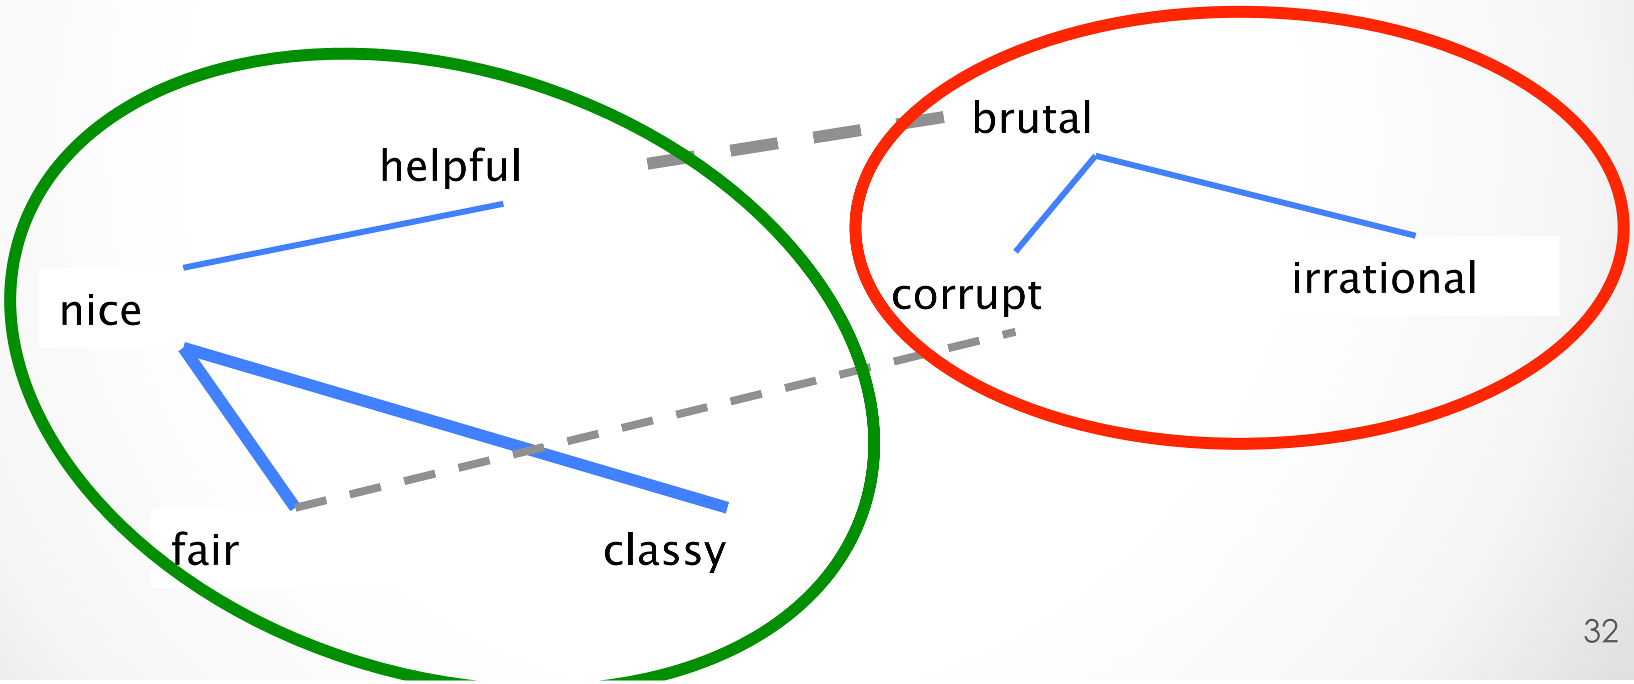

Output polarity lexicon

- 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…

errors are present

Using WordNet to learn polarity

- WordNet: online thesaurus
- 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
- Repeat, following chains of synonyms

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

How to measure polarity of a phrase?

