Sentiment Analysis & Opinion Mining

Tulasi Prasad Sariki,

SCSE,VIT University

Introduction – facts and opinions

- Two main types of information on the Web.
  - Facts(Objective) and Opinions(Subjective)

    Fact : Thursday is a day.
    Opinion : Thursday was a fun day.
    Fact : iPhone is an Apple product.
    Opinion : iPhone is good.

- Google searches for facts (currently)
- Facts can be expressed with topic keywords
- Google does not search for opinions
  - Opinions are hard to express with keywords
  - Current search ranking strategy is not appropriate for opinion search(AskUs)

Motivation

- ***What others think*** has always been an important piece of information
  - *“Which car should I buy?”*
  - *“Which schools should I apply to?”*
  - *“Which Professor to work for?”*
  - *“Whom should I vote for?”*
- ***Data Sources (User Generated Content)***
  - Blogs
  - Review Sites (amazon)
  - Micro-blogging (Twitter)

Introduction - User generated content

- Word-of-mouth on the Web
  - One can express opinions on almost anything(product, service, person, movie, location, event, organization etc.), at review sites, forums, discussion groups, blogs etc. (called user generated content.)
  - They contain valuable information
- Our interest:
  - To mine opinions expressed in user generated content
  - An intellectually very challenging problem.
  - Practically very useful.

Introduction - Applications

- Businesses and organizations: Market intelligence.
  - Business spends a huge amount of money to find consumer opinions.
- Individuals: interested in other’s opinions when
  - Purchasing a product.
  - Using a service.
  - Finding opinions on political topics.
  - Many other decision making tasks.
- Ads placements: Placing ads in user-generated content
  - Place an ad when one praises an product.
  - Place an ad from a competitor if one criticizes an product.

Introduction - Opinion Mining

- Opinion mining 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, Subjectivity Analysis, Emotion Analysis, Review Mining* are other names used in literature to identify this discipline.

Introduction - Kinds of opinions

- Two types of evaluation
  1. Mining Regular opinions(Direct/Indirect)

     This Camera is great.

     The Picture Quality of Camera is Great. } **Direct**(simpler to handle)

     After Injection of the drug, my joints felt worst <- **Indirect** (Harder to deal)

     Since my joints were painful, my doctor put me on this drug <- ***Not an Opinion***

     Crux: drug injection before or after
  2. Mining Comparative(Comparative or Superlative) opinions (Explicit/Implicit)

     Coke tastes better than Pepsi.

     Coke tastes the best.

     Explicit ← Simpler Implicit ← Harder

Introduction - Typical opinion search queries

- Find the opinion of a person or organization (opinion holder) on a particular object or a feature of an object.
  - E.g., what is the kejriwal's opinion on jan lokpal bill?
  - Find positive and/or negative opinions on a particular object (or some features of the object), e.g.,
  - customer opinions on a digital camera,
  - public opinions on a political topic.
- Find how opinions on an object change with time.
- How object A compares with Object B?
  - Gmail vs. Yahoo mail

Introduction - Find the opinion of a person on X

- In some cases, the general search engine can handle it, i.e., using suitable keywords.
  - kejriwal’s opinion on jan lokpal bill
- Reason:
  - One person or organization usually has only one opinion.
  - The opinion is likely contained in a single document.
  - Thus, a good keyword query may be sufficient.

Introduction - Find opinions on an object X

We use the product reviews as an example:

- Searching for opinions in product reviews is different from general Web search.
  - E.g., search for consumer opinions on a digital camera
- General Web search: rank pages according to some authority and relevance scores.
  - The user looks at the first page (if the search is perfect).
- Opinion search: rank is desirable, however
  - reading only the review ranked at the top is dangerous because it is only the opinion of one person.

Introduction - Search opinions (contd.)

