Extracting Opinions, Opinion Holders, and Topics Expressed in Online News Media Text

Soo-Min Kim and Eduard Hovy
USC Information Sciences Institute
4676 Admiralty Way
Marina del Rey, CA 90292-6695
{skim, hovy}@ISI.EDU

Abstract

This paper presents a method for identifying an opinion with its holder and topic, given a sentence from online news media texts. We introduce an approach of exploiting the semantic structure of a sentence, anchored to an opinion bearing verb or adjective. This method uses semantic role labeling as an intermediate step to label an opinion holder and topic using data from FrameNet. We decompose our task into three phases: identifying an opinion-bearing word, labeling semantic roles related to the word in the sentence, and then finding the holder and the topic of the opinion word among the labeled semantic roles. For a broader coverage, we also employ a clustering technique to predict the most probable frame for a word which is not defined in FrameNet. Our experimental results show that our system performs significantly better than the baseline.

1 Introduction

The challenge of automatically identifying opinions in text automatically has been the focus of attention in recent years in many different domains such as news articles and product reviews. Various approaches have been adopted in subjectivity detection, semantic orientation detection, review classification and review mining. Despite the successes in identifying opinion expressions and subjective words/phrases, there has been less achievement on the factors closely related to subjectivity and polarity, such as opinion holder, topic of opinion, and inter-topic/inter-opinion relationships. This paper addresses the problem of identifying not only opinions in text but also

holders and topics of opinions from online news articles.

Identifying opinion holders is important especially in news articles. Unlike product reviews in which most opinions expressed in a review are likely to be opinions of the author of the review, news articles contain different opinions of different opinion holders (e.g. people, organizations, and countries). By grouping opinion holders of different stance on diverse social and political issues, we can have a better understanding of the relationships among countries or among organizations.

An opinion topic can be considered as an object an opinion is about. In product reviews, for example, opinion topics are often the product itself or its specific features, such as design and quality (e.g. “*I like the design of iPod video*”, “*The sound quality is amazing*”). In news articles, opinion topics can be social issues, government’s acts, new events, or someone’s opinions. (e.g., “*Democrats in Congress accused vice president Dick Cheney’s shooting accident*”, “*Shiite leaders accused Sunnis of a mass killing of Shiites in Madaen, south of Baghdad*.”)

As for opinion topic identification, little research has been conducted, and only in a very limited domain, product reviews. In most approaches in product review mining, given a product (e.g. mp3 player), its frequently mentioned features (e.g. sound, screen, and design) are first collected and then used as anchor points. In this study, we extract opinion topics from news articles. Also, we do not pre-limit topics in advance. We first identify an opinion and then find its holder and topic. We define *holder* as an entity who holds an opinion, and *topic*, as what the opinion is about.

In this paper, we propose a novel method that employs Semantic Role Labeling, a task of identifying semantic roles given a sentence. We de-

compose the overall task into the following steps:

- Identify opinions.
- Label semantic roles related to the opinions.
- Find holders and topics of opinions among the identified semantic roles.
- Store <opinion, holder, topic> triples into a database.

In this paper, we focus on the first three subtasks.

The main contribution of this paper is to present a method that identifies not only opinion holders but also opinion topics. To achieve this goal, we utilize FrameNet data by mapping target words to opinion-bearing words and mapping semantic roles to holders and topics, and then use them for system training. We demonstrate that investigating semantic relations between an opinion and its holder and topic is crucial in opinion holder and topic identification.

This paper is organized as follows: Section 2 briefly introduces related work both in sentiment analysis and semantic role labeling. Section 3 describes our approach for identifying opinions and labeling holders and topics by utilizing FrameNet¹ data for our task. Section 4 reports our experiments and results with discussions and finally Section 5 concludes.

2 Related Work

This section reviews previous works in both sentiment detection and semantic role labeling.

2.1 Subjectivity and Sentiment Detection

Subjectivity detection is the task of identifying subjective words, expressions, and sentences (Wiebe *et al.*, 1999; Hatzivassiloglou and Wiebe, 2000; Riloff *et al.*, 2003). Identifying subjectivity helps separate opinions from fact, which may be useful in question answering, summarization, etc. Sentiment detection is the task of determining positive or negative sentiment of words (Hatzivassiloglou and McKeown, 1997; Turney, 2002; Esuli and Sebastiani, 2005), phrases and sentences (Kim and Hovy, 2004; Wilson *et al.*, 2005), or documents (Pang *et al.*, 2002; Turney, 2002).

