Extraction of Opposite Sentiments in Classified Free Format Text Reviews

Dong (Haoyuan) Li[1], Anne Laurent[2], Mathieu Roche[2], and Pascal Poncelet[1]

[1] LGI2P - École des Mines d'Alès, Parc Scientifique G. Besse, 30035 Nîmes, France
E-mail: {Haoyuan.Li, Pascal.Poncelet}@ema.fr
[2] LIRMM - Université Montpellier II, 161 rue Ada, 34392 Montpellier, France
E-mail: {laurent, mroche}@lirmm.fr

Abstract. Most of the previous approaches in opinion mining focus on the classifications of opinion polarities, *positive* or *negative*, expressed in customer reviews. In this paper, we present the problem of extracting contextual opposite sentiments in classified free format text reviews. We adapt the sequence data model to text mining with Part-of-Speech tags, and then we propose a belief-driven approach for extracting contextual opposite sentiments as unexpected sequences with respect to the opinion polarity of reviews. We conclude by detailing our experimental results on free format text movie review data.

1 Introduction

Opinion mining received much attention in finding personal opinions from user generated contents, such as customer reviews, forums, discussion groups, and blogs, where most of the previous approaches concentrate on the classifications of opinion polarities, *positive* or *negative*, in free format text reviews [9, 14, 2, 16, 5, 8, 15]. Although the positive-negative classifications are determinative, the opposite sentiments expressed in classified reviews, within the context of topic, become more and more interesting for decision making.

For instance, about a notebook computer, a positive review may contain the sentences like "however the graphics performance is not enough", or in a negative review we may also find "anyway this notebook is beautiful", and such critiques are important to improve the quality of products. However, even sentence-level sentiment classifications [2, 5, 15] extract the sentences that express the opposite sentiment with the positive-negative connotations different to document-level opinion polarity, such sentences may be not within the same context of the topic about the review.

In this paper, we present a belief-driven approach for extracting contextual opposite sentiments in classified free format text reviews. A training-extracting process is considered: Given a topic context, first a sequential pattern mining algorithm is applied to a set of classified training reviews, in order to generate the contextual models of opinion polarity with respect to current topic. Then, from such contextual models, a belief base is constructed to represent the opinion polarity by using a dictionary of antonyms[3] of the adjectives contained in the contextual models. Finally, the unexpected sequence mining process proposed in our previous work [7] is performed to the target reviews for extracting all sentences that contradict the belief base, which stand for the contextual opposite sentiments.

The rest of this paper is structured as follows. Section 2 introduces the related work. In Sect. 3 we first formalize the data model, then propose the contextual models and belief base on sentiment polarities, and then present the process of extracting contextual opposite sentiments. Section 4 details our experiments on positive-negative movie-review data. Section 5 is a short conclusion.

2 Related Work

Opinion mining in free format text contents is closely connected with the Natural Language Processing (NLP) problems, where the positive or negative connotation can be annotated by the subjective terms at document-level [9, 14, 2, 8] or sentence-level [2, 16, 5, 15].

In [2], a term frequencies based scoring system is proposed for determining both document- and sentence-level sentiment polarities. The approach proposed in [5] extracts the features of products contained in customer reviews with positive-negative polarities, which can be considered as a sentence-level opinion classification. Compared to our approach, [16] proposed a model for classifying opinion sentences as positive or negative in terms of the main perspective expressed in the opinion of document, which identifies facts and opinions, and can be considered as a contextual approach. Another contextual notion, so called contextual polarity, is proposed in [15], which is determined by the dependency tree of the structure of sentences; in our approach, we use sequential pattern mining to determine the frequent structures of contextual models for sentiment polarity.

Actually, the opinion polarities are often given by the adjectives [3, 13]. We use WordNet [4] for determining the antonyms of adjectives required for constructing the belief base, which has been used in many NLP and opinion mining approaches. For instance, in the proposal of [6], WordNet is also applied for detecting the semantic orientation of adjectives.