Ranking:

- produce two rankings
  - Positive opinions and negative opinions
  - Some kind of summary of both, e.g., # of each
- Or, one ranking but
  - The top (say 30) reviews should reflect the natural distribution of all reviews (assume that there is no spam), i.e., with the right balance of positive and negative reviews.
- Questions:
  - Should the user reads all the top reviews?
  - Or should the system prepare a summary of the reviews?

Introduction - Reviews are similar to surveys

- Reviews can be regarded as traditional surveys.
  - In traditional survey, returned survey forms are treated as raw data.
  - Analysis is performed to summarize the survey results.
    - E.g., % against or for a particular issue, etc.
- In opinion search,
  - Can a summary be produced?
    - Yes
  - What should the summary be?
    - sentiment prediction (by aggregating the sentiment scores)

Introduction-Opinion (the abstraction)

- Basic components of an opinion
  - Opinion holder: The person or organization that holds a specific opinion on a particular object.
  - Object: on which an opinion is expressed
    - Attributes / Components (Features)
  - Opinion: a view, attitude, or appraisal on an object from an opinion holder.

Introduction-Opinion Representation

An opinion is a quin-tuple, (ei; aij ; ooijkl; hk; tl), where ei is the name of an entity, aij is an aspect of ei, ooijkl is the orientation of the opinion about aspect aij of entity ei, hk is the opinion holder, and tl is the time when the opinion is expressed by hk.

- The opinion orientation ooijkl can be +ve, -ve, neutral, or be expressed with a different intensity levels.
- When an opinion is on the entity itself as a whole, we use the special aspect GENERAL to denote it.
- Ref: A Survey Of Opinion Mining And Sentiment Analysis : Bing Liu, Sentiment Analysis and Opinion Mining Handbook, April 22, 2012, Bing Liu.

OM can be done at Various Levels

Opinion Mining Levels of Granularity

- Document level sentiment classification
- Sentence level sentiment analysis
- Feature-based opinion mining and summarization
- Comparative sentence and relation extraction

Linguistic Concepts for Opinion Mining

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 nature.
- Determining the strength of term attitude (either orientation or subjectivity), as in attributing to terms (real-valued) degrees of positivity or negativity.

Linguistic Concepts for Opinion Mining

Example

- good, excellent, best – positive terms
- bad, wrong, worst – negative terms
- vertical, yellow, liquid – objective terms

Not only terms:

- Term senses, thus taking into account the fact that different senses of the same ambiguous term may have different sentiment-related properties.
  - estimable – ambiguous term with an objective sense (i.e. measurable), and a positive sense (i.e. deserving respect).
- Multi-word expressions
  - not entirely satisfactory – negative expression

Linguistic Concepts for Opinion Mining

Orientation of terms

- Determining if a subjective term has a +ve or a -ve orientation.
- Adjectives in and conjunctions usually have similar orientation, though but is used with opposite orientation.
  - Opinion Mining is Good ***and*** Interesting (positive)
  - History is Good but Boring (neutral)
  - Hitler is Bad and Cruel (negative)

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.

Linguistic Concepts for Opinion Mining

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}

Linguistic Concepts for Opinion Mining

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.

Linguistic Concepts for Opinion Mining

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 Lp and Ln 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.

