Opinion Mining

Andrea Esuli

andrea.esuli@isti.cnr.it
Istituto di Scienza e Tecnologie dell'Informazione
Consiglio Nazionale delle Ricerche, Pisa, Italy

Language and Intelligence Reading Group,
June 14, 2006, Pisa, Italy.

Outline

Outline

Opinion mining

Opinion mining (OM) is a recent discipline at the crossroads of information retrieval and computational linguistics which is concerned not with the topic a document is about, but with the opinion it expresses.

What is an opinion?

Private state – a state that is not open to objective observation or verification [Quirk et al., 1985]

Sentiment Analysis, *Sentiment Classification*, *Opinion Extraction* are other names used in literature to identify this discipline.

Example (OM problems)

- What is the general opinion on the proposed tax reform?
- How is popular opinion on the presidential candidates evolving?
- Which of our customers are unsatisfied? Why?

OM research topics

Three main OM topics:

- Development of linguistic resources for OM, e.g. automatically build a *lexicons* of *subjective terms*.
- Classification of text (entire documents, sentences) by their opinion content, e.g. classify a movie review either as Positive or Negative.
- Extraction of opinion expression from text, eventually including relations with the rest of content, e.g. recognize an opinion, who is expressing it, who/what is the target of the opinion.

Outline

Linguistic resources for OM

A linguistic resource for OM defines some sentiment-related properties of terms.
Research work on this topic deal with three main tasks:

- Determining term orientation, as in deciding if a given Subjective term has a Positive or a Negative slant.
- Determining term subjectivity, as in deciding whether a given term has a Subjective or an Objective (i.e. neutral, or factual) nature.
- Determining the strength of term attitude (either orientation or subjectivity), as in attributing to terms (real-valued) degrees of positivity or negativity.

Example

`good,excellent,best` – positive terms

`bad,wrong,worst` – negative terms

`vertical,yellow,liquid` – objective terms

Linguistic resources for OM

Not only terms:

- Tackling previous tasks for term senses, thus taking into account the fact that different senses of the same ambiguous term may have different sentiment-related properties.
- Tackling previous tasks for multi-word expressions.

Example

`estimable` – ambiguous term with an objective sense (i.e. *measurable*), and a positive sense (i.e. *deserving respect*).

`not entirely satisfactory` – negative expression

Outline

Orientation of terms

The problem:

Determining if a subjective term has a Positive or a Negative orientation.

[Hatzivassiloglou and McKeown, 1997]

Hypothesis: *adjectives* in `and` conjunctions usually have similar orientation, though `but` is used with opposite orientation.

Example (conjuction of adjectives)

1. `The tax proposal was` *`simple and well received...`*
2. `The tax proposal was` *`simplistic but well received...`*
3. `* The tax proposal was` *`simplistic and well received...`*

Method: a weighted graph of similarity of orientation is defined by analyzing conjunctions of adjectives in unprocessed text, then a minimum-cut method is applied to the graph.

Orientation of terms

[Turney and Littman, 2003]

Hypothesis: terms with similar orientation tend to co-occur in documents.

The *Semantic Orientation* (SO) of a term is estimated by combining a pointwise mutual information (PMI) measure of the term against some *paradigmatic* terms.

Pos = {`good`, `nice`, `excellent`, `positive`, `fortunate`, `correct`, `superior`}
Neg = {`bad`, `nasty`, `poor`, `negative`, `unfortunate`, `wrong`, `inferior`}

PMI is measured using the number of results returned by the AltaVista search engine.

$$
\begin{aligned}
PMI(t, t_i) &= \log \frac{\#(\text{“}t \ \texttt{NEAR} \ t_i\text{”})}{\#(\text{“}t\text{”})\#(\text{“}t_i\text{”})} \\
SO(t) &= \Sigma_{t_i \in Pos} PMI(t, t_i) - \Sigma_{t_i \in Neg} PMI(t, t_i)
\end{aligned}
$$

Orientation of terms

[Esuli and Sebastiani, 2005]

Hypothesis: terms with similar orientation have similar glosses.

Example (glosses for terms with similar orientation)

`good`: "*that which is pleasing or valuable or useful*";
"*agreeable or pleasing*".

`beautiful`: "*aesthetically pleasing*".

`pretty`: "*pleasing by delicacy or grace; not imposing*".

Each term is represented by its gloss.

A binary classifier is learned, in a semi-supervised process, using the glosses of the Positive and Negative terms in the training set.

