

# Automatic Expansion of Feature-Level Opinion Lexicons

Fermín L. Cruz, José A. Troyano, F. Javier Ortega, Fernando Enríquez

University of Seville

Avda. Reina Mercedes s/n.

41012 Seville, Spain

{fcruz,troyano,javierortega, fenros}@us.es

## Abstract

In most tasks related to opinion mining and sentiment analysis, it is necessary to compute the semantic orientation (i.e., positive or negative evaluative implications) of certain opinion expressions. Recent works suggest that semantic orientation depends on application domains. Moreover, we think that semantic orientation depends on the specific targets (*features*) that an opinion is applied to. In this paper, we introduce a technique to build domain-specific, feature-level opinion lexicons in a semi-supervised manner: we first induce a lexicon starting from a small set of annotated documents; then, we expand it automatically from a larger set of unannotated documents, using a new graph-based ranking algorithm. Our method was evaluated in three different domains (headphones, hotels and cars), using a corpus of product reviews which opinions were annotated at the feature level. We conclude that our method produces feature-level opinion lexicons with better accuracy and recall than domain-independent opinion lexicons using only a few annotated documents.

## 1 Introduction

*Sentiment analysis* is a modern subdiscipline of natural language processing which deals with subjectivity, affects and opinions in texts (a good survey on this subject can be found in (Pang and Lee, 2008)). This discipline is also known as *opinion mining*, mainly in the context of text mining and information extraction. Many classification and extraction problems have been defined, with different levels of granularity depending on applications requirements: e.g.

classification of text documents or smaller pieces of text into objective and subjective, classification of opinionated documents or individual sentences regarding the overall opinion (into “positive” and “negative” classes, or into a multi-point scale) or extraction of individual opinions from a piece of text (may include opinion target, holder, polarity or intensity of the opinions, among others). As a key in solving most of these problems, the *semantic orientation* of some opinion expressions should be computed: a numeric value, usually between  $-1$  and  $1$ , referring to the negative or positive affective implications of a given word or phrase. These values can be collected in an *opinion lexicon*, so this resource can be accessed when needed.

Many recent works (Popescu and Etzioni, 2005; Kanayama and Nasukawa, 2006; Cruz et al., 2010; Qiu et al., 2011) suggest the need for domain-specific opinion lexicons, containing semantic orientations of opinion expressions when used in a particular domain (e.g., the word “predictable” has opposite semantic orientations when used to define the driving experience of a car or the plot of a movie). Moreover, within a given domain, the specific target of the opinion is also important to induce the polarity and the intensity of the affective implications of some opinion expressions (consider for example the word “cheap” when referring to the *price* or to the *appearance* of an electronic device). This is especially important to extract opinions from product reviews, where users write their opinions about individual features of a product. These domain-specific, feature-level opinion lexicons can be manually collected, but it implies a considerable amount of time

and effort, especially if a large number of different domains are considered.

In this work, we propose a method to automatically induce feature-level, domain-specific opinion lexicons from an annotated corpus. As we are committed to reduce the time and effort, we research about the automatic expansion of this kind of lexicons, so we keep the number of required annotated documents as low as possible. In order to do so, we propose a graph-based algorithm which can be applied to other knowledge propagation problems.

In the next section, we review some related previous works to contextualize our approach. In section 3, we define the feature-level opinion lexicons and describe our method to induce and expand them in a semi-supervised manner. In section 4, we carry out some experiments over a dataset of reviews of three different domains. Finally, we discuss the results and draw some conclusions in section 5.

## 2 Related work

In this section, we briefly discuss some related works about semantic orientation induction and opinion lexicon expansion, pointing out the main differences with our contribution. We also introduce the feature-based opinion extraction task, since it is the natural application context for feature-level opinion lexicons.