- Positive phrases co-occur more with *"excellent"*
- Negative phrases co-occur more with *"poor"*

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

Polarity of exemplary phrases

Phrase from thumbs-up reviews	Polarity
online service	2.8
online experience	2.3
direct deposit	1.3
local branch	0.42
…	
low fees	0.33
true service	-0.73
other bank	-0.85
inconveniently located	-1.5
Average	**0.32**

Phrases from thumbs-down reviews	Polarity
direct deposits	5.8
online web	1.9
very handy	1.4
…	
virtual monopoly	-2.0
lesser evil	-2.3
other problems	-2.8
low funds	-6.8
unethical practices	-8.5
Average	**-1.2**

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

Sentiment versus user ratings on social media

Analyzing the polarity words 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:**

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

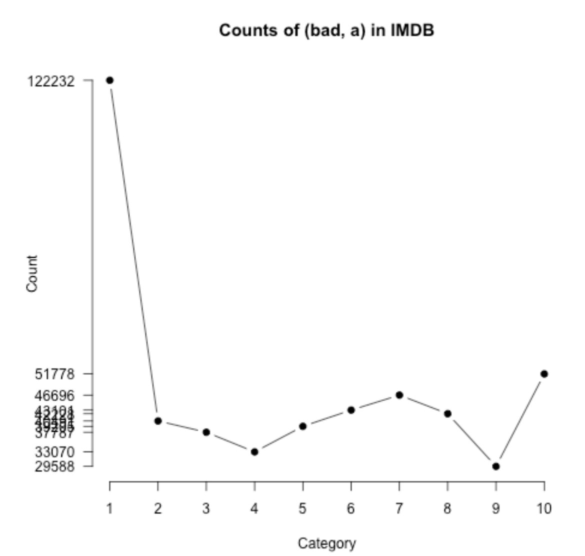

- Make them comparable between words
  - **Scaled likelihood:**

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

Analyzing the polarity words in IMDB

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

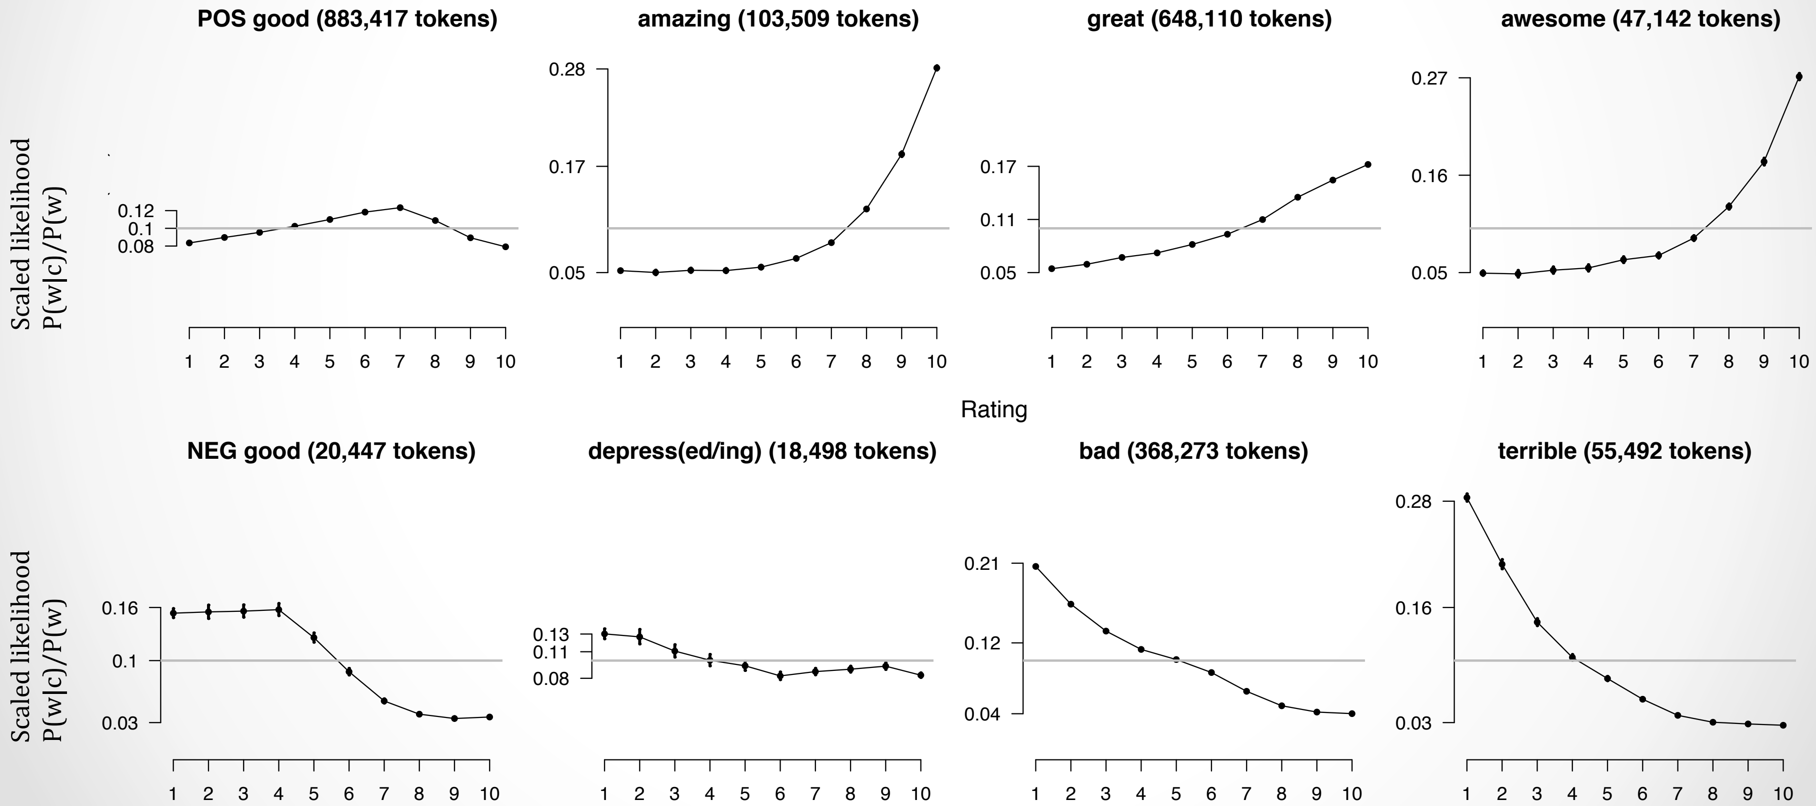

Logical negation and its polarity in IMDB

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

More negation in negative sentiment

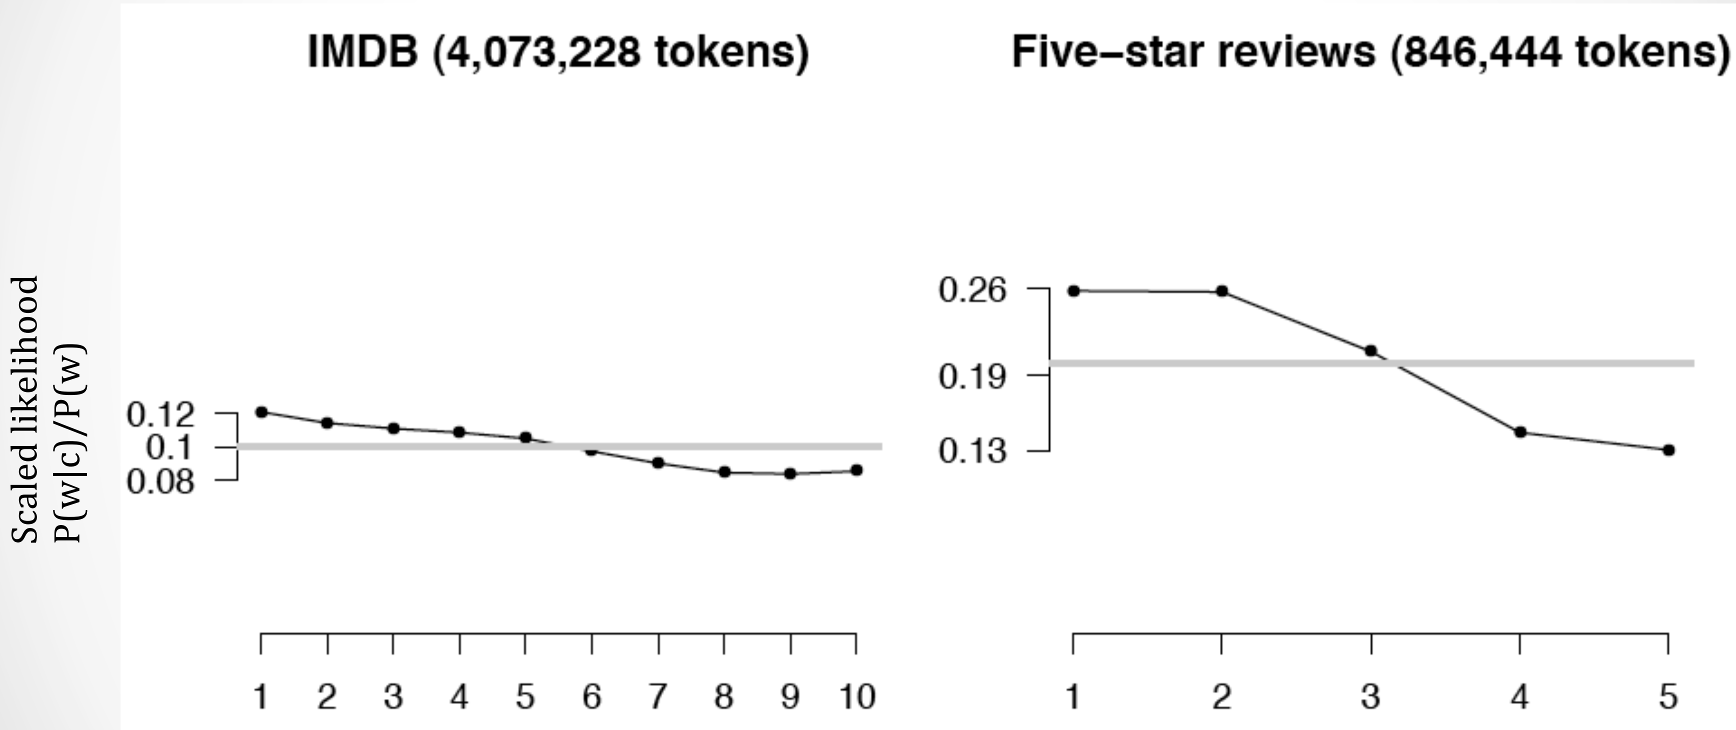

Potts, Christopher. 2011. On the negativity of negation. SALT 20, 636-659.

A baseline supervised algorithm for sentiment classification

Sentiment/rating classification in movie reviews

- 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

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

Baseline Algorithm

- Tokenization
- Feature Extraction
- Classification using different classifiers
  - Naïve Bayes
  - SVM
  - etc.

Pang, B., & Lee, L. (2006). Opinion Mining and Sentiment Analysis. Foundations and Trends® in InformatioPang, B., & Lee, L. (2006). Opinion Mining and Sentiment Analysis. Foundations and Trends in Information Retrieval, 1(2), 91–231.