3 Extracting Contextual Opposite Sentiments

3.1 Data Model

We are given a set of free format text reviews that have been already classified into positive-negative opinion polarities. Each review consists in an ordered list of sentences, and each sentence consists in an ordered list of words. In order to

[3] The antonym dictionary is based on the WordNet project, which can be found at http://wordnet.princeton.edu/.

involve the words in the context of reviews, the Part-of-Speech tag (PoS tag) introduced in the TreeTagger approach [11] is considered, and a list of such PoS tags is available in [10]. With respect to this list, we do not consider the difference between the different tags of the adjectives (J instead of JJ, JJR and JJS), of the adverbs (R in stead of R, RB, RBR and RBS), of the nouns (N in stead of N, NN, NNS, NP and NPS), and of the verbs (V in stead of V, VB, VBD, VBG, VBN, VBP and VBZ).

A *word*, denoted as w, is a lemma associated with a simplified PoS tag. For example $(be|V)$ is a word where *be* is a lemma and V is the base tag standing for the verbs. Without loss of generality, we use the wild-card $*$ and a simplified PoS tag for denoting a generalized vocabulary. For example, $(*|V)$ denotes a vocabulary that is a verb. Especially, we use (NEG) for denoting the adverb $(not|R)$, $(n't|R)$, and other negation expressions, so that by default when we say the term *word*, we do not include (NEG).

Let $\mathcal{W} = \{w_1, w_2, \ldots, w_n\}$ be a set of a limited number of distinct words, a *clause*, denoted as s, is an ordered list of words $w_1 w_2 \ldots w_k$. The *length* of a clause is the number of words contained in the clause, denoted as $|s|$. For example, $(film|N)(be|V)(good|J)$ is a clause with length 3, in the order $(film|N)$ followed by $(be|V)$ and then followed by $(good|J)$. A word could also be a clause with length 1 if it is reduced to one lemma and its associated PoS tag. An *empty clause* is denoted as $\emptyset$, we have $s = \emptyset \iff |s| = 0$. The *concatenation* of clauses is denoted as the form $s_1 \cdot s_2$.

Within the context of mining sequence patterns [1], a word is an *item* and a clause is a *sequence*. Given two clauses $s = w_1 w_2 \ldots w_m$ and $s' = w'_1 w'_2 \ldots w'_n$, if there exist integers $1 \leq i_1 < i_2 < \ldots < i_m \leq n$ such that $w_i = w'_{j_i}$ for all w_i, then s is a *sub-clause* of s', denoted as $s \sqsubseteq s'$. If we have $s \sqsubseteq s'$, we say that *s is contained in s'*, or *s' supports s*. If clause s is not contained in any other clauses, then we say that the clause s is *maximal*. For example, the clause $(film|N)(good|J)$ is contained in the clause $(film|N)(be|V)(good|J)$, but is not contained in the clause $(be|V)(good|J)(film|N)$.

A *sentence*, denoted as S, is a maximal clause that is terminated by one of the following symbols ": ; . ? !" in the given text reviews. A *document*, denoted as $\mathcal{D}$, is an ordered list of sentences. Given a document $\mathcal{D}$, the *support* or *frequency* of a clause s, denoted as $\sigma(s, \mathcal{D})$, is the total number of sentences $S \in \mathcal{D}$ that support s. Given a user specified threshold of support called *minimum support*, denoted as *min_supp*, a clause is *frequent* if $\sigma(s, \mathcal{D}) \geq min\_supp$.

3.2 Contextual Models of Sentiment Polarity

We represent sentiment polarities as rule-format on clauses, that is, $s_\alpha \Rightarrow s_\beta$, where s_α and s_β are two clauses; given a clause s, if we have $s_\alpha \cdot s_\beta \sqsubseteq s$, then we say that the clause s *supports* the rule r, denoted as $s \models r$. We therefore propose a belief system for formalizing the opposite sentiments expressed in classified reviews. A *belief* on clauses, denoted as b, consists of a rule $s_\alpha \Rightarrow s_\beta$ and a semantical constraint $s_\beta \not\sim s_\gamma$, where the clause s_γ is semantically contradicts the clause s_β. We note a belief as $b = [s_\alpha; s_\beta; s_\gamma]$. A belief constrains that if the clause s_α occurs in a clause s, i.e., $s_\alpha \sqsubseteq s$, then the clause s_β should occur in s after s_β, and the clause s_γ should not occur in s after s_α, that is,

$$[s_\alpha; s_\beta; s_\gamma] \iff s_\alpha \sqsubseteq s \Longrightarrow s_\alpha \cdot s_\beta \sqsubseteq s \wedge s_\alpha \cdot s_\gamma \not\sqsubseteq s.$$

