
Subjective Analysis of Text:
Sentiment Analysis
Opinion Analysis
(using some material from Dan Jurafsky)

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

Types of tasks:

- 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.

Types of Tasks: Google Product Search



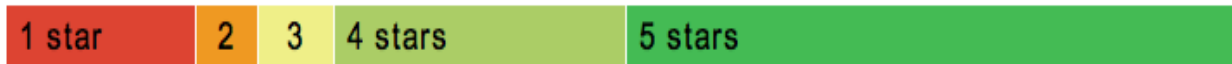
HP Officejet 6500A Plus e-All-in-One Color Ink-jet - Fax / copier / printer / scanner

\$89 online, \$100 nearby ★★★★★ 377 reviews

September 2010 - Printer - HP - Inkjet - Office - Copier - Color - Scanner - Fax - 250 sh

Reviews

Summary - Based on 377 reviews



What people are saying

ease of use		"This was very easy to setup to four computers."
value		"Appreciate good quality at a fair price."
setup		"Overall pretty easy setup."
customer service		"I DO like honest tech support people."
size		"Pretty Paper weight."
mode		"Photos were fair on the high quality mode."
colors		"Full color prints came out with great quality."

Types of tasks: Bing Shopping

HP Officejet 6500A E710N Multifunction Printer

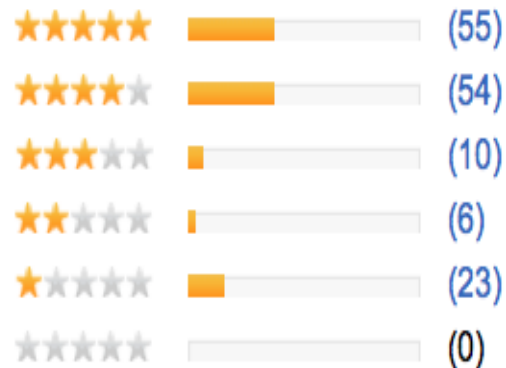
[Product summary](#) [Find best price](#) **Customer reviews** [Specifications](#) [Related items](#)



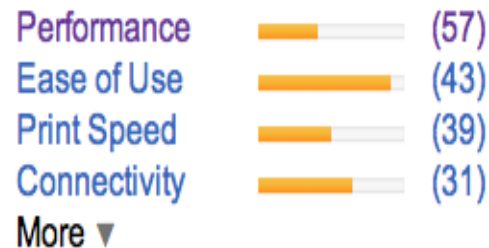
\$121.53 - \$242.39 (14 stores)

Compare

Average rating ★★★★★ (144)



