

# Detecting Opinions Using Deep Syntactic Analysis

Caroline Brun

Xerox Research Centre Europe

Meylan, France

Caroline.Brun@xrce.xerox.com

## Abstract

In this paper, we present an opinion detection system built on top of a robust syntactic parser. The goal of this system is to extract opinions associated with products but also with characteristics of these products, i.e. to perform feature-based opinion extraction. To carry out this task, and following a target corpus study, the robust syntactic parser is enriched by associating polarities to pertinent lexical elements and by developing generic rules to extract relations of opinions together with their polarity, i.e. positive or negative. These relations are used to feed an opinion representation model. A first evaluation shows very encouraging results, but numerous perspectives and developments remain to be investigated.

## 1 Introduction

Opinion mining (or sentiment analysis) arouses great interest in recent years both in academia and industry. With the emergence of discussion groups, forums, blogs, web sites compiling consumer reviews on various subjects, there is a huge mass of documents containing information expressing opinions. This constitutes a very important data source for monitoring various applications (business intelligence, product and service benchmarking, technology watch). Consequently, numerous research works at the crossroads of NLP and data mining, are focusing on the problem of opinion detection and mining. In this paper, we present an opinion detection system developed in the framework of the European Project Scoop<sup>1</sup>. This system uses a robust parser specifically adapted for opinion detection, and we focus here on recent developments made for English. Our goal is to extract opinions related to the main concepts commented in the reviews (e.g. products, movies, books...), but also on the features associated to these products (such as

certain characteristics of the products, their price, associated services, etc...).

After a brief review of related work, we describe a corpus analysis conducted on a first target corpus consisting of reviews about printers, copiers and scanners. The following section describes in details the building of the opinion detection system, which makes an intensive use of syntactic information. Finally, we present a preliminary evaluation of the performances of this system and conclude on our perspectives.

## 2 State of the Art

Besides works about lexical resources acquisition for opinion mining, discussed in section 4.3.2, two main types of works can be distinguished: those aiming at classifying texts according to an overall polarity (positive, negative and sometimes neutral), generally based on supervised approaches (such as (Pang et al. 2002), or (Charton and Acuna-Agost 2007)), and those aiming at extracting precise information about positive or negative aspects of a given product or topic. The latter consider that the main concept (e.g. a product) is related to several features (e.g. quality, print speed and resolution for a printer), that can be evaluated separately. Our system belongs to this category. In this case, the goal is to identify related features and opinions expressed about these features. Three sub-tasks are considered: feature extraction, discovery of opinions about these features, and eventually production of a summary of the information associated with a given feature. In order to extract features, methods are generally based on frequency criteria coupled with linguistically-based heuristic, see for example (Yi et al. 2003) or (Popescu and Etzioni 2005). In order to extract opinions about features, a wide range of methods have been proposed: (Hu and Liu 2004) extract the linguistic segments containing a concept and count the polarity of the polar vocabulary present in the same segment. (Vernier et al. 2009) propose a symbolic method to detect and categorize opinions locally expressed in a set of multi-domain

---

<sup>1</sup> <http://www.scoopproject.eu/overview.html>

blogs. Some systems use syntactic dependencies to link source and target of the opinion as in (Kim and Hovy 2006) or (Bloom et al. 2007). Our system belongs to this family, as we believe that syntactic processing of complex phenomena (negation, comparison and anaphora) is a necessary step to perform feature-based opinion mining. A specificity of our system is a two level architecture: it relies on a first level, general and valid across all domain and corpora, and on a second level, adapted for each sub-domain of application.

### 3 Corpus Study

In order to build our opinion detection system, we used a corpus of reviews available on the website "Epinion"<sup>2</sup>. This is a general site compiling millions of user reviews about products, movies, books, etc. As our first target application deals with consumer reviews about printers, we extracted a corpus of about 3,500 printer reviews from this site. These reviews are semi-structured and contain the following information: The product name; the overall score (from 0 to 5 stars); the review title; the creation date; the sections "Pros", "Cons" and "Bottom Line" and the content of the review in free text, with the assessment: "Recommended": "yes" or "no".