Building on this work, more sophisticated problems such as opinion holder identification have also been studied. (Bethard *et al.*, 2004) identify opinion propositions and holders. Their

work is similar to ours but different because their opinion is restricted to propositional opinion and mostly to verbs. Another related works are (Choi *et al.*, 2005; Kim and Hovy, 2005). Both of them use the MPQA corpus² but they only identify opinion holders, not topics.

As for opinion topic identification, little research has been conducted, and only in a very limited domain, product reviews. (Hu and Liu, 2004; Popescu and Etzioni, 2005) present product mining algorithms with extracting certain product features given specific product types. Our paper aims at extracting topics of opinion in general news media text.

2.2 Semantic Role Labeling

Semantic role labeling is the task of identifying semantic roles such as Agent, Patient, Speaker, or Topic, in a sentence. A statistical approach for semantic role labeling was introduced by (Gildea and Jurafsky, 2002). Their system learned semantic relationship among constituents in a sentence from FrameNet, a large corpus of semantically hand-annotated data. The FrameNet annotation scheme is based on Frame Semantics (Fillmore, 1976). *Frames* are defined as “schematic representations of situations involving various *frame elements* such as participants, props, and other conceptual roles.” For example, given a sentence “Jack built a new house out of bricks”, a semantic role labeling system should identify the roles for the verb *built* such as “[_{Agent} Jack] built [_{Created_entity} a new house] [_{Component} out of bricks]”³. In our study, we build a semantic role labeling system as an intermediate step to label opinion holders and topics by training it on opinion-bearing frames and their frame elements in FrameNet.

3 Finding Opinions and Their Holders and Topics

For the goal of this study, extracting opinions from news media texts with their holders and topics, we utilize FrameNet data. The basic idea of our approach is to explore how an opinion holder and a topic are semantically related to an opinion bearing word in a sentence. Given a sentence and an opinion bearing word, our method identifies frame elements in the sentence and

¹ <http://framenet.icsi.berkeley.edu/>

² <http://www.cs.pitt.edu/~wiebe/pubs/ardasummer02/>

³ The verb “build” is defined under the frame “Building” in which Agent, Created_entity, and Components are defined as frame elements.

Sentence: On Dec. 7, the Islamic Conference Organization (ICO) denounced the affair as a crime.

Figure 1: An overview of our algorithm

searches which frame element corresponds to the opinion holder and which to the topic. The example in Figure 1 shows the intuition of our algorithm.

We decompose our task in 3 subtasks: (1) collect opinion words and opinion-related frames, (2) semantic role labeling for those frames, and (3) finally map semantic roles to holder and topic. Following subsections describe each subtask.

3.1 Opinion Words and Related Frames

We describe the subtask of collecting opinion words and related frames in 3 phases.

Phase 1: Collect Opinion Words

In this study, we consider an opinion-bearing (positive/negative) word is a key indicator of an opinion. Therefore, we first identify opinion-bearing word from a given sentence and extract its holder and topic. Since previous studies indicate that opinion-bearing verbs and adjectives are especially efficient for opinion identification, we focus on creating a set of opinion-bearing verbs and adjectives. We annotated 1860 adjectives and 2011 verbs⁴ by classifying them into positive, negative, and neutral classes. Words in the positive class carry positive valence whereas

⁴ These were randomly selected from 8011 English verbs and 19748 English adjectives.

Table 1: Example of opinion related frames and lexical units

Frame name	Lexical units	Frame elements
Desiring	want, wish, hope, eager, desire, interested,	Event, Experiencer, Location_of_event
Emotion_directed	agitated, amused, anguish, ashamed, angry, annoyed,	Event, Topic Experiencer, Expressor,
Mental_property	absurd, brilliant, careless, crazy, cunning, foolish	Behavior, Protagonist, Domain, Degree
Subject_stimulus	delightful, amazing, annoying, amusing, aggravating,	Stimulus, Degree Experiencer, Circumstances,

those in negative class carry negative valence. Words that are not opinion-bearing are classified as neutral.

Note that in our study we treat word sentiment classification as a three-way classification problem instead of a two-way classification problem (i.e. positive and negative). By adding the third class, neutral, we can prevent the classifier assigning either positive or negative sentiment to weak opinion-bearing word. For example, the word “central” that Hatzivassiloglou and McKeown (1997) marked as a positive adjective is not classified as positive by our system. Instead we mark it as “neutral”, since it is a weak clue for an opinion. For the same reason, we did not consider “able” classified as a positive word by General Inquirer⁵, a sentiment word lexicon, as a positive opinion indicator. Finally, we collected 69 positive and 151 negative verbs and 199 positive and 304 negative adjectives.