### 2.1 Semantic orientation induction

Many methods for computing semantic orientations of words or phrases have been proposed over the last years. Some of them rely on a large set of text documents to compute semantic orientations of words in an unsupervised manner (Hatzivassiloglou and McKeown, 1997; Turney and Littman, 2003; Yu and Hatzivassiloglou, 2003). They all start from a few positive and negative seeds, and calculate the semantic orientation of target words based on conjunctive constructions (Hatzivassiloglou and McKeown, 1997) or co-occurrences (Turney and Littman, 2003; Yu and Hatzivassiloglou, 2003) of target words and seeds. These methods allow computing domain-specific semantic orientations, just using a set of documents of the selected domain, but they obtain modest values of recall and precision. We are using the observations about conjunctive constructions

from (Hatzivassiloglou and McKeown, 1997) in our approach.

Other works use the lexical resource WordNet (Fellbaum, 1998) to compute the semantic orientation of a given word or phrase. For example, in (Kamps et al., 2004), a distance function between words is defined using WordNet synonymy relations, so the semantic orientation of a word is calculated from the distance to a positive seed (“good”) and a negative seed (“bad”). Other works use a bigger set of seeds and the synonyms/antonyms sets from WordNet to build an opinion lexicon incrementally (Hu and Liu, 2004a; Kim and Hovy, 2004). In other works (Esuli and Sebastiani, 2006; Baccianella et al., 2010; Esuli and Sebastiani, 2005), the basic assumption is that if a word is semantically oriented in one direction, then the words in its gloss (i.e. textual definitions) tend to be oriented in the same direction. Two big sets of positive and negative words are built, starting from two initial sets of seed words and growing them using the synonymy and antonymy relations in WordNet. For every word in those sets, a textual representation is obtained by collecting all the glosses of that word. These textual representations are transformed into vectors by standard text indexing techniques, and a binary classifier is trained using these vectors. The same assumption about words and their glosses is made by Esuli and Sebastiani (2007), but the relation between words and glosses are used to build a graph representation of WordNet. Given a few seeds as input, two scores of positivity and negativity are computed, using a random-walk ranking algorithm similar to PageRank (Page et al., 1998). As a result of these works, an opinion lexicon named SentiWordNet (Baccianella et al., 2010) is publicly available. We are also using a ranking algorithm in our expansion method, but applying it to a differently built, domain-specific graph of terms.

The main weakness of the dictionary-based approaches is that they compute domain-independent semantic orientations. There are some manually-collected lexicons (Stone, 1966; Cerini et al., 2007), with semantic orientations of terms set by humans. However, they are also domain-independent resources.

## 2.2 Opinion lexicon expansion

There are a couple of works that deal with the more specific problem of opinion lexicon expansion. In (Kanayama and Nasukawa, 2006), the authors propose an algorithm to automatically expand an initial opinion lexicon based on *context coherency*, the tendency for same polarities to appear successively in contexts. In (Qiu et al., 2011), a method to automatically expand an initial opinion lexicon is presented. It consists of identifying the syntactic relations between opinion words and opinion targets, and using these relations to automatically identify new opinion words and targets in a bootstrapping process. Then, a polarity (positive or negative) is assigned to each of these new opinion words by applying some contextual rules. In both works, the opinion lexicons being expanded are domain-specific, but they are not taking into account the dependency between the specific targets of the opinions and the semantic orientations of terms used to express those opinions. To our knowledge, there are no previous works on inducing and expanding feature-level opinion lexicons.

## 2.3 Feature-based opinion extraction

Feature-based opinion extraction is a task related to opinion mining and information extraction. It consists of extracting individual opinions from texts, indicating the polarity and the specific target of each opinion; then, these opinions can be aggregated, summarized and visualized. It was first defined by Hu and Liu (2004b), and attempted by many others (Popescu and Etzioni (2005), Ding et al. (2008) and Cruz et al. (2010), among others), because of its practical applications. Being a key element in this task, most of these works propose algorithms to compute semantic orientations of terms, generally domain-specific orientations. We aim to build not only domain-specific but also feature-level opinion lexicons, in an attempt to improve the performance of a feature-based opinion extraction system (a description of our system can be found in (Cruz et al., 2010)).