This study revealed two important points:

(a) Complex linguistic phenomena are involved in the expression of opinions, and need to be taken into account to build an efficient extraction system:

Syntactic or lexical negation, which inverses the polarity of opinions, as in the following examples:

- *I **can't** use it **without** problems.*
- *There is **no way** I can recommend this printer*

Modality, which affects the strength of the opinion:

- *Considering the high cost of the printer, the quality **should be** outstanding.*

Comparison, which express an opinion comparatively:

- *I would be **happier** with a **better** price.*
- *Performance is **better than** many competing laser printers.*

Anaphora, impacting the detection of the topic of an opinion: In the following example, taken from a review about the "Xerox DocuPrint P8ex Laser Printer", the author refers to many other

products (underlined text) that are not the main topic of the review (bold text):

*Xerox DocuPrint P8ex Laser Printer: When my previous printer (HP LaserJet 5: it was really good at the time) did not last as long as I would like it to have lasted.... I had one functional HP remaining (this one also a good, reliable product but ancient and so slow), one NEC and then I bought **this Xerox**.*

(b) Regarding the subjective vocabulary, i.e. the vocabulary expressing whether an opinion is positive or negative, it is necessary to take into account the following problems:

Ambiguities, because the same word in a given domain can express opinions of different polarity, for example, the adjective "fast" in the domain of printers:

- *It uses ink twice as **fast**.* [Negative]
- *It is a **fast**, high quality printer.* [Positive]

Domain-dependent polarity, because the polarity of a given word can vary across domains:

- *It walks like a **lemon** and quacks like a **lemon*** [Negative]: In product reviews, "lemon" is negative, while this word is generally neutral.
- (i) *Pros: Completely **unpredictable**, Nicholas Cage is awesome.* [Positive]. (ii) *The only problem is that the HP software that runs it appears to be very flaky and **unpredictable**.* [Negative]. In the domain of movie reviews, "unpredictable" is used positively, whereas it is negative in the domain of printers.

Following this study, we designed system with a two level-architecture: the first level contains generic vocabulary, of constant polarity across domains, as well as generic extraction rules, while the second level contains domain-dependent polar vocabulary and specific extraction rules. This system, which benefits of the incremental architecture of the parser we use, is described in detail in the next section.

## 4 Our System

### 4.1 Model of an Opinion

Our goal is to develop a system for extracting opinions on product reviews. We not only aim at classifying reviews as positive or negative (document-level opinion mining), but also at extracting finer-grained opinions expressed about specific features related to a main concept (e.g. speed, print quality etc. in the case of a printer). It seems indeed very interesting to detect precise-

<sup>2</sup> <http://www.epinions.com/>

ly what users like or dislike about a given product, because an overall opinion on a review, either positive or negative, does not necessarily reflect the fact that the user likes or does not like the product as a whole. To achieve this goal, we adopt the formal representation of an opinion given proposed by (Liu, B. 2010): an opinion is a five place predicate of the form  $(o_j, f_{jk}, s_{o_{ijk}}, h_i, t_i)$ , where:

- $o_j$  is the target object of the opinion (the main concept)
- $f_{jk}$  is a feature associated to the object  $o_j$
- $s_{o_{ijk}}$  is the value (positive or negative) of the opinion expressed by the opinion holder  $h_i$  about the feature  $f_{jk}$
- $h_i$  is the opinion holder
- $t_i$  is the time when the opinion is expressed.

Our opinion extraction system is designed on top of a robust syntactic parser (XIP, see below). We use this parser to extract, from syntactic relations already extracted by a general dependency grammar, semantic relations in order to instantiate the five place predicates compliant with this model.

## 4.2 XIP in Brief

We use the Xerox Incremental Parser, XIP, (Ait-Mokhtar et al., 2002) as a fundamental component of our system, in order to extract deep syntactic dependencies, which are an intermediary step to the extraction of semantic relations of opinion. The parser also includes a module for named entity. For this project, since the first application focuses on reviews about printers, a preliminary adaptation was to integrate the recognition of printer names into the NER module.