Orientation of terms

A semi-supervised learning method to determine semantic orientation of terms:

- The training set is built by iteratively adding to it *synonyms* and *antonyms* of terms already belonging to it, starting from two small seed sets L_p and L_n of known *Positive* and *Negative* terms.
- A classifier is learned on the glosses of terms in training set and then applied to the glosses of terms in test set.

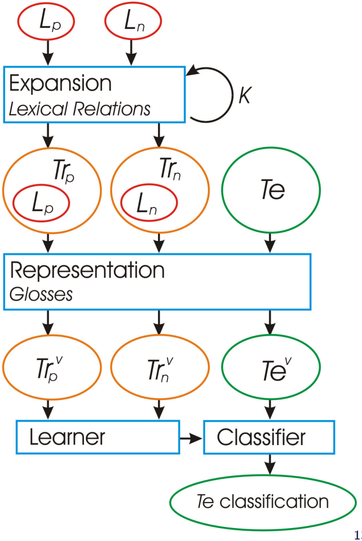

Orientation of terms

Test sets:

HM: 657 Positive / 679 Negative hand labeled adjectives, defined in [Hatzivassiloglou and McKeown, 1997].

TL: 1,614/1,982 terms extracted from the General Inquirer (GI) lexicon.

Results:

Test set	Method	Accuracy(%)
HM	[Hatzivassiloglou and McKeown, 1997]	78.08
	[Turney and Littman, 2003] AV-NEAR	87.13
	[Turney and Littman, 2003] 7M-NEAR	80.31
	[Esuli and Sebastiani, 2005]	**87.38**
TL	[Turney and Littman, 2003] AV-NEAR	82.84
	[Turney and Littman, 2003] 7M-NEAR	76.06
	[Turney and Littman, 2003] AV-AND	67.00
	[Esuli and Sebastiani, 2005]	**83.09**

Outline

Subjectivity of terms

The problem:

Determining if a term expresses subjectivity (Subjective) or not (Objective).

[Baroni and Vegnaduzzo, 2004]

Turney's PMI method is used to identify subjective adjectives.

Seed terms set is composed by 35 subjective terms, selected by human judges.

The method produces a ranking by subjectivity of the 3,047 test terms (972 Subjective, 31.9%).

Results: Precision/Recall table (AltaVista with NEAR operator).

recall	precision	recall	precision
.100	.882	.700	.604
.300	.768	.900	.476
.500	.710	1	.319

Subjectivity of terms

[Esuli and Sebastiani, 2006a]

The method of [Esuli and Sebastiani, 2005] is adapted to classify terms as either Positive, Negative or Objective.

Hypothesis:

- (from previous work) terms with similar orientation have similar glosses.
- terms without orientation have *non-oriented* glosses.

Example

`yellow`: "*similar to the color of an egg yolk*".

`vertical`: "*at right angles to the plane of the horizon or a base line*".

Test set: the whole GI lexicon (1,614 Pos/1,982 Neg/5,009 Obj).

Results: 67.6% accuracy on classification on Subjective vs Objective, 66.0% on classification on the three categories.

Outline

Subjectivity and orientation of term senses

[Esuli and Sebastiani, 2006b]

Previous experiences on terms showed that:

- Variation in the parameters of the classifiers do not affect accuracy but *distribution of terms among categories*.
- "Diffult" terms are those that have multiple senses with different sentiment properties (e.g. `bright`, `high`).

The method of [Esuli and Sebastiani, 2006a] has been adapted to classify each synset of WordNet, using various configuration of the classifier.

SentiWordNet is a lexical resource that assigns to each synset of WordNet three sentiment scores: positivity, negativity, objectivity.

The sum of the scores for a synset is always one.

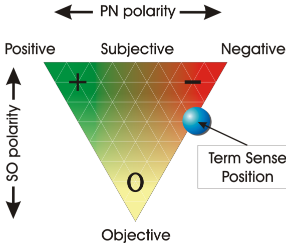

[LREC '06] SentiWordNet interface

estimable | Search word | show position

Adjective

3 senses found.

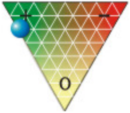 P = 0.75, N = 0, O = 0.25	estimable(1) *deserving of respect or high regard*
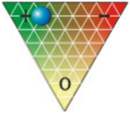 P = 0.625, N = 0.25, O = 0.125	honorable(5) good(4) respectable(2) estimable(2) *deserving of esteem and respect; "all respectable companies give guarantees"; "ruined the family's good name"*
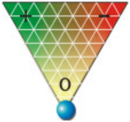 P = 0, N = 0, O = 1	computable(1) estimable(3) *may be computed or estimated; "a calculable risk"; "computable odds"; "estimable assets"*

main page

Outline

The Appraisal theory

[Martin and White, 2005] – The Appraisal theory.