Most mentioned



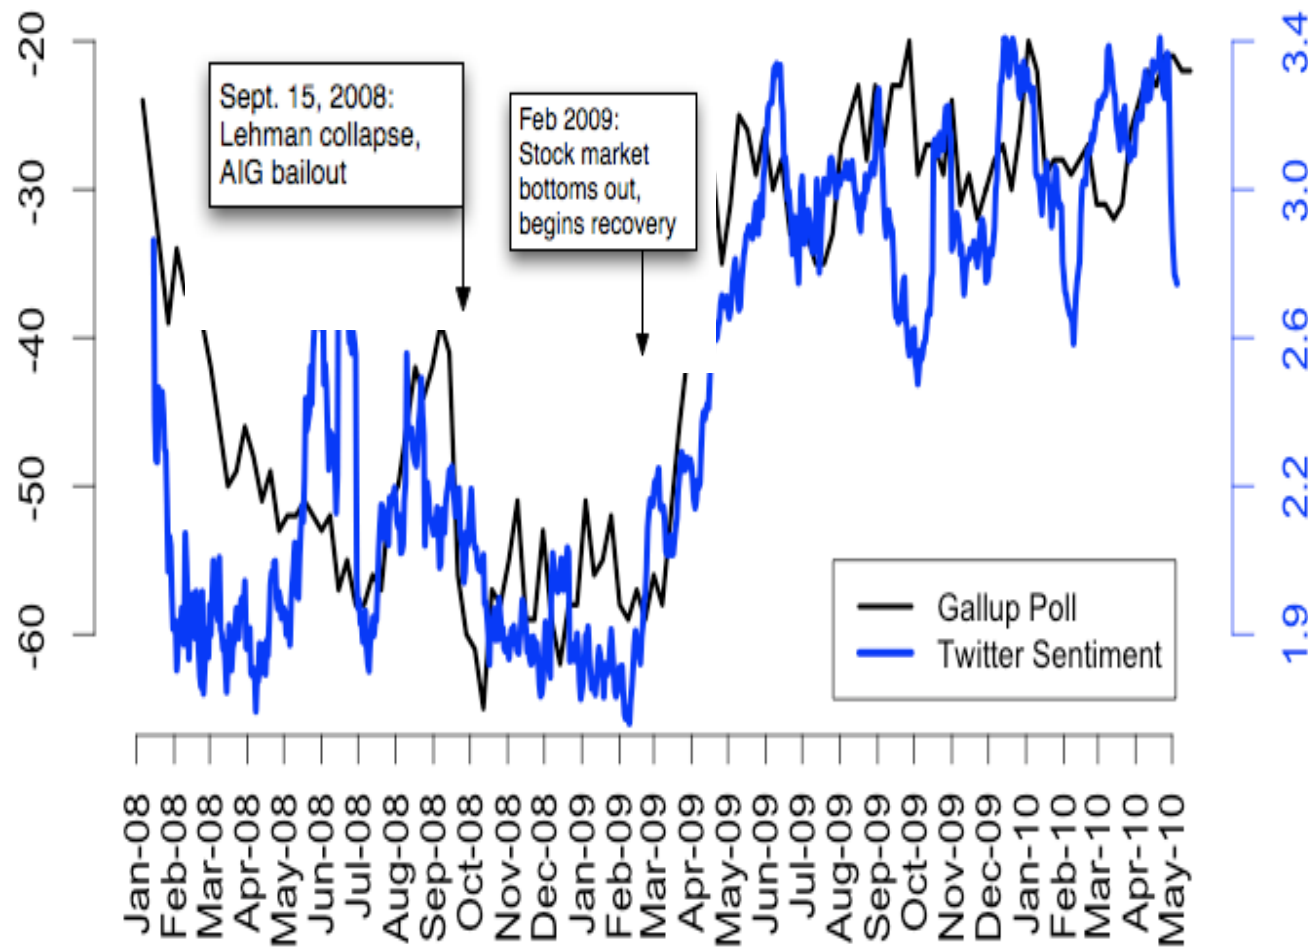
Show reviews by source



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

window = 15, $r = 0.804$



Target Sentiment on Twitter

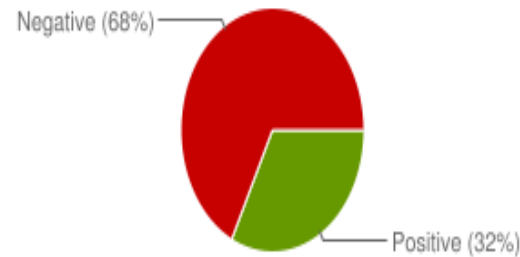
- [Twitter Sentiment App](#)
- Alec Go, Richa Bhayani, Lei Huang. 2009. Twitter Sentiment Classification using Distant Supervision

Type in a word and we'll highlight the good and the bad

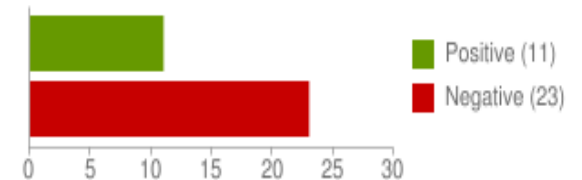
[Save this search](#)

Sentiment analysis for "united airlines"

Sentiment by Percent



Sentiment by Count



[jjacobson](#): OMG... Could **@United airlines** have worse customer service? W8g now 15 minutes on hold 4 questions about a flight 2DAY that need a human.
Posted 2 hours ago