## 4.3 Design of the System

As said before, we aim at extracting from customer reviews, semantic relations to instantiate five place predicates modeling an opinion. In the context of our application, we can simplify the extraction of the required information, considering that the moment in time when the opinion is expressed is the date of creation of the document and that the opinion holder is the review's author. Moreover, if not mentioned explicitly in the sentence, by default, the object of an opinion is the main topic of the review, i.e., in our case, the product reviewed. For reasons of implementation, we also model the polarity of an opinion as a feature (whose value is "positive" or "negative") associated with the sentiment semantic

relation. Finally, an argument of the sentiment relation is the predicate carrying the opinion. This information can be useful for a subsequent phase of normalization. So we want to extract semantic relations of the form:

SENTIMENT[POLARITY](MAIN-CONCEPT, FEATURE, PREDICATE), for example:

(1) "This printer is slow":

→ SENTIMENT[NEG](printer, \_, slow)

(2) "The laser print quality is great"

→ SENTIMENT[POS](Default, print quality, great)

In the first example, the predicate carrying the opinion is "slow", the object is "printer", the opinion relates to this object entirely and the sentiment is negative. In the second example, the predicate carrying the opinion is "great", the associated feature is "print quality", and as it is not explicitly mentioned, the object of the opinion is the main topic of the review (*Default*).

In order to extract such semantic relationships, we have first extracted the associated features from our corpus, then implemented a polar lexicon, and finally design hand-crafted sentiment extraction rules, according to the two-level architecture mentioned before. These different development steps are now described in detail.

## 4.4 Associated Feature Extraction

The main concepts of our first application are the topic discussed in customer reviews about printers: the vocabulary denoting these concepts is: *printer, copier, scanner, machine, and product*.

To extract the associated features related to these concepts, we use a method partly similar to what is proposed in (Popescu and Etzioni 2005): They seek meronymy relationships (part-whole) to identify related features. We use our parser to extract, from our corpus, the most frequent nouns modifying a main concept, i.e. matching the two following syntactic relations:

• MODIFIER-PRE(MAIN-CONCEPT, CANDIDATE-FEATURE), which matches for example "*printer quality*", where "*quality*" would be extracted as a feature candidate.

• MODIFIEUR\_PREP[OF](CANDIDATE-FEATURE, MAIN-CONCEPT), which matches for example "*the speed of the machine*" for which "*speed*" would be extracted as a candidate feature.

We calculate the frequencies for each candidate feature, and get a list of 736 feature candidates. To filter the noise, we apply the following heuristic: we consider that a candidate is actually a related feature if it is in attributive syntactic rela-

tion at least once with the adjectives “good” or “bad” in the corpus. These syntactic relations are again extracted automatically using the parser.

At the end, we get a list of 76 related features, the most frequent being: *quality, speed, photo, color, software, cartridge, price, resolution...* A manual verification reveals that these words are indeed related features: they refer either to hardware parts of the products (*cartridge, drum*), functional characteristics (*resolution, speed*) or related concepts (*price, support, warranty*).