A clause s that verifies a belief b is *expected*, denoted as $s \models b$; that violates a belief b is *unexpected*, denoted as $s \not\models b$. Given a belief $b = [s_\alpha; s_\beta; s_\gamma]$ and a clause s such that $s_\alpha \sqsubseteq s$, the unexpectedness is considered as:

$$s_\alpha \cdot s_\beta \not\sqsubseteq s \wedge s_\alpha \cdot s_\gamma \sqsubseteq s \Longrightarrow s \not\models b.$$

Example 1. Given a belief $b = [(be|V); (good|J); (bad|J)]$ and two clauses $s_1 = (be|V)(a|DT)(good|J)(film|N)$, $s_2 = (be|V)(bad|J)(actor|N)$, we have $s_1 \models b$ and $s_2 \not\models b$. □

Let M^+ be the positive sentiment and M^- be the negative sentiment, a sentiment $M \in \{M^+, M^-\}$ can be expressed in documents (denoted as $\mathcal{D} \models M$), sentences (denoted as $S \models M$), clauses (denoted as $s \models M$) or vocabularies (denoted as $v \models M$). In addition, we denote the negation of a sentiment M as $\overline{M}$, so that we have $\overline{M^+} = M^-$ and $\overline{M^-} = M^+$. The negation is taken into account in other text-mining applications (for instance for synonym/antonym extraction process [13]).

Proposition 1. *Given a sentiment $M \in \{M^+, M^-\}$, if a document $\mathcal{D} \models M$, then there exists at least one sentence $S \in \mathcal{D}$ such that $S \models M$; if a sentence $S \models M$, then there exists at least one word $w \sqsubseteq S$ such that $w \models M$ or at least one clause $(NEG)v \sqsubseteq S$ (or $w(NEG) \sqsubseteq S$) such that $w \models \overline{M}$.*

Contextual Model	Sentiment Rule	Belief Pattern
J-N model	$(*\|J) \Rightarrow (*\|N)$	$[(\overline{*}\|J); \emptyset; (*\|N)]$ $[(NEG)(*\|J); \emptyset; (*\|N)]$
N-J model	$(*\|N) \Rightarrow (*\|J)$	$[(*\|N); (*\|J); (\overline{*}\|J)]$ $[(*\|N); (*\|J); (NEG)(*\|J)]$
V-J model	$(*\|V) \Rightarrow (*\|J)$	$[(*\|V); (*\|J); (\overline{*}\|J)]$ $[(*\|V); (*\|J); (NEG)(*\|J)]$ $[(*\|V)(NEG); (\overline{*}\|J); (*\|J)]$
J-V model	$(*\|J) \Rightarrow (*\|V)$	$[(*\|J); (*\|V); (*\|V)(NEG)]$
NEG-J-N model	$(NEG)(*\|J) \Rightarrow (*\|N)$	$[(NEG)(\overline{*}\|J); \emptyset; (*\|N)]$
N-NEG-J model	$(*\|N)(NEG) \Rightarrow (*\|J)$	$[(*\|N)(NEG); (*\|J); (\overline{*}\|J)]$
V-NEG-J model	$(*\|V)(NEG) \Rightarrow (*\|J)$	$[(*\|V)(NEG); (*\|J); (\overline{*}\|J)]$
J-V-NEG model	$(*\|J) \Rightarrow (*\|V)(NEG)$	$[(\overline{*}\|J); \emptyset; (*\|V)(NEG)]$

Fig. 1. Contextual models of sentiment polarity.