## 3 Proposed method

In this section we define *feature-level opinion lexicons* and propose a semi-supervised method to obtain it. The method consists of two main steps. First,

a small lexicon is induced from a set of annotated documents. Then, the lexicon is automatically expanded using a set of unannotated documents.

### 3.1 Definitions

A domain  $D$  is a class of entities with a fixed set of opinable features  $F_D$ . A *feature* is any component, part, attribute or property of an entity. A *feature-based opinion* is any piece of text with positive or negative implications on any feature of an entity. We name *opinion words* to the minimum set of words from an opinion from which you can decide the *polarity* (i.e., if it is a positive or a negative opinion). A *feature-level opinion lexicon*  $L_D$  for a given domain  $D$  is a function  $T \times F_D \rightarrow [-1.0, 1.0]$ , where  $T$  is a set of *terms* (i.e., individual words or phrases), and  $F_D$  is the set of opinable features for the domain  $D$ .  $L_D$  assign a semantic orientation to each term from  $T$  when used as opinion words in an opinion on a particular feature from  $F_D$ .

### 3.2 Induction

In order to generate a feature-based opinion lexicon to be used as seed in our expansion experiments, we collect a set of text reviews  $R_D$  on a particular domain  $D$ , and annotate all the feature-based opinions we encounter. Each opinion is a tuple  $(polarity, f, opW)$ , where *polarity* is + (positive) or - (negative),  $f$  is a feature from  $F_D$ , and  $opW$  is a set of opinion words from the text. Each annotated opinion gives information about the semantic orientation of the opinion words. Most of the times, the polarity of the opinion implies the polarity of the opinion words. But sometimes, the opinion words include some *special expressions* that have to be considered to induce the polarity of the rest of opinion words, as *negation expressions*<sup>1</sup>, which invert the polarity of the rest of opinion words; and *dominant polarity expressions*<sup>2</sup>, which completely determine the polarity of an opinion, no matter which other opinion words take part. For each opinion term observed (individual words or phrases included as opinion words, once negation and dominant polarity

<sup>1</sup>*Negation expressions*: barely, hardly, lack, never, no, not, not too, scarcely.

<sup>2</sup>*Dominant polarity expressions*: enough, sufficient, sufficiently, reasonably, unnecessarily, insufficiently, excessively, excessively, overly, too, at best, too much.

expressions been removed), the final semantic orientation for a given feature is the mean of the semantic orientations suggested by each annotated opinion on that feature containing the opinion expression (we take 1.0/-1.0 for each positive/negative annotation).

### 3.3 Expansion

Starting from a big set of unannotated text reviews  $R'_D$ , we use the information provided by conjunctive constructions to expand the lexicon previously induced. As explained by Hatzivassiloglou and McKeown (1997), two opinion terms appearing in a conjunctive constructions tend to have semantic orientations with the same or opposite directions, depending on the conjunction employed. Based on this principle, we build a graph linking those terms appearing in a conjunctive expression. We compute the semantic orientation of each term spreading the information provided by those terms in the initial lexicon through the graph. In order to do that, we propose a new random-walk ranking algorithm with the ability to deal with graphs containing positively and negatively weighted edges.

#### 3.3.1 Building the graph

The graph is built from  $R'_D$ , searching for conjunctive constructions between terms. Two terms participate in a conjunctive construction if they appear consecutively in the text separated by a conjunction *and* or *but*, or the punctuation mark *comma* (.). There are two types of conjunctive constructions, *direct* and *inverse*, depending on the conjunction and the negation expressions participating. In a direct conjunctive construction, both terms seems to share the same semantic orientation; in a reverse one, they might have opposite semantic orientations. Some examples are shown next:

- **Direct conjunctive constructions**

The camera has a **bright and accurate** len.  
 It is a **marvellous, really entertaining** movie.  
 ... **clear and easy to use** interface.  
 ... **easy to understand, user-friendly** interface.

- **Inverse conjunctive constructions**

The camera has a **bright but inaccurate** len.  
 It is a **entertaining but typical** film.  
 The driving is **soft and not aggressive**.