#### 4.5 Building the Lexicon

The vocabulary encoding the polarity (positive or negative) associated with subjective words contains adjectives (“beautiful” (positive), “ugly” (negative)), nouns (“talent” (positive), “nuisance” (negative)), verbs (“love” (positive), “hate” (negative)) adverbs, (“admirably” (positive), “annoyingly” (negative)). Many studies address this problem. For example, (Agarwal and Bhattacharyaa 2006) are classifying adjectives according to their polarity by using a small set of “seed” adjectives, of known polarity, and calculate their degree of association with other adjectives in a large corpus, the underlying idea being that close adjectives tend to co-occur. (Vegna-duzzo 2004) also classifies adjectives according to their polarities using seed adjectives and a method based on the distributional similarity of the syntactic context. (Esule and Sebastiani 2006) develop SentiWordnet: they carry out a quantitative analysis of definitions (“glosses”) associated with Wordnet synsets using different statistical classifiers to provide three measures for each synset: positivity, negativity and objectivity. This work is particularly challenging and interesting; however, we could not use it in our application, because the ambiguity of each Wordnet lexical entry is preserved. Moreover, this is a very general resource that would not fulfill completely our application needs: for example, the adjective “fast”, mainly considered as objective by SentiWordnet, it is either positive (“fast printer”) or negative (“fast ink consumption”) in printer’s domain. As we do not have at our disposal a manually opinion-annotated corpus, we once again used the syntactic dependencies provided by the parser. We automatically extract a set of syntactic relations, on the entire corpus of reviews to select the vocabulary which is potentially subjective. These relationships are filtered according to the presence of a main concept, or an associated feature, or the personal pronoun “I” in a syntactic relationship.

We extract the following relations:

- ATTRIBUTE(CONCEPT|FEATURE,CANDIDATE) to extract nouns and adjectives in attributive position with a main concept or an associated feature, as in “*the size of the printer is huge*”;
- ATTRIBUTE(PRON\_PERS(I), CANDIDATE), to extract nouns and adjectives in attributive position with the personal pronoun “I” as in “*I am extremely unhappy*”;
- MODIFIER(CONCEPT|FEATURE, CANDIDATE) to extract adjectives modifying a main concept or an associated feature, as in “*It prints great photos*”;
- SUBJECT-VERB\_OBJECT(PRON\_PERS(I), CANDIDATE, CONCEPT|FEATURE), to extract verbs whose subject is “I” and direct object is a main concept or an associated feature, as in “*I appreciate the speed of the printer*”;
- SUBJECT-VERB-OBJECT(CONCEPT|FEATURE, CANDIDATE, PRON\_PERS(I)), to extract verbs whose subject is a main concept or a related feature, and object is the personal pronoun “I” as in “*I am disappointed with this product*”;
- SUBJECT-VERB(CONCEPT|FEATURE, CANDIDATE), to extract verbs whose subject is a main concept or a related feature, as in “*this printer stinks!*”.

The results of the extraction are then filtered according to the syntactic category of the candidate and its number of occurrences in the corpus. Then these candidates are analyzed manually to attach to them the appropriate polarity (positive or negative) and to include them in the general or in the domain-dependant vocabulary. We then use WordNet to find synonyms and antonyms of the selected words. Finally, we obtain 130 verbs in the general lexicon and 42 in the specialized lexicon, 465 adjectives in the general lexicon and 230 in the specialized lexicon, and 145 nouns in the general lexicon and 42 in the specialized lexicon. We thus constructed a “generic” lexicon of polarity, valid for any application and a specialized lexicon, related to the domain of printers.

Moreover, as we work with a robust parser adapted to extract semantic relations of sentiment, the mere mention of polarities in the lexicon is not completely adequate for the development of sentiment extraction rules. It is also necessary to encode information within the predicates in order to be able to detect the scope of the opinions. Typically, we associate semantic features to verbs, indicating if the scope of the opinion is the subject (1), or the direct object of verbs, (2), or on a prepositional complement, (3): (1) “*These printers never cease to amaze me.*”

- (2) “*I appreciate the swiftness of this machine.*”  
 (3) “*We have had several problems with a LaserJet.*”

We needed also to add semantic features to domain specific vocabulary occurring in some specific opinion expressions. Indeed, in the domain of printers, examples of type (4) or (5) are frequent:

- (4) “*This machine was very easy to setup.*”  
 (5) “*It is so easy to operate.*”

Here, it is the combination [easy + to + verb expressing a functional characteristic of the printer] that denotes a positive opinion. We therefore assigned semantic features for verbs of this type in the specialized lexicon.