[12345clumsy6789](#): I hate **United Airlines** Ceiling!!! Fukn impossible to get my conduit in this damn mess! ?
Posted 2 hours ago

[EMLandPRGbelgiu](#): EML/PRG fly with Q8 **united airlines** and 24seven to an exotic destination. <http://t.co/Z9QloAjF>
Posted 2 hours ago

[CountAdam](#): FANTASTIC customer service from **United Airlines** at XNA today. Is tweet more, but cell phones off now!
Posted 4 hours ago

Terminology

- For the more subtle aspects of the semantics of text:
- **Affective aspects of text** is that which is “influenced by or resulting from emotions”
 - One aspect of non-factual aspects of text
- **Subjective aspects of text**
“The **linguistic** expression of somebody’ s **opinions, sentiments**, emotions, evaluations, beliefs, speculations (*private states*)”
 - A private state is not open to objective observation or verification
 - Subjectivity analysis would classify parts of text as to whether it was subjective or objective

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

- **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 Task Levels

- Simplest task is polarity:
 - Is the attitude of this text positive or negative?
 - Negative / positive attitude of reporter / blogger
 - Favorable / unfavorable review of a product
 - Right / left political leaning of speaker
- More complex:
 - Rank the attitude of this text from 1 to 5
- Advanced:
 - Detect the target, source, or complex attitude types
- May also be referred to as
opinion extraction, opinion mining, or sentiment mining

What's the problem?

- Consider classifying a subjective text unit as either positive or negative.
 - Example: The most thoroughly joyless and inept film of the year, and one of the worst of the decade. [Mick LaSalle, describing *Gigli*]
- Can't we just look for words like *great* or *terrible* ?
 - Yes, but ...
 - This laptop is a **great deal**.
 - A **great deal** of media attention surrounded the release of the new laptop.
 - This laptop is a **great deal** ... and I've got a nice bridge you might be interested in.

What's the problem?

- Subtlety or sarcasm:
 - 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.

Domain Adaptation

- Certain sentiment-related indicators seem domain-dependent.
 - .Read the book.: good for book reviews, bad for movie reviews
 - .Unpredictable.: good for movie plots, bad for a car's steering
[Turney '02]
- In general, sentiment classifiers (especially those created via supervised learning) have been shown to often be domain dependent
 - [Turney '02, Engström '04, Read 05, Aue & Gamon '05, Blitzer, Dredze & Pereira '07].
- But let's take a closer look at the types of problems . . .

Sentiment polarity

- Classic Sentiment polarity task from Pang and Lee:
 - Is an IMDB movie review positive or negative?
 - Data: *Polarity Data 2.0: (people indicate polarity of own review)*
 - <http://www.cs.cornell.edu/people/pabo/movie-review-data>
- **Treat as a document classification task**
 - Positive, negative, and (possibly) neutral
- Similar but different from topic-based text classification.
 - In topic-based text classification, topic words are important.
 - In sentiment classification, sentiment words are more important, e.g., great, excellent, horrible, bad, worst, etc.

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

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 .

Treat as a Classification Problem

- Tokenization
- Feature Extraction
 - The most important part!
- Classification using different classifiers
 - Naïve Bayes
 - MaxEnt
 - SVM
 - It turns out that MaxEnt and SVM are better than Naïve Bayes at some sentiment domains.

