
Subjective Analysis of Text: Sentiment Analysis

(credits to Liz Liddy and Lillian Lee)

Opinion Analysis

(credits to Bing Liu)

Certainty

(credits to Vicky Rubin and Liz Liddy)

Multi-Perspective Question Answering

(credits to Jan Wiebe)

Terminology

- 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

Elusive Aspects of Text Semantics

- In addition to representing documents, email, blogs, etc or answering questions on just the basis of thematic content
- Recognition of more subtle aspects of what is being conveyed in language
- Includes affective, emotive, opinion, certainty & evaluative aspects of meaning

Task Description

- Simplest level - Measuring **polarity** of text
 - Negative / positive attitude of reporter / blogger
 - Favorable / unfavorable review of a product
 - Right / left political leaning of speaker
 - Certainty / uncertainty about what's reported
- Huge amounts of text available
 - Blogs
 - Message boards
 - Discussion groups
 - eCommerce product sites
 - Email

What Could Be Done Now

- Business
 - Gauge reactions to new products
 - Understand which features of products have emotion-invoking affects on customers
 - Compare to competitors' products
 - Contribute to company's reputation management
- Consumers
 - Summarize key pros and cons in product reviews
- Government
 - Track attitudes towards government policies
 - Understand trends in the public's views
 - Gauge public reaction to campaign ads
 - Predict election outcomes

What is Possible Now [cont'd]

- Learn '*the buzz*' on the street
 - Extracting market sentiment from stock message boards to predict impact on stock price
- Identify financial scams
- Estimate political orientation of documents / sites / authors / blogs
- Understand the nature of relationships between cited and citing documents
- Discern level of certainty about events / statements

Sentiment classification

- Classify documents (e.g., reviews) based on the overall sentiments expressed by authors,
 - 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.

Easy case: sentiment polarity

- 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*]
- One application: summarizing reviews (thumbs up or thumbs down)
- Can't we just look for words like .great. or .terrible. ?
 - Yes, but ...
 - ... learning a sufficient set of such words or phrases is an active challenge.
 - [Hatzivassiloglou&McKeown '97, Turney '02, Wiebe et al. '04, and more than a dozen others, at least]

One experiment in creating polarity words

- Human 1
 - Positive: dazzling, brilliant, phenomenal, excellent, fantastic
 - Negative: suck, terrible, awful, unwatchable, hideous
 - 58% (on movie reviews)
- Human 2
 - Positive: gripping, mesmerizing, riveting, spectacular, cool, awesome, thrilling, badass, excellent, moving, exciting ,
 - Negative: bad, cliched, sucks, boring, stupid, slow
 - 64%
- Statistics-based
 - Positive: love, wonderful, best, great, superb, beautiful, still
 - Negative: bad, worst, stupid, waste, boring, ?, !
 - 69%

Issues

- 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.
 - This film should be brilliant. It sounds like a great plot, the actors are rst grade, and the supporting cast is good as well, and Stallone is attempting to deliver a good performance. However, it can't hold up.

Domain Adaptation

- Certain sentiment-related indicators seem domain-dependent.
 - .Read the book..
 - .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].
- This holds even though the set of categories (e.g., 1-5 stars) is generally considered constant across domains.

Additional Ideas for Techniques

- Relationships between items can be a rich source of information about for performing classification on the items.
 - Nearby sentences can share the same subjectivity status, subjective or objective [Pang&Lee '04]
- Mentions separated by .and. usually have similar sentiment labels; those separated by .but. usually have contrasting labels
 - [Popescu&Etzioni '05, Snyder&Barzilay '07];
 - similar reasoning holds for synonyms and antonyms [Hu&Liu '04]
- In some domains, references to other speakers generally indicate disagreement
 - [Agrawal et al '03, Mullen&Malouf '06, Goldberg, Zhu & Wright '07] (cf. Adamic&Glance ['05])

Finding polarity (aka paradigm) words

- SO-PMI - Turney & Littman (2003)
- Algorithm that estimates the Semantic Orientation (+ or -) of a given word
 - Evaluative character and degree of a word
- Infers SO of a word from its statistical association with a set of positive and negative paradigm terms
- Measures of Association
 - Pointwise Mutual Information (PMI)
 - Latent Semantic Analysis (LSA)
 - Semantic distance in WordNet

Polarity or Paradigm Words

- Studies use wide range of # of paradigm words
 - 1,300 for Hatzivassiloglou & McKeown, 1997
 - 14 for Turney & Littman, 2003
 - + *good, nice, excellent, positive, fortunate, correct, superior*
 - *bad, nasty, poor, negative, unfortunate, wrong, inferior*
 - Applications may retrain with specialized words
- Humans - high agreement on SO of words
 - 97% agreement amongst 4 subjects (Hatzivassiloglou & McKeown, 1997)