At the end of this step, we have an attested list of polar words, enriched with syntactico-semantic information. In order to extend the coverage of the lexicon for adjectives, which are intensively used to express opinions, we combined the methods proposed by (Hatzivassiloglou and McKeown 97) and (Monceau et al 2009). They both use information about syntactic conjunction of adjectives to statistically predict their polarity, the underlying idea being that conjunctions give information about the orientation of adjectives: we use our attested list of polar adjectives enriched with 300 hand-coded objective adjectives to train a standard SVM classifier (SVM-multiclass, (Joachim 1999)). In order to do this, we extract from the British National Corpus<sup>3</sup>, with the robust parser, all conjunction relations involving attested polar and objective adjectives, for all types of conjuncts (“and”, “or”, “neither nor” and “but”). For each adjective (negative, positive or objective), we count the number of times it is coordinated with a negative, positive, or objective adjective, for the four type of conjuncts. These numbers of occurrences are used as the values of 12<sup>4</sup> features to train the 3 classes SVM. We used about 350 attested polar adjectives, and 200 objective adjectives for training, and keep about 100 polar adjectives and 100 objective adjectives for validation. We use the resulting model to classify all unknown adjectives appearing in a coordination relation with an attested adjective within the BNC. We end up with 9692 new adjectives, among which 1777 are classified as negative, 1329 as positive and 6586 as objective. From these results, we manually validated 1302 negative adjectives and 995 posi-

tive adjectives, and integrate them into the general polar lexicon.

#### 4.6 Rule Development

Once encoded the polar vocabulary, we developed a set of hand-crafted rules, on top of the output of the deep syntactic parser, to extract semantic relationships denoting opinions. The rules are also divided into two subsets: generic rules and domain-specific rules.

The generic rules are testing, for a semantico-syntactic pattern detected by the parser, the presence of polar vocabulary within the arguments of syntactic relationships. For example:

If(SUBJ-N(#1[polarity,!polarity:!,topic-subj], #2[main-concept]))

→ SENTIMENT[polarity](#2,#1)

Indicates that if the parser has detected that the subject (#2) of a verb (#1) expressing an opinion (feature *polarity*, either positive or negative) is a main concept (feature *main-concept*) then a relationship of sentiment is created using percolation (!polarity:!). This rule associates a positive or negative value to the output relation according to the orientation of the verb. It matches:

- “*These printers#2 never cease to amaze#1 me*”
- “*I was quite disappointed#1 with this machine#2*”

Similar rules are also developed if the scope of the opinion is an associated feature. Moreover, when neither a main concept nor an associated feature is mentioned in the sentence, relations with default values are extracted:

Very nice! → SENTIMENT[POS](default,\_nice)

Do not buy! → SENTIMENT[NEG](default,\_buy).

In the current system, about 60 generic rules are developed to cover the majority of structures identified from the corpus study.

We have also focused on the treatment of negation, since this phenomenon reverses the polarity of opinions. This treatment follows two axes. First, we developed rules to deal with the very frequent cases of negation in telegraphic style (“*Not quite as fast as HP says*”), to deal with the interaction between quantification and negation (“*I never had so many problems*”) or double negation (“*I cannot say I do not appreciate this printer*”). Then, we developed rules reversing the polarity according to the scope of the negation, built on top of the sentiment relations extracted in the previous processing step. These rules allow to affect the proper polarity to examples like “*I really do not like this feature*”; “*This is not a good photo printer*”.

<sup>3</sup> About 100539584 words.

<sup>4</sup> 4 coordination types \* 3 classes of adjectives.

In addition, a layer of domain specific rules has been developed, to handle expressions such as:

- *It is easy to set up.* [Positive]
- *It uses a lot of ink.* [Negative]

which are specific to the domain: generally "easy" can not be considered as a positive word ("It is easy to lose money" has a negative connotation). However, the association [easy + to + verb indicating a functional characteristic of the printer] expresses a positive opinion. Similarly, a verb of consumption ("consume", "use", "eat" ...) with a consumable item of the printer ("ink", "paper", "cartridge",...) as direct object denotes a negative opinion. About twenty such domain-dependent rules, based on the semantic features encoded in the domain-dependent lexicon, have been developed.