Linguistic Concepts for Opinion Mining

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

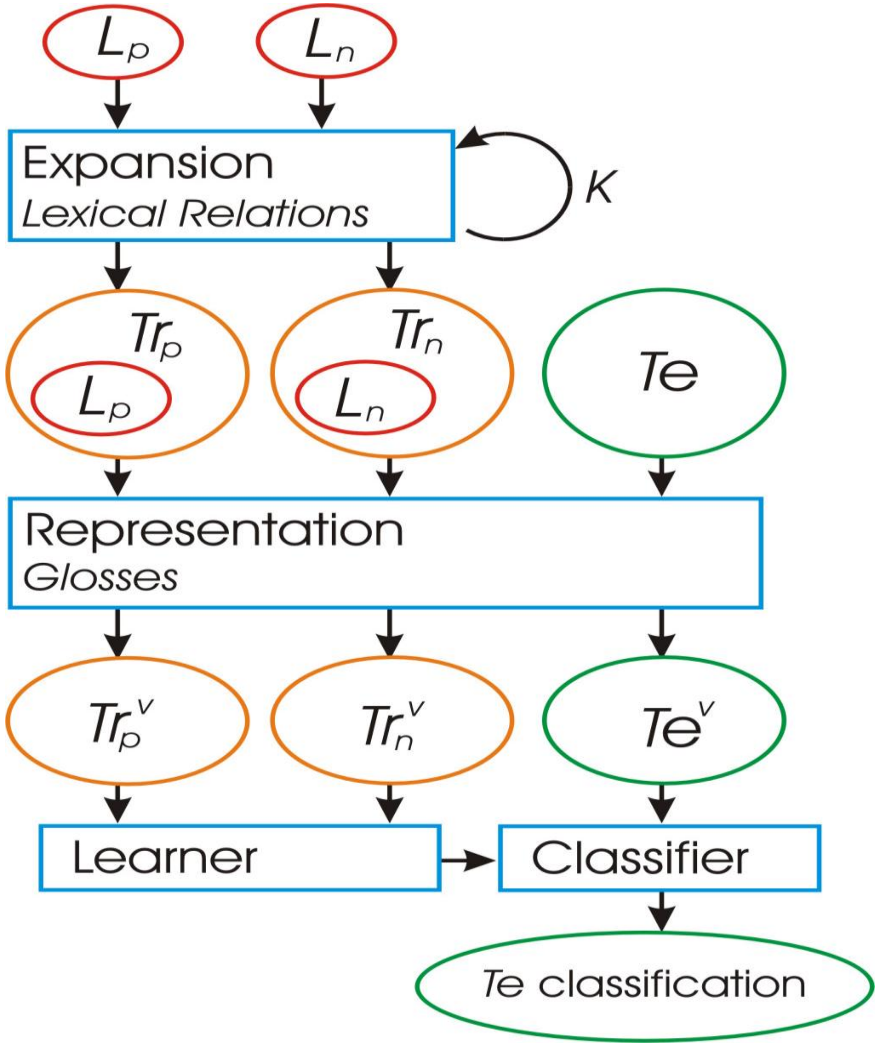

Determining the overall sentiment

- The orientation of the whole document is the sum of the orientation of all its parts.
- PMI method has been applied to classify(Semantic Orientation)
- Learners: Naıve Bayes, Maximum Entropy, SVM.
- Features: unigrams, bigrams, adjectives, POS, position.
- Preprocessing: negation propagation.
- Representation binary, frequency.

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 (not true in many discussion posts) and contains opinion from a single opinion holder.
  - 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.

Various Studies at Document level

Paper	Technique	Features	Dataset	Type
[1]	SVM, NB	Unigram, Bigrams, Trigrams	Restaurant Review	Supervised
[2]	SVM, NB, ME	Unigram, Bigrams, Adjective Positions	Movie Review	Supervised
[3]	SVM, NB, N-gram	Unigram Frequency	Travel Destination	Supervised
[4]	SVM, Rule based	POS tagging, Ngrams	Product Review	Supervised
[5]	SVM	Unigram, Bigram	Movie Review	Supervised
[6]	SVM	Adjective word frequency, percentage of appraisal groups	Movie Review	Supervised
[7]	PMI-IR	Adjectives, Adverbs	Automobile	Unsupervised
[8]	ARM	Adjectives, Adverbs	Movie Review	Unsupervised
[9]	Dictionary based	Adjectives ,Nouns, verbs , Adverbs, Intensifier, Negation	Movie, Cemara	Unsupervised

Conclusions Drawn from the above Studies

- Applied several machine learning and data mining techniques to classify the reviews into positive and negative.
- In supervised learning 3 techniques were tried:
  - Naïve Bayes
  - Maximum entropy
  - Support vector machine
  - SVM: the best accuracy 83%
  - Limitation with supervised learning is that it is sensitive to the quantity & quality of the training data(preprocessing).