Sentiment Tokenization Issues

- For text from web, deal with HTML and XML markup
- Or Twitter mark-up (names, hash tags)
- Capitalization (preserve for

words in all caps)

- Phone numbers, dates
- Emoticons

Potts emoticons

```
[<>]?           # optional hat/brow
[:;=8]          # eyes
[\-o\*\'\ ]?   # optional nose
[\]\)\)\(\[dDpP/\:\}\{\@\|\]\] # mouth
|               ##### reverse orientation
[\]\)\)\(\[dDpP/\:\}\{\@\|\]\] # mouth
[\-o\*\'\ ]?   # optional nose
[:;=8]          # eyes
[<>]?           # optional hat/brow
```

- Useful code:
 - [Christopher Potts sentiment tokenizer](#)
 - [Brendan O'Connor twitter tokenizer](#)

Extracting Features for Sentiment Classification

- Which words to use?
 - Only adjectives
 - All words
 - All words turns out to work better, at least on this data
- Syntax is not used as often
 - Constituent or dependency parses are occasionally used
 - Particularly at phrase level to find dependencies of opinion words
 - Can be used to shift the “valence”
 - For negation, intensification and diminution
 - Very good, deeply suspicious
 - Should have been good
 - He is a great actor, *however* this performance . . .
 - » However changes the valence of great to be negative

Handling negation is important!

- How to handle negation:
 - I **didn't** like this movie vs I really like this movie
 - Pang and Lee simple approximation to 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

- Negation has both scope and focus
 - These may be represented in more complex structures
 - Details in Wilson “Fine-grained sentiment analysis”

Sentiment Lexicons

- One of the early approaches to sentiment analysis was to just count the words in each document that had either a positive or negative polarity from a (hand-built) sentiment lexicon.
 - This approach usually not very accurate on individual documents, but it's easy because doesn't need training data.
 - May be useful over aggregate collections or to show trends over time.
- Now we use either presence or frequencies of sentiment words as features of the classifier

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.

- Gives a list of words that have been judged to be weakly or strongly positive, negative or neutral in subjectivity
- Home page: http://www.cs.pitt.edu/mpqa/subj_lexicon.html
- 6885 words from 8221 lemmas
 - 2718 positive, 4912 negative
 - GNU GPL license
 - Examples:

type=weaksubj len=1 word1=abandoned pos1=adj stemmed1=n priorpolarity=negative

type=weaksubj len=1 word1=abandonment pos1=noun stemmed1=n priorpolarity=negative

type=weaksubj len=1 word1=abandon pos1=verb stemmed1=y priorpolarity=negative

type=strongsubj len=1 word1=abase pos1=verb stemmed1=y priorpolarity=negative

type=strongsubj len=1 word1=abacement pos1=anypos stemmed1=y priorpolarity=negative

type=strongsubj len=1 word1=abash pos1=verb stemmed1=y priorpolarity=negative

type=weaksubj len=1 word1=abate pos1=verb stemmed1=y priorpolarity=negative

type=strongsubj len=1 word1=absolve pos1=verb stemmed1=y priorpolarity=positive

type=strongsubj len=1 word1=absolute pos1=adj stemmed1=n priorpolarity=neutral

LIWC (Linguistic Inquiry and Word Count)

- Linguistic Inquiry and Word Count
 - Text analysis software based on dictionaries of word dimensions
 - Dimensions can be syntactic
 - Pronouns, past-tense verbs
 - Dimensions can be semantic
 - Social words, affect, cognitive mechanisms
 - Other categories
 - See <http://www.liwc.net/comparedicts.php>
- James Pennebaker, Univ. of Texas at Austin
 - <http://www.liwc.net/>
 - \$30 - \$90 fee for software (make sure to get dictionaries)
- Often used for positive and negative emotion words in opinion mining

ANEW (Affective Norms for English Words)

- Provides a set of emotional ratings for a large number of words in the English language
- Participants gave graded reactions from 1-9 on three dimensions
 - Good/bad, psychological valence
 - Active/passive, arousal valence
 - Strong/weak, dominance valence
- From the NIMH Center for the Study of Emotion and Attention at the University of Florida
 - <http://csea.phhp.ufl.edu/Media.html>
 - See also the paper by Dodds and Danforth on Happiness of Large-Scale Written Expressions
 - Free for research use

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

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

Which Sentiment Lexicon to use?