Finally, a few rules take into account the structure of the reviews of the site "Epinion", using some structural clues ("*Cons*", "*Pro*", "*Recommended*"...) to calculate the opinions. For example, "*Cons: none*" indicate a very positive opinion.

The set of sentiment-related rules is now fairly stable. We must continue our development efforts to address the problems of modality, comparisons, and integrate a coreference module to the system.

#### 4.7 Evaluation

As we do not have a corpus of annotated printer reviews, in terms of positive or negative opinion relations, we used the structure of the "Epinion" reviews, in order to assess the performance of our system in a "coarse" way: since the user explicitly states whether he recommends or not the printer, we consider the corpus as annotated for classification. We then use the relations of opinions extracted by our system to train a SVM binary classifier (SVMlight, Joachims 1999) in order to classify the reviews as positive (i.e. recommended) or negative (i.e. not recommended). The experimental setup consists in 313 reviews extracted randomly from the initial corpus to train the SVM classifier, 293 reviews extracted randomly for validation and 2735 reviews extracted randomly for testing. The SVM features are the relations of opinion on a given target and their values are the frequencies of these relations, e.g. OPINION-POSITIVE-on-SPEED:2, OPINION-NEGATIVE-on-PRICE:1, etc. We calculated the baseline using simple keyword-based SVM classification, without syntactic analysis. We therefore evaluate the system ability to classify documents according to an overall opinion.

Table 1 shows the results obtained on the test corpus (2735 test reviews).

	Favorable reviews	Unfavorable reviews	Total reviews
Number	2066	669	2735
Classified as "positive"	1996	128	2124
Classified as "negative"	70	541	611
Accuracy	97%	81%	93%
Baseline Accuracy	87%	51%	79%

**Table 1: Coarse evaluation on the printer corpus**

These results are very encouraging since they are in line with state of the art results, obtained for similar classification tasks, cf. (Pang et al. 2002) or (Paroubek et al. 2007).

To validate the quality of our general grammar, and to assess its portability, we conducted a similar second evaluation, in the domain of movie reviews. For this experiment we only use the general opinion extraction grammar and no specialized grammar. The experimental conditions are otherwise exactly the same as before. The results are given in Table II. Results are also very satisfactory. The slight difference is probably due to the lack of specialized grammar rules.

	Favorable Reviews	Unfavorable reviews	Total
Number	1343	420	1763
Classified as positive	1281	122	1403
Classified as negative	62	298	360
Accuracy	95%	71%	89%
Baseline Accuracy	83%	46%	74%

**Table 2: Coarse evaluation on the movie corpus**

In both cases, the system shows the same trend, it has more difficulty to properly classify unfavorable reviews. There are several explanations for this: first, as the polar vocabulary is partly extracted from the reviews themselves, and as the proportion of unfavorable reviews is small, there might be a coverage problem for the negative vocabulary. Moreover, it seems that the authors use a different discourse whether they recommend or not a product or a film: a brief analysis of the errors on unfavorable reviews shows that authors tend to use comparison with other products or movies they have preferred. For now, our system deals only partially with comparison and does not yet integrate a coreference module,

many positive opinions about other products or films are incorrectly credited to the account of the main topic of the review.

In conclusion, we are aware that this is a preliminary evaluation, since our final goal is fine-grained opinion extraction. We plan to make another evaluation, using a reference corpus manually annotated with sentiment relations.

## 5 Conclusion

In this paper, we present a system extracting opinions on online product reviews. This system uses deep syntactic relations provided by a robust dependency parser in order to extract sentiment relationships. These relationships are intended to instantiate a formal model of representation of the opinions. We have developed semi-automatically a dedicated lexicon associating polarities and semantic features to words. We have then developed a set of generic and domain-dependant hand-crafted rules for extracting relations of opinions. The evaluation of the performances of the system on coarse-grained classification of reviews is very encouraging. We will pursue the developments in order to take account complex linguistic phenomena not yet well covered, namely comparative constructions, modality and coreference. The coreference module pre-exists but requires modifications for its integration within our system. We then plan to conduct a fine-grained evaluation.