Classifying Movie Reviews

- Processing of new reviews
 - POS tagging
 - Phrase recognition with ADJs or ADVs
 - Calculate SO-PMI of each phrase
 - Compute average SO of extracted phrases
 - Assign + or – to the review
- Typical performance is in range of 80 – 95%
- Challenge
 - Positive reviews may have words with negative SO, although subjective evaluation is positive and vice versa

Opinion Mining

- 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

Two forms of opinions

- **Direct Opinions:** sentiment expressions on some objects, e.g., products, events, topics, persons
 - E.g., “the picture quality of this camera is great”
 - Subjective
- **Comparisons:** relations expressing similarities or differences of more than one object. Usually expressing an ordering.
 - E.g., “car x is cheaper than car y.”
 - Objective or subjective

Opinion mining tasks

- At the document (or review) level:
 - **Task:** sentiment classification of reviews
 - Classes: positive, negative, and neutral
 - **Assumption:** each document (or review) focuses on a single object O (not true in many discussion posts) and contains opinion from a single opinion holder.
- At the sentence level:
 - **Task 1:** identifying subjective/opinionated sentences
 - Classes: objective and subjective (opinionated)
 - **Task 2:** sentiment classification of sentences
 - **Classes:** positive, negative and neutral.
 - **Assumption:** a sentence contains only one opinion
 - not true in many cases.
 - Then we can also consider clauses.

Opinion mining tasks (contd)

- At the feature level:
 - Task 1*: Identifying and extracting object features that have been commented on in each review.
 - Task 2*: Determining whether the opinions on the features are positive, negative or neutral in the review.
 - Task 3*: Grouping feature synonyms.
 - Produce a feature-based opinion summary of multiple reviews (**more on this later**).
- **Opinion holders**: identify holders is also useful, e.g., in news articles, etc, but they are usually known in user generated content, i.e., the authors of the posts.

Let us go further?

- Sentiment classifications at both document and sentence (or clause) level are useful, **but**
 - They do not find what the opinion holder liked and disliked.
- An negative sentiment on an object
 - does not mean that the opinion holder dislikes everything about the object.
- A positive sentiment on an object
 - does not mean that the opinion holder likes everything about the object.
- **We need to go to the feature level.**

But before we go further

- **Opinion Words or Phrases** (also called polar words, opinion bearing words, etc). E.g.,
 - **Positive**: beautiful, wonderful, good, amazing,
 - **Negative**: bad, poor, terrible, cost someone an arm and a leg (idiom).
- They are instrumental for opinion mining (obviously)
- Three main ways to compile such a list:
 - **Manual approach**: not a bad idea, only an one- time effort
 - **Corpus-based approaches**
 - **Dictionary-based approaches**
- **Important to note:**
 - **Some opinion words are context independent.**
 - **Some are context dependent.**

Feature-based opinion mining and summarization (Hu and Liu, KDD-04)

- Again focus on reviews (easier to work in a concrete domain!)
- Objective: find what reviewers (opinion holders) liked and disliked
 - Product features and opinions on the features
- Since the number of reviews on an object can be large, an opinion summary should be produced.
 - Desirable to be a structured summary.
 - Easy to visualize and to compare.
 - Analogous to multi-document summarization.

Different review format

Format 1 - Pros, Cons and detailed review: The reviewer is asked to describe Pros and Cons separately and also write a detailed review. [Epinions.com](#) uses this format.

Format 2 - Pros and Cons: The reviewer is asked to describe Pros and Cons separately. [C|net.com](#) used to use this format.

Format 3 - free format: The reviewer can write freely, i.e., no separation of Pros and Cons. [Amazon.com](#) uses this format.

Format 1

My SLR is on the shelf

by [camerafun4](#). Aug 09 '04

Pros: Great photos, easy to use, very small

Cons: Battery usage; included memory is stingy.

I had never used a digital camera prior to purchasing have always used a SLR ... [Read the full review](#)

Format 3

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.

Format 2



out of 10

"It is a great digitbal still camera for this century"

September 1, 2004

Pros:

It's small in size, and the rotatable lens is great. It's very easy to use, and has fast response from the shutter. The LCD has increased from 1.5 in to 1.8, which gives bigger view. It has lots of modes to choose from in order to take better pictures.

Cons:

It almost has no cons, it would be better if the LCD is bigger and it's going to be best if the model is designed to a smaller size.

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

Feature Based 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. ...

....