- An area of active research in the sentiment analysis community
- It is now recognized that the amount of overlap between the lexicons is small!
 - But in general, where there is overlap, the sentiment polarity of the words is in agreement, 2% or less disagreement.
 - Except for SentiWordNet, which disagrees up to 25%
 - Chris Potts seminar
- How to represent features from sentiment words still under research:
 - Frequency of all positive and all negative words
 - Presence of positive or negative words (particularly for twitter)
 - Using the sentiment words individually as features, instead of all words

Build a sentiment lexicon?

- For some domains, it has been shown that the best lexicon is one built for that domain
- Automatic lexicon building
 - bootstrapping
 - Identify a number of seed words of positive and negative polarity
 - Search for text involving those words that also have connecting words, such as “and”
 - Other words that occur with the connecting word are added to the lexicon with the appropriate polarity
 - Trained from annotated text by associating words with the sentiment labels that they occur with
 - Using Mutual Information scores or other measures

Other Tasks: Subjectivity Detection

- For many applications, first decide if the document contains subjective information or which parts are subjective
 - Focus of TREC 2006 Blog track
 - Label sentences or documents as to whether they are objective or subjective

Opinion Mining

- The **third level of sentiment analysis** is sometimes called opinion mining because you are finding sentiment towards aspects or attributes
- Businesses spend a huge amount of money to find consumer sentiments and opinions.
 - Consultants, surveys and focused groups, etc
 - Text in the form of transcripts of interviews or survey responses
- Opinions also available on the web
 - product reviews
 - blogs, discussion groups

Sentence Level Detection

- Sentence level or sub-sentence level detection of subjectivity
 - Wiebe, many projects
 - Pang and Lee – for movie reviews, first determine which sentences express opinions and then label for opinion polarity
- Clause level opinion strength
 - Wilson, “How mad are you?”
- Important for finding aspects or attributes
 - *The food was great but the service was awful.*