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 or phrases.

Various Studies at Sentence level

- [12] Used graphical models in which document level sentiment is linked to several paragraph level sentiments and each paragraph level sentiment is linked to several sentence level sentiments.
- [13] Developed a conditional random field model structured like the dependency parse trees of sentences, by considering opinionated words and sentence shifters.
- [14] Developed a system for computing sentiment of sentences based on the words in the sentence using [11],[15] (appraisal theory, some rules)

Various Studies at Sentence level

[16] Corpus based methods by considering syntactic patterns and co-occurrences of patterns.

[17][18][19][20][21] By constructing Lexicons and by computing PMI (Using sentiwordnet, synonyms, antonyms).

OM Using Sentiwordnet

- Subjectivity and orientation of term senses
- SentiWordNet is a lexical resource that assigns to each synset of WordNet 3 sentiment scores: positivity, negativity, objectivity. (Approx. 1.7 Million words)
- The sum of the scores for a synset is alwaysone.
- Drawback is : Domain independent

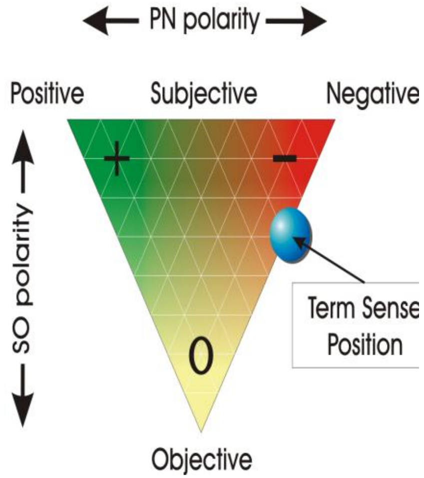

Ref: [Esuli and Sebastiani, 2006]

OM Using Sentiwordnet

Very comfortable, but straps go *loose quickly*.

comfortable

- Positive: 0.75
- Objective: 0.25
- Negative: 0.0

loose

- Positive: 0.0
- Objective: 0.375
- Negative: 0.625

Overall - Positive

- Positive: 0.75
- Objective: 0.625
- Negative: 0.625

At the feature level

- *Task* 1: Identify and extract object features that have been commented on by an opinion holder
- *Task* 2: Determine whether the opinions on the features are positive, negative or neutral.
- *Task* 3: Group feature synonyms.
- Produce a feature-based opinion summary of multiple reviews.

Various Studies at the feature level

Paper	Technique	Type
[25]	ARM and PMI	Unsupervised
[26]	ARM	Unsupervised
[27]	Likely hood Ratio	Unsupervised
[28]	CRF Approch	Supervised
[29]	SVM	Supervised
[30]	HMM	Supervised
[31]	Double Propagation- Syntactic relation	Semi supervised