The terms observed in conjunctive constructions (in bold type in the previous examples) are the nodes of the graph. If two terms participate in a conjunctive construction, the corresponding nodes are linked by an edge. Each edge is assigned a weight equal to the number of direct conjunctive constructions minus the number of inverse conjunctive constructions observed between the linked terms.

#### 3.3.2 PolarityRank

We propose a new random-walk ranking algorithm, named PolarityRank. It is based on PageRank (Page et al., 1998). In summary, PageRank computes the relevance of each node in a graph based on the incoming edges and the relevance of the nodes participating in those edges; an edge is seen as a recommendation of one node to another. PolarityRank generalizes the concept of vote or recommendation, allowing edges with positive and negative weights. A positive edge still means a recommendation, more strongly the greater the weight of the edge. By contrast, a negative edge represents a negative feedback, more strongly the greater the absolute value of the weight. PolarityRank calculates two scores for each node, a positive and a negative one ( $PR^+$  and  $PR^-$ , respectively). Both scores are mutually dependent: the positive score of a node  $n$  is increased in proportion to the positive score of the nodes linked to  $n$  with positively weighted edges; in addition, the positive score of  $n$  is also increased in proportion to the negative score of the nodes linked to  $n$  with negatively weighted edges. The same principles apply to the calculus of the negative scores of the nodes.

The algorithm definition is as follows. Let  $G = (V, E)$  be a directed graph where  $V$  is a set of nodes and  $E$  a set of directed edges between pair of nodes. Each edge of  $E$  has an associated real value or weight, distinct from zero, being  $p_{ji}$  the weight associated with the edge going from node  $v_j$  to  $v_i$ . Let us define  $Out(v_i)$  as the set of indices  $j$  of the nodes for which there exists an outgoing edge from  $v_i$ . Let us define  $In^+(v_i)$  and  $In^-(v_i)$  as the sets of indices  $j$  of the nodes for which there exists an incoming edge to  $v_i$  whose weight is positive or negative, respectively. We define the positive and negative PolarityRank of a node  $v_i$  (equation 1), where the values  $e^+$  and  $e^-$  are greater than zero for certain nodes acting as positive or negative seeds, re-

spectively. The parameter  $d$  is a damping factor that guarantees convergence; in our experiments we use a value of 0.85 (as recommended in the original definition of PageRank). The computation of  $PR^+$  and  $PR^-$  is done iteratively as described by Page et al. (1998).

$$\begin{aligned}
PR^+(v_i) &= (1-d)e_i^+ + \\
&+ d \left( \sum_{j \in In^+(v_i)} \frac{p_{ji}}{\sum_{k \in Out(v_j)} |p_{jk}|} PR^+(v_j) + \right. \\
&+ \left. \sum_{j \in In^-(v_i)} \frac{-p_{ji}}{\sum_{k \in Out(v_j)} |p_{jk}|} PR^-(v_j) \right) \\
PR^-(v_i) &= (1-d)e_i^- + \\
&+ d \left( \sum_{j \in In^+(v_i)} \frac{p_{ji}}{\sum_{k \in Out(v_j)} |p_{jk}|} PR^-(v_j) + \right. \\
&+ \left. \sum_{j \in In^-(v_i)} \frac{-p_{ji}}{\sum_{k \in Out(v_j)} |p_{jk}|} PR^+(v_j) \right)
\end{aligned} \tag{1}$$

The sum of the values of  $e^+$  and  $e^-$  must be equal to the number of nodes in the graph.