Phase 2: Find Opinion-related Frames

We collected frames related to opinion words from the FrameNet corpus. We used FrameNet II (Baker *et al.*, 2003) which contains 450 semantic frames and more than 3000 frame elements (FE). A frame consists of lexical items, called *Lexical Unit* (LU), and related frame elements. For instance, LUs in *ATTACK* frame are verbs such as *assail*, *assault*, and *attack*, and nouns such as *invasion*, *raid*, and *strike*. FrameNet II contains

⁵ <http://www.wjh.harvard.edu/~inquirer/homecat.htm>

approximately 7500 lexical units and over 100,000 annotated sentences.

For each word in our opinion word set described in Phase 1, we find a frame to which the word belongs. 49 frames for verbs and 43 frames for adjectives are collected. Table 1 shows examples of selected frames with some of the lexical units those frames cover. For example, our system found the frame *Desiring* from opinion-bearing words *want*, *wish*, *hope*, etc. Finally, we collected 8256 and 11877 sentences related to selected opinion bearing frames for verbs and adjectives respectively.

Phase 3: FrameNet expansion

Even though Phase 2 searches for a correlated frame for each verb and adjective in our opinion-bearing word list, not all of them are defined in FrameNet data. Some words such as *criticize* and *harass* in our list have associated frames (*Case 1*), whereas others such as *vilify* and *maltreat* do not have those (*Case 2*). For a word in Case 2, we use a clustering algorithms CBC (Clustering By Committee) to predict the closest (most reasonable) frame of undefined word from existing frames. CBC (Pantel and Lin, 2002) was developed based on the distributional hypothesis (Harris, 1954) that words which occur in the same contexts tend to be similar. Using CBC, for example, our clustering module computes lexical similarity between the word *vilify* in Case 2 and all words in Case 1. Then it picks *criticize* as a similar word, so that we can use for *vilify* the frame *Judgment_communication* to which *criticize* belongs and all frame elements defined under *Judgment_communication*.

3.2 Semantic Role Labeling

To find a potential holder and topic of an opinion word in a sentence, we first label semantic roles in a sentence.

Modeling: We follow the statistical approaches for semantic role labeling (Gildea and Jurafsky, 2002; Fleischman *et. al*, 2003) which separate the task into two steps: identify candidates of frame elements (Step 1) and assign semantic roles for those candidates (Step 2). Like their intuition, we treated both steps as classification problems. We first collected all constituents of the given sentence by parsing it using the Charniak parser. Then, in Step 1, we classified candidate constituents of frame elements from non-candidates. In Step 2, each selected candidate was thus classified into one of frame ele-

Table 2: Features used for our semantic role labeling model.

Feature	Description
target word	A predicate whose meaning represents the frame (a verb or an adjective in our task)
phrase type	Syntactic type of the frame element (e.g. NP, PP)
head word	Syntactic head of the frame element phrase
parse tree path	A path between the frame element and target word in the parse tree
position	Whether the element phrase occurs <i>before</i> or <i>after</i> the target word
voice	The voice of the sentence (<i>active</i> or <i>passive</i>)
frame name	one of our opinion-related frames

ment types (e.g. Stimulus, Degree, Experiencer, etc.). As a learning algorithm for our classification model, we used Maximum Entropy (Berger *et al.*, 1996). For system development, we used MEGA model optimization package⁶, an implementation of ME models.

Data: We collected 8256 and 11877 sentences which were associated to opinion bearing frames for verbs and adjectives from FrameNet annotation data. Each sentence in our dataset contained a frame name, a target predicate (a word whose meaning represents aspects of the frame), and frame elements labeled with element types. We divided the data into 90% for training and 10% for test.

Features used: Table 2 describes features that we used for our classification model. The target word is an opinion-bearing verb or adjective which is associated to a frame. We used the Charniak parser to get a phrase type feature of a frame element and the parse tree path feature. We determined a head word of a phrase by an algorithm using a tree head table⁷, position feature by the order of surface words of a frame element and the target word, and the voice feature by a simple pattern. Frame name for a target

⁶ <http://www.isi.edu/~hdaume/megam/index.html>

⁷ <http://people.csail.mit.edu/mcollins/papers/heads>

Table 3. Precision (P), Recall (R), and F -score (F) of Topic and Holder identification for opinion verbs (V) and adjectives (A) on Testset 1.