Evolution of Sentiment Classification

- Based on four Indexes (Confusion Matrix)
  - Accuracy

    $$TP + TN / TP + TN + FP + FN$$

  - Precision

    $$TP / TP + FP$$

  - Recall

    $$TP / TP + FN$$

  - F1-score

    Harmonic mean of precession and recall

Sentiment Summary and Visualization

- Summary of reviews of

Digital camera 1

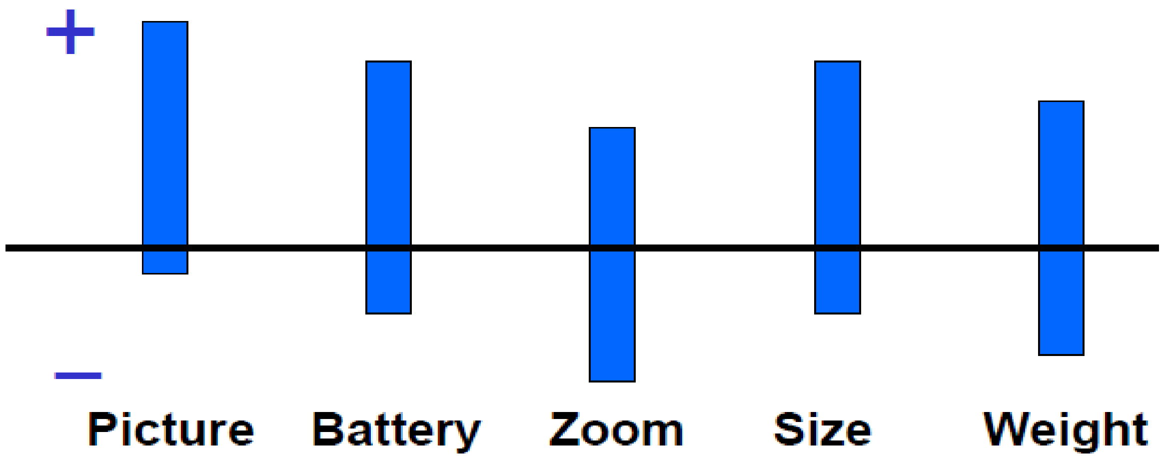

- Comparison of reviews of

Digital camera 1

Digital camera 2

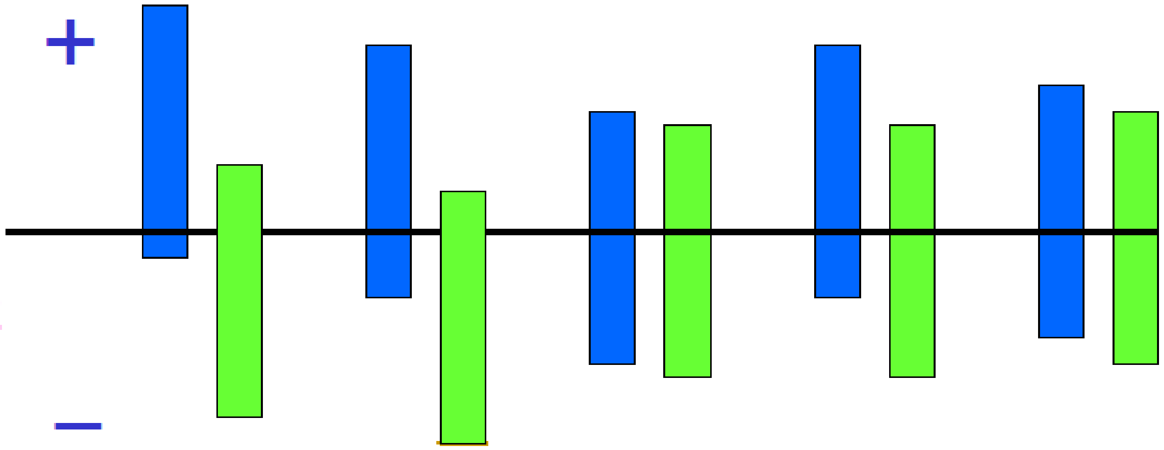

Challenges in Sentiment Analysis

look at reviews on one site only…”

- Problems?
  - Biased views
  - all reviewers on one site may have the same opinion
- Fake reviews/Spam
  - people post good reviews about their own product OR services
  - some posts are plain spams

Challenges in Sentiment Analysis

- Word sense disambiguation
- Preprocessing and Cleaning
- Multi Class classification
- Dealing with Noise
  - Smiles, Special Symbols
- Negation handling
  - I didn’t like apple products
- Unstructured Data Slangs/ Abbreviations
  - Lol (loughing out loud), omg(Oh My god)

Challenges in Sentiment Analysis

- Sentence Segmentation
- Feature Extraction (implicit vs explicit)
- Ambiguous Words
    - This product is literal waste of time
    - Throw your waste material here

References

[1] Z. Zhang, Q. Ye, Z. Zhang, and Y. Li, "Sentiment classification of Internet restaurant reviews written in Cantonese,"Expert Syst. Appl., vol. 38, pp. 7674-7682, 2011.