### 3.3.3 Extending the lexicon

Based on a seed lexicon  $L_D$ , and a set of unannotated reviews  $R'_D$ , the expanded lexicon  $L'_D$  is obtained following these steps:

1. Build a graph  $G = (V, E)$  representing the conjunctive relations observed in  $R'_D$ .
2. For each feature  $f$  from  $F_D$ :
  - (a) For each  $v_i$  from  $V$  with associated term  $t_i$ , such that  $L_D(t_i, f)$  is defined, assign that value to  $e_i^+$  if it is greater than 0, else assign it to  $e_i^-$ .
  - (b) Linearly normalize the values of  $e_i^+$  and  $e_i^-$ , so that the sum of the values is equal to  $|V|$ .
  - (c) Compute  $PR^+$  and  $PR^-$ .
  - (d) For each  $v_i$  from  $V$  with associated term  $t_i$ , assign  $SO(v_i)$  to  $L'_D(t_i, f)$ , where:

$$SO(v_i) = \frac{PR^+(v_i) - PR^-(v_i)}{PR^+(v_i) + PR^-(v_i)}$$

Note that these values are contained in the interval  $[-1.0, 1.0]$ .

## 4 Experiments

In this section we report the results of some experiments aimed to evaluate the quality of the feature-level opinion lexicons obtained by our method.

### 4.1 Data

We used a set of reviews of three different domains (*headphones*, *hotels* and *cars*). We retrieved them from Epinions.com, a website specialized in product reviews written by customers. Some reviews from the dataset were labeled, including the polarity, the feature and the opinion words of each individual opinion found. Some information of the dataset is shown in table 1. The dataset is available for public use<sup>3</sup>.

Domain	Reviews	Opinions	Features
Headphones	587 (2591)	3897	31
Hotels	988 (6171)	11054	60
Cars	972 (23179)	8519	91

Table 1: Information of the dataset. The number of unannotated reviews available for each domain is shown in parenthesis.

### 4.2 Experimental setup

All the experiments were done using 10-fold cross-validation. Each annotated dataset was randomly partitioned into ten subsets. The results reported for each experiment are the average results obtained in ten different runs, taking a different subset as testing set and the remaining nine subsets as training set (to induce seed lexicons). To evaluate the lexicons, we compute recall and precision over the terms participating as opinion words in the opinions annotated in the testing set. Recall is the proportion of terms which are contained in the lexicon; precision is the proportion of terms with a correct sentiment orientation in the lexicon.

### 4.3 Results

Table 2 shows the results of the evaluation of the induced and expanded lexicons. In order to figure out the gain in precision and recall obtained by our expansion method, we induced lexicons for each domain using different numbers of annotated reviews

<sup>3</sup><http://www.lsi.us.es/~fermin/index.php/Datasets>

Domain	$ R_D $	Induced Lexicon			Expanded Lexicon			$\delta(p)$	$\delta(r)$	$\delta(F_1)$
		p	r	$F_1$	p	r	$F_1$			
Headphones	9	0.9941	0.4479	0.6176	0.9193	0.7332	0.8158	-0.0748	<b>+0.2853</b>	<b>+0.1982</b>
	45	0.9821	0.7011	0.8181	0.9440	0.8179	0.8764	-0.0381	+0.1168	+0.0583
	108	0.9665	0.8038	0.8777	0.9525	0.8562	0.9018	-0.0140	+0.0524	+0.0241
	531	<b>0.9554</b>	0.9062	0.9302	0.9526	<b>0.9185</b>	<b>0.9352</b>	-0.0028	+0.0123	+0.0051
Hotels	9	0.9875	0.3333	0.4984	0.9416	0.8131	0.8726	-0.0459	<b>+0.4798</b>	<b>+0.3743</b>
	117	0.9823	0.7964	0.8796	0.9716	0.8802	0.9236	-0.0107	+0.0838	+0.0440
	324	0.9822	0.8732	0.9245	0.9775	0.9128	0.9440	-0.0047	+0.0396	+0.0195
	891	<b>0.9801</b>	0.9449	0.9622	0.9792	<b>0.9507</b>	<b>0.9647</b>	-0.0009	+0.0058	+0.0026
Cars	9	0.9894	0.4687	0.6361	0.9536	0.8262	0.8853	-0.0358	<b>+0.3575</b>	<b>+0.2493</b>
	117	0.9868	0.8008	0.8841	0.9712	0.8915	0.9296	-0.0156	+0.0907	+0.0455
	279	0.9849	0.8799	0.9294	0.9786	0.9116	0.9439	-0.0063	+0.0317	+0.0145
	882	<b>0.9847</b>	0.9300	0.9566	0.9831	<b>0.9408</b>	<b>0.9615</b>	-0.0016	+0.0108	+0.0049