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

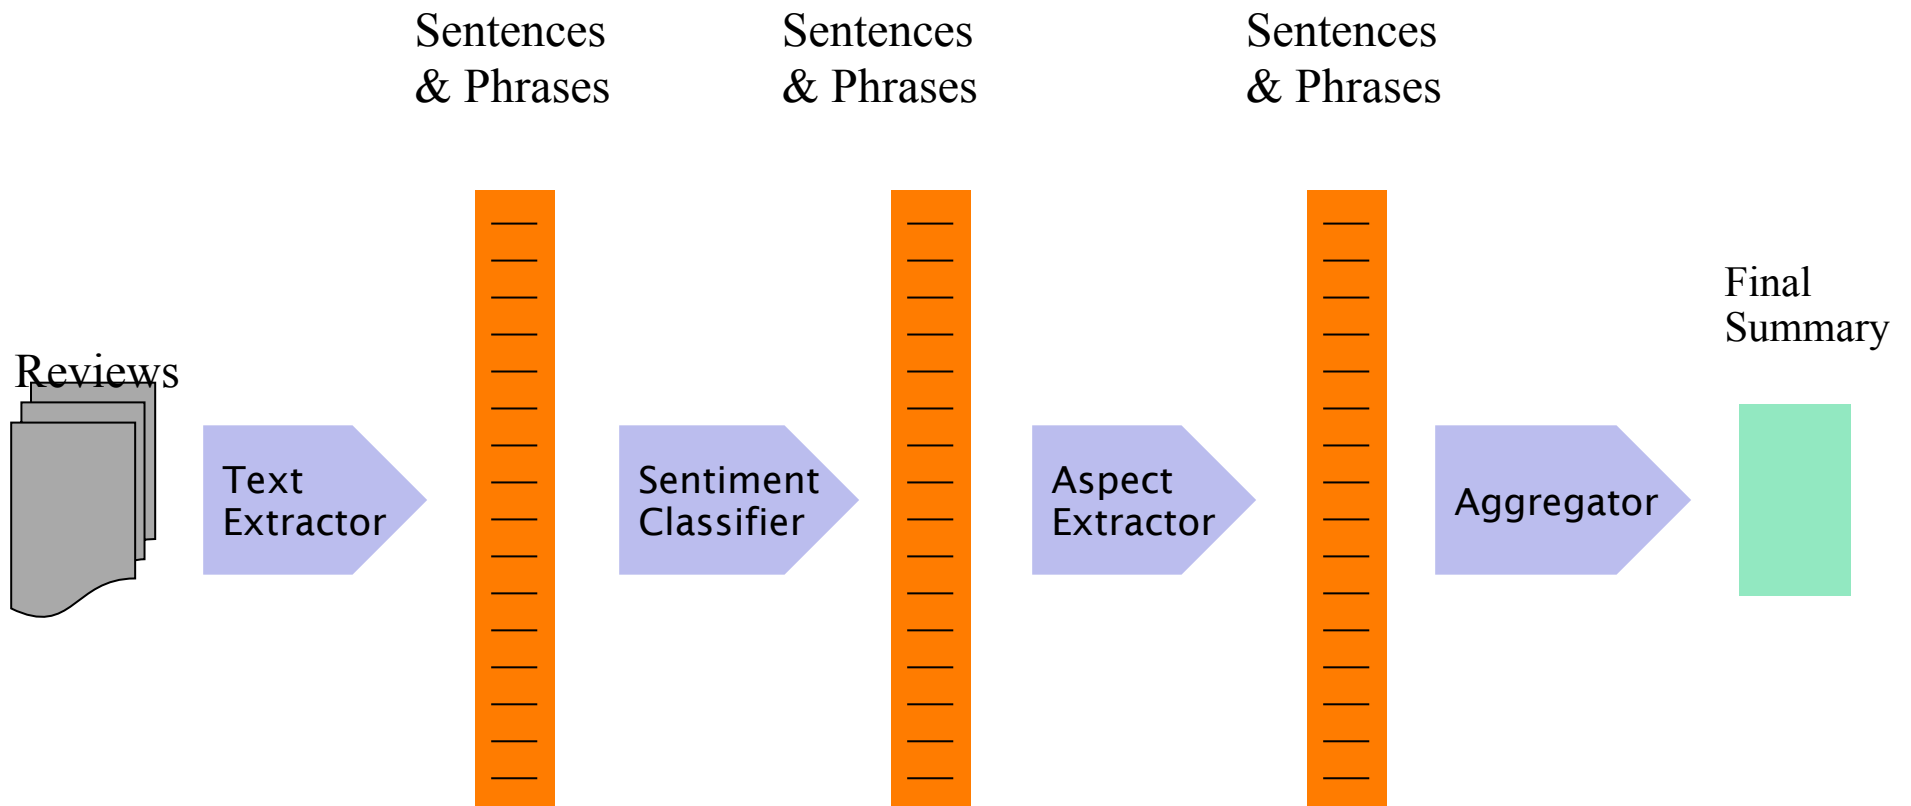
Casino	casino, buffet, pool, resort, beds
Children’s Barber	haircut, job, experience, kids
Greek Restaurant	food, wine, service, appetizer, lamb
Department Store	selection, department, sales, shop, clothing

Finding aspect/attribute/target of sentiment

- The aspect name may not be in the sentence
- Other methods to find aspects:
 - 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/opinion 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



Joint Topic/Sentiment Analysis

- An alternative approach to first finding the aspect or attribute and then the opinion or sentiment is to find them both in the same classification
 - Comparative studies of related products
 - Topics that have various features and attributes
 - Consumers
 - Political areas

Example results for opinion of aspect

Rooms (3/5 stars, 41 comments)

- (+) The room was clean and everything worked fine – even the water pressure ...
- (+) We went because of the free room and was pleasantly pleased ...
- (-) ...the worst hotel I had ever stayed at ...

Service (3/5 stars, 31 comments)

- (+) Upon checking out another couple was checking early due to a problem ...
- (+) Every single hotel staff member treated us great and answered every ...
- (-) The food is cold and the service gives new meaning to SLOW.