## References

- A. Agarwal, P. Bhattacharya. 2006. Augmenting Wordnet with Polarity Information on Adjectives. *3rd International Wordnet Conference*, Jeju Island, Korea, South Jeju (Seogwipo)
- Salah Ait-Mokthar, Jean-Pierre Chanod. 2002, Robustness beyond Shallowness: Incremental Dependency Parsing. *Special Issue of NLE Journal*.
- E. Charton, R. Acuna-Agosity. 2007. Quel modèle pour détecter une opinion? Trois propositions pour généraliser l'extraction d'une idée dans un corpus, *Deft'07*, Grenoble.
- X. Ding, B. Liu, P.Yu. 2008. A Holistic Lexicon-Based Approach to Opinion Mining. *Proceedings of the international conference on Web search and web data mining*, WSDM '08, ACM.
- E. Dubreil, M. Vernier, L. Monceaux, B. Daille. 2008. Annotating Opinion – Evaluation of Blogs, Workshop on LREC 2008 Conference, Sentiment Analysis: Metaphor, Ontology and Terminology (EMOT-08), Marrakech, Morocco.
- A. Esuli, F. Sebastiani. 2006. SENTIWORDNET: A Publicly Available Lexical Resource for Opinion Mining. In *5th Conference on Language Resources and Evaluation (LREC'06)*, pp. 417-422.
- C. Hagège, C. Roux. 2002. Entre syntaxe et sémantique: normalisation de la sortie de l'analyse syntaxique en vue de l'amélioration de l'extraction d'information à partir de textes. *TALN 2003*, Batsur-Mer, France, 11-14 Juin.
- M. Hu, B. Liu. 2004. Mining and summarizing customer reviews. *ACM SIGKDD International Conference on Knowledge Discovery & Data Mining (KDD-2004)*, Seattle, Washington, USA.
- T. Joachims. 1999: Making large-Scale SVM Learning Practical. *Advances in Kernel Methods - Support Vector Learning*, B. Schölkopf and C. Burges and A. Smola (ed.), MIT Press.
- V. Hatzivassiloglou, K. McKeown. 1997. Predicting the semantic orientation of adjectives. In *Proceedings of the Joint ACL/EACL Conference*, pp. 174–181.
- B. Liu. 2010. Sentiment Analysis and Subjectivity, Chapter of *Handbook of Natural Language Processing*, 2<sup>nd</sup> edition.
- L. Monceaux, B. Daille, E. Dubreil 2009. Catégorisation des évaluations dans un corpus de blogs multi-domaine. *Revue des nouvelles technologies de l'information (RNTI)*.
- B. Pang, L. Lee, S. Vaithyanathan. 2002. Thumbs up? Sentiment classification using machine learning techniques. In *Proceedings of the Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 79–86.
- P. Paroubek Berthelin J.B., El Ayari S., Grouin C., Heitz T., Hurault-Plantet M., Jardino M., Khalis Z., Lastes M. 2007. Résultats de l'édition 2007 du DÉfi Fouille de Textes, *Deft'07*, Grenoble.
- A. Popescu, O. Etzioni. 2005. Extracting product features and opinions from reviews. *Actes de Conference on Empirical Methods in Natural Language Processing (EMNLP)*.
- S. Vegnaduzzo. 2004. Acquisition of subjective adjectives with limited resources. *Actes de AAAI spring symposium on exploring attitude and affect in text: Theories and Applications*, Stanford, US.
- J. Yi, T. Nasukawa, A. Valerio, H. Zhang. 2003. Sentiment Analyzer: Extracting Sentiments about a Given Topic Using natural Language Processing Techniques. *ICDM'03: 3<sup>rd</sup> IEEE International Conference on Data Mining*, pp. 427.
- Hong Yu, Vassileos Hazivassiloglou. 2003. Towards Answering Opinion Questions: Separating Facts from Opinions and Identifying the Polarity of Opinions, *EMNLP 2003*, Sapporo, Japan.