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

...

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

Four-Dimensional Relational Model for Certainty Categorization

D1: PERSPECTIVE	D2: FOCUS	D3: TIMELINE	D4: LEVEL
<p>Writer's Point of View</p>	<p>Abstract Information (e.g. opinions, judgments, attitudes, beliefs, emotions, assessments, predictions)</p>	<p>Past Time (i.e. completed, recent in the past)</p>	<p>Absolute</p>
<p>Reported Point of View</p> <p>Directly involved 3rd parties (e.g. witnesses, victims)</p> <p>Indirectly involved 3rd parties (e.g. experts, authorities)</p>		<p>Factual Information (e.g. concrete facts, events, states)</p>	<p>Present Time (i.e. immediate, current, incomplete, habitual)</p>
		<p>Future Time (i.e. predicted, scheduled)</p>	<p>Moderate</p>
			<p>Low</p>

Rubin, Kando & Liddy. Certainty Categorization Model. AAAI-EAAT Symposium, 2004.

Dimension 1: Perspective

D1: PERSPECTIVE

D2: FOCUS

D3: TIMELINE

D4: LEVEL

Writer's Point of View

Reported Point of View

Directly involved 3rd parties (e.g. witnesses, victims)

Indirectly involved 3rd parties (e.g. experts, authorities)

- point of view, voice, or experiencer of certainty

- **the writer** is the author of the article

*More evenhanded coverage of the presidential race would help enhance the legitimacy of the eventual winner, which **now appears likely to be Putin.***
(ID=e8.14)

- **people or organizations, direct participants**

*The Dutch recruited settlers with an advertisement that **promised** to provide them with slaves who "**would accomplish** more work for their masters, ...*
(ID=e27.13)

- **tangentially related to the event in the professional or other capacities**

*The historian Ira Berlin, author of "Many Thousands Gone," **estimates** that one slave perished for every one who survived capture in the African interior... (ID=e27.8)*

Dimension 2: Focus

D1: PERSPECTIVE

D2: FOCUS

D3: TIMELINE

D4: LEVEL

An idea that does not represent an external reality but rather a hypothesized world, existing in the mind, separated from embodiment or object of nature.

*In Iraq, the first steps **must be taken** to put a hard-won new security council resolution on arms inspections into effect. (ID=e8.12)*

Abstract Information
(e.g. opinions, judgments, attitudes, beliefs, emotions, assessments, predictions)

Factual Information
(e.g. concrete facts, events, states)

Based on, characterized by, or contains facts, i.e. has actual existence in the world of events.

*The settlement **may not fully compensate** survivors for the delay in justice, ... (ID=e14.19)*

Dimension 3: Timeline

D1: PERSPECTIVE

D2: FOCUS

D3: TIMELINE

D4: LEVEL

- accounts for relevance of time to the moment when the article was written
- the past includes completed or recent states or events;
- the present is current, immediate, and incomplete states of affairs;
- the future is predictions, plans, warnings, and suggested actions.

Past Time
(i.e. completed, recent in the past)

Present Time
(i.e. immediate, current, incomplete, habitual)

Future Time
(i.e. predicted, scheduled)

→ *The failure lasted only about 30 minutes and had no operational effect, the FAA said, adding that **it was not even clear** that the problem was caused by the date change.*

(ID=n4.19)

Dimension 4: Level

D1: PERSPECTIVE

D2: FOCUS

D3: TIMELINE

D4: LEVEL

Eventually, however, auditors will almost certainly have to form a tough self-regulatory body that can oversee its members' actions... (ID=e24.18)

... but clearly an opportunity is at hand for the rest of the world to pressure both sides to devise a lasting peace based on democratic values and respect for human rights. (ID=e22.6)

That fear now seems exaggerated, but it was not entirely fanciful. (ID=e4.8)

So far the presidential candidates are more interested in talking about what a surplus might buy than in the painful choices that lie ahead. (ID=e3.7)

Absolute

High

Moderate

Low

- currently, a 4-way distinction
- only sentences with explicit indication of certainty are in scope
 - low certainty and uncertainty are lumped together

Potential Applications

- Alerting intelligence analysts to level above or below normal and associating certainty with its source
- Searching by level and point of view parameter
 - *What does Pres. Bush sound most certain about in his speeches?*
- Ordering retrieved information by certainty of authors or author's reports of certainty of others
 - Decreases amount of uncertain information presented
 - Prioritizes sources that provide highly certain information
- Summarizing per document, across documents, per topic
- Inferring true state of affairs based on high level certainty statements from multiple sources