	Topic			Holder		
	P (%)	R (%)	F (%)	P (%)	R (%)	F (%)
V	69.1	64.0	66.5	81.9	75.7	78.7
A	67.5	73.4	70.3	66.2	77.9	71.6

Table 4. Baseline system on Testset 1.

	Topic			Holder		
	P (%)	R (%)	F (%)	P (%)	R (%)	F (%)
V	85.5	18.5	30.4	73.7	46.4	56.9
A	68.2	26.5	38.2	12.0	49.1	19.3

word was selected by methods described in Phase 2 and Phase 3 in Subsection 3.1.

3.3 Map Semantic Roles to Holder and Topic

After identifying frame elements in a sentence, our system finally selects holder and topic from those frame elements. In the example in Table 1, the frame “*Desiring*” has frame elements such as Event (“*The change that the Experiencer would like to see*”), Experiencer (“*the person or sentient being who wishes for the Event to occur*”), Location_of_event (“*the place involved in the desired Event*”), Focal_participant (“*entity that the Experiencer wishes to be affected by some Event*”). Among these FEs, we can consider that Experiencer can be a holder and Focal_participant can be a topic (if any exists in a sentence). We manually built a mapping table to map FEs to holder or topic using as support the FE definitions in each opinion related frame and the annotated sample sentences.

4 Experimental Results

The goal of our experiment is first, to see how our holder and topic labeling system works on the FrameNet data, and second, to examine how it performs on online news media text. The first data set (Testset 1) consists of 10% of data described in Subsection 3.2 and the second (Testset 2) is manually annotated by 2 humans. (see Subsection 4.2). We report experimental results for both test sets.

4.1 Experiments on Testset 1

Gold Standard: In total, Testset 1 contains 2028 annotated sentences collected from FrameNet data set. (834 from frames related to opinion verb and 1194 from opinion adjectives) We measure the system performance using precision (the percentage of correct holders/topics among system’s labeling results), recall (the percentage of correct holders/topics that system retrieved), and F -score.

Baseline: For the baseline system, we applied two different algorithms for sentences which have opinion-bearing verbs as target words and for those that have opinion-bearing adjectives as target words. For **verbs**, baseline system labeled a subject of a verb as a holder and an object as a topic. (e.g. “[_{holder} He] *condemned* [_{topic} the lawyer].”) For **adjectives**, the baseline marked the subject of a predicate adjective as a holder (e.g. “[_{holder} I] *was happy*”). For the topics of adjectives, the baseline picks a modified word if the target adjective is a modifier (e.g. “That was a *stupid* [_{topic} mistake].”) and a subject word if the adjective is a predicate. ([_{topic} The view] *is breathtaking* in January.)

Result: Table 3 and 4 show evaluation results of our system and the baseline system respectively. Our system performed much better than the baseline system in identifying topic and holder for both sets of sentences with verb target words and those with adjectives. Especially in recognizing topics of target opinion-bearing words, our system improved F -score from 30.4% to 66.5% for verb target words and from 38.2% to 70.3% for adjectives. It was interesting to see that the intuition that “*A subject of opinion-bearing verb is a holder and an object is a topic*” which we applied for the baseline achieved relatively good F -score (56.9%). However, our system obtained much higher F -score (78.7%). Holder identification task achieved higher F -score than topic identification which implies that identifying topics of opinion is a harder task.

We believe that there are many complicated semantic relations between opinion-bearing words and their holders and topics that simple relations such as subject and object relations are not able to capture. For example, in a sentence “Her letter *upset* me”, simply looking for the subjective and objective of the verb *upset* is not enough to recognize the holder and topic. It is necessary to see a deeper level of semantic rela-

Table 5. Opinion-bearing sentence identification on Testset 2. (P: precision, R: recall, F: F-score, A: Accuracy, H1: Human1, H2: Human2)

	P (%)	R (%)	F (%)	A (%)
H1	56.9	67.4	61.7	64.0
H2	43.1	57.9	49.4	55.0

Table 6: Results of Topic and Holder identification on Testset 2. (Sys: our system, BL: baseline)

		Topic			Holder		
		P(%)	R(%)	F(%)	P(%)	R(%)	F(%)
Sys	H1	64.7	20.8	31.5	47.9	34.0	39.8
	H2	58.8	7.1	12.7	36.6	26.2	30.5
BL	H1	12.5	9.4	10.7	20.0	28.3	23.4
	H2	23.2	7.1	10.9	14.0	19.0	16.1

tions: “Her letter” is a stimulus and “me” is an experiencer of the verb *upset*.