We focus on the sentiments expressed by the sentences that contain adjectives and nouns/verbs, such as "*this is a good film*". The sentiment expressed by sentences like "*this film is well produced*" is currently not considered in our approach. Note that we extract basic words relations without the use of syntactic analysis tools [12] to avoid the silence in the data (i.e. syntactic relations not extracted by the natural language systems).

With the adoption of rules and beliefs, we can extract the contextual information from reviews by finding the most frequent clauses that consist of at adjectives and nouns/verbs by sequential pattern mining algorithms, where the frequent nouns and verbs reflect topic of reviews, and the sentence-level sentiment polarities are expressed by frequent adjectives.

We propose a set of contextual models for constructing the belief base of opinion polarities within the context of review topic, listed in Fig. 1, where the word $(\overline{*}|J)$ stands for each antonym of the word $(*|J)$. Given a review, each sentence violating a belief generated from one of the belief patterns listed in Fig. 1 stands for an opposite sentiment.

3.3 Extracting Contextual Opposite Sentiments

We now introduce the training-extracting process of our approach. Let $\mathcal{V}$ be a set of adjectives expressing the sentiment M, we denote $\overline{\mathcal{V}}$ the set that contains the antonym(s) of each word contained in $\mathcal{V}$. Thus, for each $(*|J) \in \mathcal{V}$, we have $(*|J) \models M$ and $(\overline{*}|J) \in \overline{\mathcal{V}}$. Given a *training document* $\mathcal{D}_L$ such that for each sentence $S \in \mathcal{D}_L$, there exist at least one adjective $(*|J) \in \mathcal{V}$ or there exist (NEG) and at least one adjective $(*|J) \in \overline{\mathcal{V}}$. In order to construct the belief base of contextual models, we first apply a sequential pattern mining algorithm for discovering all maximal frequent clauses from $\mathcal{D}_L$ with respect to a minimum support threshold, denoted as $\mathcal{D}_F$. For each clause $s \in \mathcal{D}_F$, if s verifies a contextual model listed in Fig. 1 with the listing-order, then a set of beliefs can be generated from s corresponding to the belief pattern(s) of each contextual model. A belief base $\mathcal{B}_M$ can therefore be constructed with respect to the topic of reviews.

Example 2. Given a clause $s = (this|DT)(be|V)(a|DT)(good|J)(film|N)$, we have that s supports the J-N and V-J models, and the sentiment rules are $(good|J) \Rightarrow (film|N)$ and $(be|V) \Rightarrow (good|J)$. We have the priority of J-N model is higher than V-J model, so that $(good|J) \Rightarrow (film|N)$ is used for generating beliefs. Let $(bad|J)$ be the antonym of $(good|J)$, we have two beliefs generated: $[(bad|J); \emptyset; (film|N)]$ and $[(NEG)(good|J); \emptyset; (film|N)]$. □

Given a classified review $\mathcal{D}_M$ and a belief base $\mathcal{B}_M$ corresponding to the sentiment polarity M, the procedure of extracting unexpected sentences can be briefly described as follows. For each sentence $S \in \mathcal{D}_M$ and for each belief $b \in \mathcal{B}_M$ such that $b = [s_\alpha; s_\beta; s_\gamma]$, s_α is first matched for improving the performance; if $s_\alpha \sqsubseteq S$, and then if $s_\alpha \cdot s_\beta \not\sqsubseteq S$ and $s_\alpha \cdot s_\gamma \sqsubseteq S$, then S is an unexpected sentence expressing the contextual opposite sentiment $\overline{M}$. A detailed description of the representation of belief base and the unexpected sequence mining process can be found in [7].

4 Experiments

The data sets we use for evaluating our approach are the movie-review data[4] introduced in [8]. We combined these reviews into two documents $\mathcal{D}^+$ (containing 1,000 positive reviews, 75,740 sentences, and 21,156 distinct words) and $\mathcal{D}^-$ (containing 1,000 negative reviews, 67,425 sentences, and 19,714 distinct words). The two dictionaries $\mathcal{V}^+$ and $\mathcal{V}^-$ are generated from $\mathcal{D}^+$ and $\mathcal{D}^-$, by finding most frequent positive/negative adjectives.