Generally, 'Affective' Next Steps

- Currently, seem to do well on extremes
 - I.e. movie & product reviews,
 - *But isn't that the easy stuff?*
 - *What about the more subtle challenges?*
- Recognize that affect is not just lexical
 - Need represent propositions, not just words
 - Nuances of syntactic & discourse level features
- Affect may vary by genre, register, and / or audience
- If web-based, combine with link information

Expanding Evidence Sources

- Wider range of lexical items as an indicator of attitude of the writer towards events and situations described in a text
 - verbs: *delight, approve, ensure*
 - nouns: *favor, credit, freedom*
 - adjectives: *brave, bright, good*
 - adverbs: *bravely*
- Base attitudinal valence of a lexical item can be modified by context
 - Sentence level
 - Discourse level

From Reviews to Opinions

- Automatic recognition of opinion-indicating features of text
- Frequently in form of indirect statements
 - Non-Opinion
 - *Investment banking revenue fell 32% to \$205 million, in part reflecting the continued slowdown of the underwriting business.*
 - Opinion-Proposition
 - *The U.S. government in recent years has accused overseas companies of [PROPOSITION excessively slashing prices on semiconductors and supercomputers -- products Fujitsu and NEC make].*
 - Opinion-Reference
 - *But for the most part, officials have avoided expressing [PROPOSITION those views] in public, fearing they would unnecessarily antagonize the Fed.*

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 AAIL Spring Symposium on Exploring Attitude and Affect in Text,

Syntax & Discourse Matter

- Negation
 - John is **clever**. - John is *not* **clever**.
- Modals
 - The film is **brilliant**. - The film *should* be **brilliant**.
- Intensifiers
 - They are **suspicious**. - They are *deeply* **suspicious**.
- Presuppositions
 - He **got into Harvard**. - He *barely* **got into Harvard**.
 - He *even* **got into Harvard**.
- Discourse connectors
 - *Although* Boris is **brilliant** at math, he is a **horrible** teacher.

Example of Valence Shifting

The film should be brilliant $+ \Rightarrow 0$ *brilliant* within scope of *should*

The characters are appealing

Stallone plays a happy man $+ \Rightarrow 0$ *appealing* under scope of *characters*

It sounds like a great story,

however, as a movie it $+ \Rightarrow 0$ *happy* part of story world

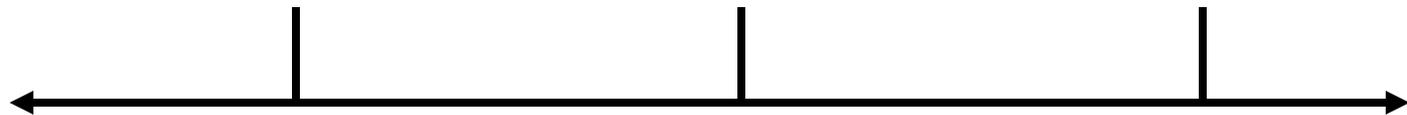
is a failure $+ \Rightarrow 0$ *great* within the scope of *sounds like*

$+ \Rightarrow 0$ *however* reverses the “+”

$+ \Rightarrow -$ valence of *great*

Estimating Political Orientation

- Left- or right- leaning of Web documents
 - Particularly salient dimension of author and reader opinion in weblogs
 - Yahoo's Politics blog category is largest of 18
 - Method is applicable to any topic/domain where opinion is polarized
- Uses PMI-IR to estimate strength of association between document & left- / right- groups



Estimating Political Orientation

- Uses PMI-IR at the document level
- Established 19 ‘exemplar’ documents each for left- and right- political orientations
- Then evaluates new document as to whether
 - It contains words from exemplar documents
 - Documents it links to contain words from exemplar documents
- Tested on 695 web documents - 94% correct
- Link structure adds useful new evidence to lexical features

Future Applications of Interest

- Understanding ‘man on the street’ views
 - Business considers web chats / blogs as valuable sources of insight about attitudes & opinions
 - **But**, want to know demographics of customers
- Given huge number of human reviews, want:
 - Reporting
 - Analysis
 - Prediction
- Need to have business cases that prove that ‘affect mining’ provides added value

Sources for Further Information

- AAAI – 2004 Spring Symposium Series on *Exploring Attitude & Affect in Text: Theories and Applications*
 - Final Proceedings available from AAAI
 - Book on topic
 - Edited by Janyce Wiebe, Jimi Shanahan & Yan Qu Computing Attitude and Affect in Text. Springer, Dordrecht, The Netherlands.
- *Proceedings of the Workshop on Sentiment and Subjectivity in Text* at the joint COLING-ACL conference. Sydney, Australia, July 2006