4.2 Experiments on Testset 2

Gold Standard: Two humans⁸ annotated 100 sentences randomly selected from news media texts. Those news data is collected from online news sources such as The New York Times, UN Office for the Coordination of Humanitarian Affairs, and BBC News⁹, which contain articles about various international affairs. Annotators identified opinion-bearing sentences with marking opinion word with its holder and topic if they existed. The inter-annotator agreement in identifying opinion sentences was 82%.

Baseline: In order to identify opinion-bearing sentences for our baseline system, we used the opinion-bearing word set introduced in Phase 1 in Subsection 3.1. If a sentence contains an opinion-bearing verb or adjective, the baseline system started looking for its holder and topic. For holder and topic identification, we applied the

same baseline algorithm as described in Subsection 4.1 to Testset 2.

Result: Note that Testset 1 was collected from sentences of opinion-related frames in FrameNet and therefore all sentences in the set contained either opinion-bearing verb or adjective. (i.e. All sentences are opinion-bearing) However, sentences in Testset 2 were randomly collected from online news media pages and therefore not all of them are opinion-bearing. We first evaluated the task of opinion-bearing sentence identification. Table 5 shows the system results. When we mark all sentences as opinion-bearing, it achieved 43% and 38% of accuracy for the annotation result of Human1 and Human2 respectively. Our system performance (64% and 55%) is comparable with the unique assignment.

We measured the holder and topic identification system with precision, recall, and *F*-score. As we can see from Table 6, our system achieved much higher precision than the baseline system for both Topic and Holder identification tasks. However, we admit that there is still a lot of room for improvement.

The system achieved higher precision for topic identification, whereas it achieved higher recall for holder identification. In overall, our system attained higher *F*-score in holder identification task, including the baseline system. Based on *F*-score, we believe that identifying topics of opinion is much more difficult than identifying holders. It was interesting to see the same phenomenon that the baseline system mainly assuming that subject and object of a sentence are likely to be opinion holder and topic, achieved lower scores for both holder and topic identification tasks in Testset 2 as in Testset 1. This implies that more sophisticated analysis of the relationship between opinion words (e.g. verbs and adjectives) and their topics and holders is crucial.

4.3 Difficulties in evaluation

We observed several difficulties in evaluating holder and topic identification. First, the boundary of an entity of holder or topic can be flexible. For example, in sentence “Senator Titus Olupitan who sponsored the bill wants the permission.”, not only “Senator Titus Olupitan” but also “Senator Titus Olupitan who sponsored the bill” is an eligible answer. Second, some correct holders and topics which our system found were evaluated wrong even if they referred the same entities in the gold standard because human annotators marked only one of them as an answer.

⁸ We refer them as Human1 and Human2 for the rest of this paper.

⁹ www.nytimes.com, www.irinnews.org, and www.bbc.co.uk

In the future, we need more annotated data for improved evaluation.

5 Conclusion and Future Work

This paper presented a methodology to identify an opinion with its holder and topic given a sentence in online news media texts. We introduced an approach of exploiting semantic structure of a sentence, anchored to an opinion bearing verb or adjective. This method uses semantic role labeling as an intermediate step to label an opinion holder and topic using FrameNet data. Our method first identifies an opinion-bearing word, labels semantic roles related to the word in the sentence, and then finds a holder and a topic of the opinion word among labeled semantic roles.

There has been little previous study in identifying opinion holders and topics partly because it requires a great amount of annotated data. To overcome this barrier, we utilized FrameNet data by mapping target words to opinion-bearing words and mapping semantic roles to holders and topics. However, FrameNet has a limited number of words in its annotated corpus. For a broader coverage, we used a clustering technique to predict a most probable frame for an unseen word.

Our experimental results showed that our system performs significantly better than the baseline. The baseline system results imply that opinion holder and topic identification is a hard task. We believe that there are many complicated semantic relations between opinion-bearing words and their holders and topics which simple relations such as subject and object relations are not able to capture.

In the future, we plan to extend our list of opinion-bearing verbs and adjectives so that we can discover and apply more opinion-related frames. Also, it would be interesting to see how other types of part of speech such as adverbs and nouns affect the performance of the system.