Table 2: Results of expansion of lexicons induced from different numbers of annotated reviews. The second and third experiments for each domain are done selecting the number of annotated reviews needed to achieve  $F_1$  scores for the induced lexicon similar to the  $F_1$  scores for the expanded lexicon from the previous experiment.

and expanding them using the whole set of unannotated reviews. For each domain, we show the results of experiments using only nine annotated reviews (one from each subset of reviews of the cross-validation process), and using all the available annotated reviews. The second and third experiments for each domain are those where  $F_1$  scores for the induced lexicon is similar to the  $F_1$  scores for the expanded lexicon from the previous experiment. Thus, we can measure the number of additional annotated reviews needed to obtain similar results without expansion. Using only nine annotated reviews, the expanded feature-level opinion lexicon achieves 0.8158 of  $F_1$  for the *headphones* domain, 0.8764 for the *hotels* domain and 0.8853 for the *cars* domain, a far better result than using a domain-independent opinion lexicon<sup>4</sup>. To obtain similar  $F_1$  scores without using the expansion method, you should annotate between six and thirteen times more reviews.

## 5 Conclusions

There is evidence that the semantic orientation of an opinion term not only depends on the domain, but also on the specific feature which that term is applied to. In this paper, we propose a method to automatically induce domain-specific, feature-level

opinion lexicons from annotated datasets. We research about the automatic expansion of this kind of lexicons, so we keep the number of required annotated documents as low as possible. The results of the experiments confirm the utility of feature-level opinion lexicons in opinion mining tasks such as feature-based opinion extraction, reaching 0.9538 as average of  $F_1$  in three tested domains. Even though if only a few annotated reviews are available, the lexicons produced by our automatic expansion method reach an average  $F_1$  of 0.8592, which is far better than using domain-independent opinion lexicon. Our expansion method is based on the representation of terms and their similarities and differences in a graph, and the application of a graph-based algorithm (PolarityRank) with the ability to deal with positively and negatively weighted graphs. The same algorithm can be applied to other knowledge propagation problems, whenever a small amount of information on some of the entities involved (and about the similarities and differences between the entities) is available. For example, we applied the same algorithm to compute trust and reputation in social networks (Ortega et al., 2011).

## References

Stefano Baccianella, Andrea Esuli, and Fabrizio Sebastiani. 2010. Sentiwordnet 3.0: An enhanced lexical resource for sentiment analysis and opinion mining. In Nicoletta Calzolari (Conference Chair),

<sup>4</sup>We perform some experiment using the domain-independent opinion lexicon SentiWordNet (Baccianella et al., 2010), obtaining  $F_1$  values equal to 0.7907, 0.8199 and 0.8243 for the *headphones*, *hotels* and *cars* domains.