[2] B. Pang, L. Lee, and S. Vaithyanathan, "Thumbs up?: sentiment classification using machine learning techniques,“ presented at the Proceedings of the ACL-02 conference on Empirical methods in natural language processing - Volume 10, 2002.

[3] Q. Ye, Z. Zhang, and R. Law, "Sentiment classification of online reviews to travel destinations by supervised machine learning approaches," Expert Systems with Applications, vol. 36, pp. 6527-6535, 2009.

[4] R. Prabowo and M. Thelwall, "Sentiment analysis: A combined approach," Journal of Informetrics, vol. 3, pp.143-157, 2009.

[5] E. Riloff, S. Patwardhan, and J. Wiebe, "Feature subsumption for opinion analysis," presented at the Proceedings of the 2006 Conference on Empirical Methods in Natural Language Processing, Sydney, Australia, 2006.

References

[6] C. Whitelaw, N. Garg, and S. Argamon, "Using appraisal groups for sentiment analysis," presented at the Proceedings of the 14th ACM international conference on Information and knowledge management, Bremen, Germany, 2005.

[7] P. D. Turney, "Thumbs up or thumbs down?: semantic orientation applied to unsupervised classification of reviews,“ presented at the Proceedings of the 40th Annual Meeting on Association for Computational Linguistics, Philadelphia, Pennsylvania, 2002.

[8] A. Harb, M. Planti, G. Dray, M. Roche, Fran, o. Trousset, and P. Poncelet, "Web opinion mining: how to extract opinions from blogs?," presented at the Proceedings of the 5th international conference on Soft computing as transdisciplinary science and technology, Cergy-Pontoise, France, 2008.

[9] M. Taboada, J. Brooke, M. Tofiloski, K. Voll, and M.Stede, "Lexicon-based methods for sentiment analysis,"Comput. Linguist., vol. 37, pp. 267-307, 2011.

[10] H. Cui, V. Mittal, and M. Datar, "Comparative experiments on sentiment classification for online product reviews,“ presented at the proceedings of the 21st national conference on Artificial intelligence - Volume 2, Boston, Massachusetts, 2006.

References

[11] J. R. Martin and P. R. R. White, The Language Of Evaluation: Appraisal In English. London: Palgrave, 2005.

[12] R. McDonald, K. Hannan, T. Neylon, M. Wells, and J. Reynar, "Structured Models for Fine-to-Coarse Sentiment Analysis," in Proceedings of the 45th Annual Meeting of the Association of Computational Linguistics, 2007.

[13] T. Nakagawa, K. Inui, and S. Kurohashi, "Dependency tree-based sentiment classification using CRFs with hidden variables," presented at the Human Language Technologies: The 2010 Annual Conference of the North American Chapter of the Association for Computational Linguistics, Los Angeles, California, 2010.

[14] A. Neviarouskaya, H. Prendinger, and M. Ishizuka, "Recognition of affect, judgment, and appreciation in text,"presented at the Proceedings of the 23rd InternationalConference on Computational Linguistics, Beijing, China, 2010.

[15] C. E. Izard, The Face of Emotion: Appleton Century Crofts, 1971.

References

[16] V. Hatzivassiloglou and K. R. McKeown, "Predicting the semantic orientation of adjectives," presented at the Proceedings of the eighth conference on European chapter of the Association for Computational Linguistics, Madrid, Spain, 1997.

[17] P. D. Turney and M. L. Littman, "Measuring praise and criticism: Inference of semantic orientation from association,“ ACM Trans. Inf. Syst., vol. 21, pp. 315-346, 2003.