Sentiment tokenization Issues

- Deal with HTML and XML markup
- Twitter markup (names, hashtags)
- Capitalization (preserve for words in all caps)
- Phone numbers, dates
- Emoticons
- Useful code:
    - Christopher Potts sentiment tokenizer
    - Brendan O'Connor twitter tokenizer

Naïve Bayes

$$c_{NB} = \underset{c_j \in C}{\operatorname{argmax}} P(c_j) \prod_{i \in positions} P(w_i \mid c_j)$$

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

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 or all words
    - All words turns out to work better, at least on this data

- Frequency of words or presence?

Negation

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
```

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.

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."'
  - A movie revie:
    - "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."

More complex tasks

Finding aspect/attribute/target of sentiment

M. Hu and B. Liu. 2004. Mining and summarizing customer reviews. In Proceedings of KDD.
S. Blair-Goldensohn, K. Hannan, R. McDonald, T. Neylon, G. Reis, and J. Reynar. 2008. Building a Sentiment Summarizer for Local Service Reviews. WWW Workshop.

- Frequent phrases + rules
  - Find all highly frequent phrases across reviews ("`fish tacos`")
  - Filter by rules like "occurs right after sentiment word"
    - "…`great fish tacos`" means `fish tacos` a likely aspect

Finding aspect/attribute/target of sentiment

- The aspect name may not be in the sentence
- For restaurants/hotels, aspects are well-understood
- Supervised classification
  - Hand-label a small corpus of restaurant review sentences with aspect
    - food, décor, service, value, NONE
  - Train a classifier to assign an aspect to a sentence
    - "Given this sentence, is the aspect *food, décor, service, value,* or *NONE*"

Putting it all together: Finding sentiment for aspects

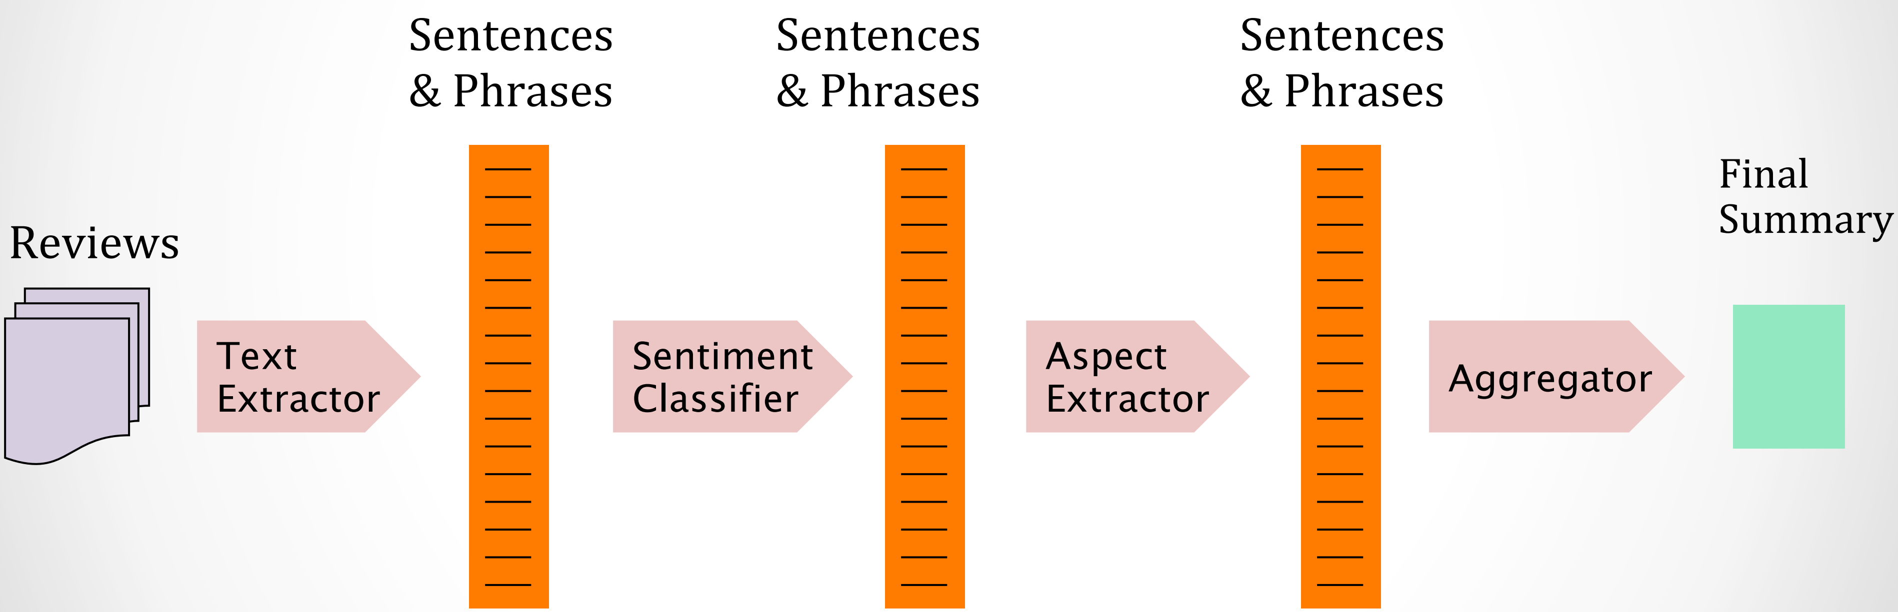