Appraisal theory is a framework of linguistic resources which describe how writers and speakers express inter-subjective and ideological positions.

- *appraisal*
  - *attitude*
    - *affect*
    - *appreciation*
    - *judgement*
  - *graduation*
    - *force*
    - *focus*
  - *orientation*
    - *positive*
    - *negative*
  - *polarity*
    - *marked*
    - *unmarked*

	`happy`	`very`	“`very happy`”	`not`	“`not very happy`”
attitude:	affect	–	affect	–	affect
orientation:	positive	–	positive	**negate**	**negative**
force:	neutral	**increase**	**high**	**reverse**	**low**
focus:	neutral	–	neutral	–	neutral
polarity:	unmarked	–	unmarked	**marked**	**marked**

[Whitelaw et al., 2005] semi-automatically have produced a lexicon of 1,329 appraisal entities from 400 seed terms, in around twenty man-hours.

Outline

Sentiment classification

The problem:

Determining the *overall* sentiment properties of a text.

Applications:

- Split *reviews* of a movie into the sets "thumbs up" and "thumbs down".
- Alert a customer service for very dissatisfied customers.
- When searching for opinions on a product on the web, filter search results to obtain only Subjective web pages.
- Monitor bloggers mood trend of along time.

Sentiment classification

[Turney, 2002]

Hypothesis: the orientation of the whole document is the sum of the orientation of all its parts.

The PMI method has been applied to classify reviews as either Positive or Negative.

The SO of a reviews is computed as the average of the SO of adjectives and adverbs contained in the review.

The average accuracy on 410 reviews is 74%, ranging from 84% for automobile reviews to 66% for movie reviews.

Sentiment classification

[Pang et al., 2002]

Application some standard supervised automatic text classification methods to the problem to classify orientation of movie reviews.

Learners:	Naïve Bayes, MaxEnt, SVM.
Features:	unigrams, bigrams, adjectives, POS, position.
Preprocessing:	negation propagation.
Representation	binary, frequency.

82.9% accuracy, on a 10-fold cross validation experiments on 1,400 movie reviews (SVM, unigrams, binary).

In *[Pang and Lee, 2004]* a sentence subjectivity classifier is applied, as preprocessing, to reviews, to filter out Objective sentences. Accuracy on movie reviews classification raises to 86.4%.

[Whitelaw et al., 2005] added appraisal features to the Movie Review corpus and obtained a 90.2% classification accuracy.

Outline

Analysis of opinions in text

The problem:

Identify expression of opinions in text, and eventually:

- their sentiment properties (e.g. orientation, strenght);
- who is expressing them;
- their target.

Example

- `I'm` *`not very happy`* `with this car.`
- `The` *`CEO of XX`* `said that XX stocks are` *`healty`*`.`
- *`Market analysts`* `said that XX stocks are` *`rubbish`*`.`

Analysis of opinions in text

[Wiebe et al., 2005] – An annotation scheme for Multi-Perspective Question Answering (MPQA).

Example

`What are the predictions on XX's stocks?`

- A fine-grained annotation scheme, annotating text at the word- and phrase-level.
- For every expression of a private state in each sentence, a private state frame is defined (frames can be nested).
- A private state frame includes the source of the private state (i.e., whose private state is being expressed), the target (i.e., what the private state is about), and various properties involving intensity, significance, and type of attitude.

Opinion annotation

```
''The report is full of absurdities'', Xirao-Nima said.
```

Objective speech event:	
Text anchor:	the entire sentence
Source:	<writer>
Direct subjective:	
Text anchor:	said
Source:	<writer,Xirao-Nima>
Intensity:	high
Expression intensity:	neutral
Target:	report
Attitude type:	negative
Expressive subjective element:	
Text anchor:	full of absurdities
Source:	<writer, Xirao-Nima>
Intensity:	high
Attitude type:	negative

The results of this work is the “MPQA Corpus of Opinion Annotations”, which contains 535 news articles (11,114 sentences) manually annotated (Version 1.2).