Dining (3/5 stars, 18 comments)

- (+) our favorite place to stay in biloxi.the food is great also the service ...
- (+) Offer of free buffet for joining the Play

Feature-based Summary (Hu and Liu, KDD-04)

- From reviews, extract a summary:

GREAT Camera., Jun 3, 2004

Reviewer: **jprice174** from Atlanta, Ga.

I did a lot of research last year before I bought this camera... It kinda hurt to leave behind my beloved nikon 35mm SLR, but I was going to Italy, and I needed something smaller, and digital.

The **pictures** coming out of this camera are amazing. The **'auto'** feature takes great pictures most of the time. And with digital, you're not wasting film if the picture doesn't come out. ...

....

Feature Based Summary:

Feature1: **picture**

Positive: 12

- The **pictures** coming out of this camera are amazing.
- Overall this is a good camera with a really good **picture** clarity.

...

Negative: 2

- The **pictures** come out hazy if your hands shake even for a moment during the entire process of taking a picture.
- Focusing on a display rack about 20 feet away in a brightly lit room during day time, **pictures** produced by this camera were blurry and in a shade of orange.

Feature2: **battery life**

...

How to deal with 7 stars?

Bo Pang and Lillian Lee. 2005. Seeing stars: Exploiting class relationships for sentiment categorization with respect to rating scales. *ACL*, 115–124

- The second level of sentiment analysis deals is a similar classification task, but needs to find levels of strength
 1. Map to binary, or
 2. Use linear or ordinal regression
 - Or specialized models like metric labeling

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:**
 - Detection of flirtation or friendliness in conversations
- **Personality traits:**
 - Detection of extroverts
- **Certainty:**
 - How certain the holder is of the opinion or sentiment

Detection of Friendliness

Ranganath, Jurafsky, McFarland

- Friendly speakers use collaborative conversational style
 - Laughter
 - Less use of negative emotional words
 - More sympathy
 - That's too bad I'm sorry to hear that
 - More agreement
 - I think so too
 - Less hedges
 - kind of sort of a little ...

Viewpoints and Perspectives

- In some types of documents, the authors are not necessarily discussing opinions on particular topics, but are revealing general attitudes or sometimes a set of bundled attitudes and beliefs
 - Classifying political blogs as liberal, conservative, libertarian, etc.
 - Identifying Israeli vs. Palestinian viewpoints
- One type of this is Multi-perspective Question Answering
 - On next slide . . .

MPQA

- Multi-Perspective Question Answering
 - *What does Bush think about Hillary Clinton?*
 - *How does the US regard the latest terrorist attacks in Baghdad?*
- Sentence, or part of a sentence, that answers the question:
 - *“How does X feel about Y?”*
 - *“It makes the system more flexible,” argues a Japanese businessman.*
- Looking for opinion linked to opinion-holder

Stoyanov, Cardie, Wiebe, & Litman,

Evaluating an Opinion Annotation Scheme Using a Multi-Perspective Question and Answer Corpus. 2004 AAI Spring Symposium on Exploring Attitude and Affect in Text,

Stance and Argumentation

- Some forms of online discourse takes the form of trying to argue a viewpoint or opinion, or taking a stance in a particular debate
 - Ideological Debates
 - Somasundaram and Wiebe – look at argumentation
 - Abbot, Walker, et al – classifying stance in on-line debates
 - “Cats rule, dogs drool!” is much easier to classify than debates on abortion, religion, politics

Certainty Recognition

- *Certainty*
 - the quality / state of being free from doubt, especially on the basis of evidence
- Related work:
 - Types of subjectivity (Liddy et al. 1993; Wiebe 1994, 2000; Wiebe et al. 2001)
 - Adverbs and modality (Hoye, 1997)
 - Hedging in different kinds of discourse
 - Expressions of (un)certainty in English (from applied linguistics)
- Goal – characterize ‘certainty’ of textual statements

Summary on Sentiment

- Understanding semantics of less factual aspects of text
- Generally modeled as classification or regression task
 - predict a binary label for polarity
 - or predict an ordinal label for the level of sentiment
- Features:
 - Negation is important
 - Using all (filtered) words 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