S. Blair-Goldensohn, K. Hannan, R. McDonald, T. Neylon, G. Reis, and J. Reynar. 2008. Building a Sentiment Summarizer for Local Service Reviews. WWW Workshop

Summary on Sentiment

- Generally modeled as classification or regression task
  - predict a binary or ordinal label
- Features:
  - Negation is important
  - Using all words (in naïve bayes) works well for some tasks
  - Finding subsets of words may help in other tasks
    - Hand-built polarity lexicons
    - Use seeds and semi-supervised learning to induce lexicons

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*

Computational work on other affective states

- **Emotion**:
  - Detecting annoyed callers to dialogue system
  - Detecting confused/frustrated versus confident students
- **Mood**:
  - Finding traumatized or depressed writers
- **Interpersonal stances**:
  - Positive and negative ties and structural balance theory
- **Personality traits**:
  - Predicting personality traits from Facebook likes

Labeled interpersonal stances

Positive and negative interpersonal actions in a massive multiplayer online game: (alliance, trade) vs. (aggression, punishment).

Structural balance theory:

- A friend of a friend tends to be a friend
- A friend of an enemy tends to be an enemy

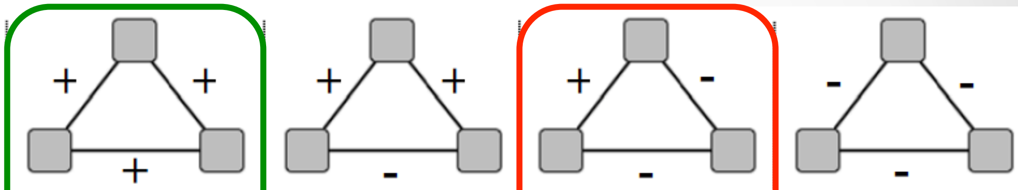

Strong formulation of balance	B	U	B	U
Weak formulation of balance	B	U	B	B
N_{Δ}	26,329	4,428	39,519	8,032
N_{Δ}^{rand}	10,608	30,145	28,545	9,009

Szell, M., Lambiotte, R., & Thurner, S. (2010). Multirelational organization of large-scale social networks in an online world. Proceedings of the National Academy of Sciences, 107(31).

Leskovec, J., Huttenlocher, D., & Kleinberg, J. (2010). Predicting positive and negative links in online social networks. In Proceedings of the 19th international conference on World wide web (pp. 641–650).

Sentiment Analysis

These slides are based on:

Dan Jurafsky and James H. Martin, Speech and Language Processing (3rd ed. draft)
https://web.stanford.edu/~jurafsky/slp3/
(Chapters 7 and 21)