Reference

- Baker, Collin F. and Hiroaki Sato. 2003. The FrameNet Data and Software. *Poster and Demonstration at Association for Computational Linguistics*. Sapporo, Japan.
- Berger, Adam, Stephen Della Pietra, and Vincent Della Pietra. 1996. A maximum entropy approach to natural language processing, *Computational Linguistics*, (22-1).
- Bethard, Steven, Hong Yu, Ashley Thornton, Vasileios Hatzivassiloglou, and Dan Jurafsky. 2004. Automatic Extraction of Opinion Propositions and their Holders, *AAAI Spring Symposium on Exploring Attitude and Affect in Text: Theories and Applications*.
- Choi, Y., Cardie, C., Riloff, E., and Patwardhan, S. 2005. Identifying Sources of Opinions with Conditional Random Fields and Extraction Patterns. *Proceedings of HLT/EMNLP-05*.
- Esuli, Andrea and Fabrizio Sebastiani. 2005. Determining the semantic orientation of terms through gloss classification. *Proceedings of CIKM-05, 14th ACM International Conference on Information and Knowledge Management*, Bremen, DE, pp. 617-624.
- Fillmore, C. Frame semantics and the nature of language. 1976. In *Annals of the New York Academy of Sciences: Conferences on the Origin and Development of Language and Speech*, Volume 280: 20-32.
- Fleischman, Michael, Namhee Kwon, and Eduard Hovy. 2003. Maximum Entropy Models for FrameNet Classification. *Proceedings of EMNLP*, Sapporo, Japan.
- Gildea, D. and Jurafsky, D. Automatic Labeling of semantic roles. 2002. In *Computational Linguistics*. 28(3), 245-288.
- Harris, Zellig. 1954. *Distributional structure*. Word, 10(23) :146--162.
- Hatzivassiloglou, Vasileios and Kathleen McKeown. 1997. Predicting the Semantic Orientation of Adjectives. *Proceedings of 35th Annual Meeting of the Assoc. for Computational Linguistics (ACL-97)*: 174-181
- Hatzivassiloglou, Vasileios and Wiebe, Janyce. 2000. Effects of Adjective Orientation and Gradability on Sentence Subjectivity. *Proceedings of International Conference on Computational Linguistics (COLING-2000)*. Saarbrücken, Germany.
- Hu, Minqing and Bing Liu. 2004. Mining and summarizing customer reviews". *Proceedings of the ACM SIGKDD International Conference on Knowledge Discovery & Data Mining (KDD-2004)*, Seattle, Washington, USA.
- Kim, Soo-Min and Eduard Hovy. 2004. Determining the Sentiment of Opinions. *Proceedings of COLING-04*. pp. 1367-1373. Geneva, Switzerland.
- Kim, Soo-Min and Eduard Hovy. 2005. Identifying Opinion Holders for Question Answering in Opinion Texts. *Proceedings of AAAI-05 Workshop on Question Answering in Restricted Domains*
- Pang, Bo, Lillian Lee, and Shivakumar Vaithyanathan. 2002. Thumbs up? Sentiment Classification using Machine Learning Techniques, *Proceedings of EMNLP-2002*.

- Pantel, Patrick and Dekang Lin. 2002. Discovering Word Senses from Text. *Proceedings of ACM Conference on Knowledge Discovery and Data Mining. (KDD-02)*. pp. 613-619. Edmonton, Canada.
- Popescu, Ana-Maria and Oren Etzioni. 2005. Extracting Product Features and Opinions from Reviews , *Proceedings of HLT-EMNLP 2005*.
- Riloff, Ellen, Janyce Wiebe, and Theresa Wilson. 2003. Learning Subjective Nouns Using Extraction Pattern Bootstrapping. *Proceedings of Seventh Conference on Natural Language Learning (CoNLL-03)*. ACL SIGNLL. Pages 25-32.
- Turney, Peter D. 2002. Thumbs up or thumbs down? Semantic orientation applied to unsupervised classification of reviews, *Proceedings of ACL-02*, Philadelphia, Pennsylvania, 417-424
- Wiebe, Janyce, Bruce M., Rebecca F., and Thomas P. O'Hara. 1999. Development and use of a gold standard data set for subjectivity classifications. *Proceedings of ACL-99*. University of Maryland, June, pp. 246-253.
- Wilson, Theresa, Janyce Wiebe, and Paul Hoffmann. 2005. Recognizing Contextual Polarity in Phrase-Level Sentiment Analysis. *Proceedings of HLT/EMNLP 2005*, Vancouver, Canada