- Khalid Choukri, Bente Maegaard, Joseph Mariani, Jan Odijk, Stelios Piperidis, Mike Rosner, and Daniel Tapias, editors, *Proceedings of the Seventh conference on International Language Resources and Evaluation (LREC'10)*, Valletta, Malta, may. European Language Resources Association (ELRA).
- S. Cerini, V. Compagnoni, A. Demontis, M. Formentelli, and G. Gandini, 2007. *Language resources and linguistic theory: Typology, second language acquisition, English linguistics.*, chapter Micro-WNOp: A gold standard for the evaluation of automatically compiled lexical resources for opinion mining. Franco Angeli Editore, Milano, IT.
- Fermín L. Cruz, José A. Troyano, Fernando Enríquez, Javier Ortega, and Carlos G. Vallejo. 2010. A knowledge-rich approach to feature-based opinion extraction from product reviews. In *Proceedings of the 2nd International Workshop on Search and Mining User-Generated Contents*, pages 13–20. ACM.
- Xiaowen Ding, Bing Liu, and Philip S. Yu. 2008. A holistic lexicon-based approach to opinion mining. In *WSDM '08: Proceedings of the international conference on Web search and web data mining*, pages 231–240, New York, NY, USA. ACM.
- Andrea Esuli and Fabrizio Sebastiani. 2005. Determining the semantic orientation of terms through gloss analysis. In *Proceedings of the ACM SIGIR Conference on Information and Knowledge Management (CIKM)*.
- Andrea Esuli and Fabrizio Sebastiani. 2006. Determining term subjectivity and term orientation for opinion mining. In *Proceedings of the European Chapter of the Association for Computational Linguistics (EACL)*.
- Andrea Esuli and Fabrizio Sebastiani. 2007. Pageranking wordnet synsets: An application to opinion mining. In *Proceedings of ACL-07, the 45th Annual Meeting of the Association of Computational Linguistics*, pages 424–431. Association for Computational Linguistics.
- Christiane Fellbaum, editor. 1998. *WordNet: An Electronic Lexical Database*. MIT Press.
- Vasileios Hatzivassiloglou and Kathleen R. McKeown. 1997. Predicting the semantic orientation of adjectives. In *Proceedings of the eighth conference on European chapter of the Association for Computational Linguistics*, pages 174–181, Morristown, NJ, USA. Association for Computational Linguistics.
- Minqing Hu and Bing Liu. 2004a. Mining and summarizing customer reviews. In *Proceedings of the ACM SIGKDD Conference on Knowledge Discovery and Data Mining (KDD)*, pages 168–177.
- Minqing Hu and Bing Liu. 2004b. Mining and summarizing customer reviews. In *KDD '04: Proceedings of the tenth ACM SIGKDD international conference on Knowledge discovery and data mining*, pages 168–177, New York, NY, USA. ACM.
- Jaap Kamps, Maarten Marx, Robert J. Mokken, and Maarten De Rijke. 2004. Using wordnet to measure semantic orientation of adjectives. In *National Institute for*, volume 26, pages 1115–1118.
- Hiroshi Kanayama and Tetsuya Nasukawa. 2006. Fully automatic lexicon expansion for domain-oriented sentiment analysis. In *Proceedings of the Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 355–363, Sydney, Australia, July. Association for Computational Linguistics.
- Soo-Min Kim and Eduard Hovy. 2004. Determining the sentiment of opinions. In *Proceedings of the International Conference on Computational Linguistics (COLING)*.
- Javier Ortega, José Troyano, Fermín Cruz, and Fernando Enríquez de Salamanca. 2011. PolarityTrust: measuring trust and reputation in social networks. In *Fourth International Conference on Internet Technologies and Applications (ITA 11)*, Wrexham, North Wales, United Kingdom, 9.
- Lawrence Page, Sergey Brin, Rajeev Motwani, and Terry Winograd. 1998. The pagerank citation ranking: Bringing order to the web. Technical report, Stanford Digital Library Technologies Project.
- Bo Pang and Lillian Lee. 2008. Opinion mining and sentiment analysis. *Foundations and Trends in Information Retrieval*, 2(1-2):1–135.
- Ana-Maria Popescu and Oren Etzioni. 2005. Extracting product features and opinions from reviews. In *Proceedings of the Human Language Technology Conference and the Conference on Empirical Methods in Natural Language Processing (HLT/EMNLP)*.
- Guang Qiu, Bing Liu, Jiajun Bu, and Chun Chen. 2011. Opinion word expansion and target extraction through double propagation. *Computational Linguistics*, 37(1).
- Philip J. Stone. 1966. *The General Inquirer: A Computer Approach to Content Analysis*. The MIT Press.
- Peter D. Turney and Michael L. Littman. 2003. Measuring praise and criticism: Inference of semantic orientation from association. *ACM Transactions on Information Systems*, 21:315–346.
- Hong Yu and Vasileios Hatzivassiloglou. 2003. Towards answering opinion questions: Separating facts from opinions and identifying the polarity of opinion sentences. In *Proceedings of the Conference on Empirical Methods in Natural Language Processing (EMNLP)*.