[18] A. Esuli and F. Sebastiani, "Determining the semantic orientation of terms through gloss classification," presented at the Proceedings of the 14th ACM international conference on Information and knowledge management, Bremen, Germany, 2005.

[19] A. Esuli and F. Sebastiani, "Determining Term Subjectivity and Term Orientation for Opinion Mining Andrea Esuli,“ In Proceedings of the 11th Conference of the European Chapter of the Association for Computational Linguistics (EACL’06), 2006.

[20] S. Baccianella, A. Esuli, and F. Sebastiani, "SentiWordNet 3.0: An Enhanced Lexical Resource for Sentiment Analysis and Opinion Mining," in Proceedings of the Seventh conference on International Language Resources and Evaluation (LREC'10), 2010, pp. 2200-2204.

References

[21] S.-M. Kim and E. Hovy, "Determining the sentiment of opinions," presented at the Proceedings of the 20th international conference on Computational Linguistics, Geneva, Switzerland, 2004.

[22] M. Hu and B. Liu, "Mining and summarizing customer reviews," presented at the Proceedings of the tenth ACM SIGKDD international conference on Knowledge discovery and data mining, Seattle, WA, USA, 2004.

[23] J.-Y. Yang, H.-J. Kim, and S.-G. Lee, "Feature-based product review summarization utilizing user score," Journal of information science and engineering, vol. 26, pp.1973-1990, 2010.

[24] B. Liu and L. Zhang, "A Survey of Opinion Mining and Sentiment Analysis," C. C. Aggarwal and C. Zhai, Eds., ed: Springer US, 2012, pp. 415-463.

[25] A.-M. Popescu and O. Etzioni, "Extracting product features and opinions from reviews," presented at the Proceedings of the conference on Human Language Technology and Empirical Methods in Natural Language Processing, Vancouver, British Columbia, Canada, 2005.

[26] J. Yi, T. Nasukawa, R. Bunescu, and W. Niblack, "Sentiment Analyzer: Extracting Sentiments about a Given Topic using Natural Language Processing Techniques," presented

at the Proceedings of the Third IEEE International Conference on Data Mining, 2003.

References

[27] N. Jakob and I. Gurevych, "Extracting opinion targets in a single- and cross-domain setting with conditional random fields," presented at the Proceedings of the 2010 Conference on Empirical Methods in Natural Language Processing, Cambridge, Massachusetts, 2010.

[28] J. S. Kessler and N. Nicolov, "Targeting sentiment expressions through supervised ranking of linguistic configurations,“ in Proceedings of the Third International AAAI Conference on Weblogs and Social Media,, San Jose, California, USA, 2009, pp. 90-97.

[29] W. Jin, H. H. Ho, and R. K. Srihari, "OpinionMiner: a novel machine learning system for web opinion mining and extraction," presented at the Proceedings of the 15th ACM SIGKDD international conference on Knowledge discovery and data mining, Paris, France, 2009.

[30] V. Stoyanov and C. Cardie, "Topic identification for fine-grained opinion analysis," presented at the Proceedings of the 22nd International Conference on Computational Linguistics - Volume 1, Manchester, United Kingdom, 2008.

[31] G. Qiu, B. Liu, J. Bu, and C. Chen, "Opinion word expansion and target extraction through double propagation,“ Comput. Linguist., vol. 37, pp. 9-27, 2011.

[32] K. Dave, S. Lawrence, and D. M. Pennock, "Mining the peanut gallery: opinion extraction and semantic classification of product reviews," presented at the Proceedings of the 12th international conference on World Wide Web, Budapest, Hungary, 2003.

[33] N. Godbole, M. Srinivasaiah, and S. Skiena, "Large-scale sentiment analysis for news and blogs," in International Conference on Weblogs and Social Media (ICWSM),2007, pp. 219-222.