Outline

An emerging discipline

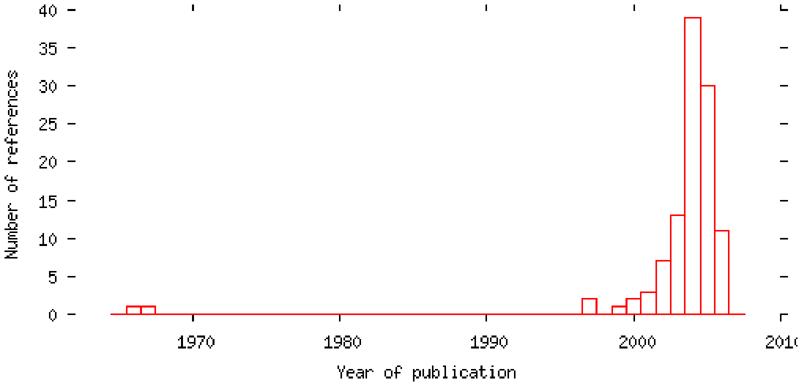

Resources

- The Sentiment Bibliography
  http://liinwww.ira.uka.de/bibliography/Misc/Sentiment.html
- The Sentiment & Affect Yahoo! Group
  http://groups.yahoo.com/group/SentimentAI
- The General Inquirer
  http://www.wjh.harvard.edu/~inquirer
- SentiWordNet
  http://patty.isti.cnr.it/~esuli/software/SentiWordNet
- Movie Review corpus
  http://www.cs.cornell.edu/people/pabo/movie-review-data
- MPQA opinion corpus
  http://www.cs.pitt.edu/mpqa/databaserelease
- The Appraisal website
  http://grammatics.com/appraisal

Questions?


SERGIO LEONE
ELI WALLACH
CLINT EASTWOOD
LEE VAN CLEEF
THE GOOD
THE BAD
AND THE UGLY

Baroni, M. and Vegnaduzzo, S. (2004).
Identifying subjective adjectives through web-based mutual information.
In *Proceedings of KONVENS-04*, pages 17–24, Vienna, AU.

Esuli, A. and Sebastiani, F. (2005).
Determining the semantic orientation of terms through gloss analysis.
In *Proceedings of CIKM-05, the ACM SIGIR Conference on Information and Knowledge Management*, Bremen, DE.

Esuli, A. and Sebastiani, F. (2006a).
Determining term subjectivity and term orientation for opinion mining.
In *Proceedings EACL-06, the 11rd Conference of the European Chapter of the Association for Computational Linguistics*, Trento, IT.

Esuli, A. and Sebastiani, F. (2006b).
Sentiwordnet: A publicly available lexical resource for opinion mining.
In *Proceedings of LREC-06, the 5th Conference on Language Resources and Evaluation*, Genova, IT.

Hatzivassiloglou, V. and McKeown, K. R. (1997).
Predicting the semantic orientation of adjectives.
In *Proceedings of ACL-97, 35th Annual Meeting of the Association for Computational Linguistics*, pages 174–181, Madrid, ES. Association for Computational Linguistics.

Martin, J. R. and White, P. R. R. (2005).
The Language of Evaluation: Appraisal in English.

Palgrave, London, UK.

Pang, B. and Lee, L. (2004).
A sentimental education: Sentiment analysis using subjectivity summarization based on minimum cuts.
In *Proceedings of ACL-04, 42nd Meeting of the Association for Computational Linguistics*, pages 271–278, Barcelona, ES. Association for Computational Linguistics.

Pang, B., Lee, L., and Vaithyanathan, S. (2002).
Thumbs up? sentiment classification using machine learning techniques.
In *Proceedings of EMNLP-02, the Conference on Empirical Methods in Natural Language Processing*, pages 79–86, Philadelphia, US. Association for Computational Linguistics.

Turney, P. (2002).
Thumbs up or thumbs down? Semantic orientation applied to unsupervised classification of reviews.
In *Proceedings of ACL-02, 40th Annual Meeting of the Association for Computational Linguistics*, pages 417–424. Association for Computational Linguistics.

Turney, P. D. and Littman, M. L. (2003).
Measuring praise and criticism: Inference of semantic orientation from association.
ACM Transactions on Information Systems, 21(4):315–346.

Whitelaw, C., Garg, N., and Argamon, S. (2005).
Using appraisal taxonomies for sentiment analysis.
In *Proceedings of MCLC-05, the 2nd Midwest Computational Linguistic Colloquium*, Columbus, US.

Wiebe, J., Wilson, T., and Cardie, C. (2005).
Annotating expressions of opinions and emotions in language.
Language Resources and Evaluation, 1(2):0–0.