#	Positive	Frequency	Negative	Frequency
1	good	2146	bad	1414
2	great	882	stupid	214
3	funny	441	poor	152
4	special	282	awful	109
5	perfect	244	silly	97
6	beautiful	202	horrible	71
7	nice	184	suck	65
8	entertaining	179	violent	64
9	wonderful	165	sad	56
10	excellent	146	ugly	44

Fig. 2. The dictionaries $\mathcal{V}^+$ and $\mathcal{V}^-$.

To not make our experiments too complex, we selected ten most frequent adjectives for each dictionary, listed as Fig. 2. The training documents $\mathcal{D}_L^+$ (contains 1,678 sentences) and $\mathcal{D}_L^-$ (contains 3,842 sentences) are therefore generated from $\mathcal{D}^+$ and $\mathcal{D}^-$ by gathering the sentences containing at least one adjective from $\mathcal{V}^+$ and $\mathcal{V}^-$.

The maximal frequent clauses (standing for $\mathcal{D}_F^+$ and $\mathcal{D}_F^-$) and the sentiment rules (standing for $\mathcal{P}^+$ and $\mathcal{P}^-$) extracted by the sequential pattern mining algorithm are shown in Fig. 3. For instance, with $min\_supp = 0.001$, we find 160 distinct sentiment rules from 572 discovered maximal frequent clauses in positive reviews, however with $min\_supp = 0.01$, only 8 distinct sentiment rules are found from 19 frequent clauses. The 10 most frequent sentiment rules are listed in Fig. 4. The antonym dictionaries for constructing the belief bases are given by WordNet. For respecting the size limit of this paper, we list a small part of the two belief bases in Fig. 5.

In order to analyze the accuracy of our approach, we randomly select a number of beliefs for extracting the sentences that express the sentiment opposite to the documents $\mathcal{D}^+$ and $\mathcal{D}^-$. For instance, as the beliefs listed in Fig. 5, the 5 beliefs of positive sentiment produced totally 304 unexpected sentences, and 236 of them express the negative sentiment; the 5 beliefs of negative sentiment produced totally 136 unexpected sentences, and 97 of them express the positive sentiment. Within these beliefs, the average accuracy is about 74.48%.

[4] http—//www.cs.cornell.edu/People/pabo/movie-review-data/

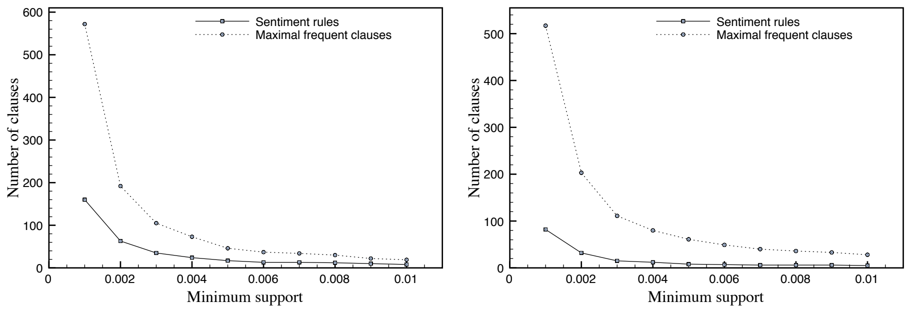

(a) (b)

Fig. 3. (a) Maximal frequent clauses and sentiment rules of positive reviews. (b) Maximal frequent clauses and sentiment rules of negative reviews..

Positive Sentiment Rules	**Negative Sentiment Rules**
⟨(be|V)⟩ ⇒ ⟨(good|J)⟩	⟨(bad|J)⟩ ⇒ ⟨(guy|N)⟩
⟨(good|J)⟩ ⇒ ⟨(film|N)⟩	⟨(bad|J)⟩ ⇒ ⟨(be|V)⟩
⟨(good|J)⟩ ⇒ ⟨(be|V)⟩	⟨(bad|J)⟩ ⇒ ⟨(movie|N)⟩
⟨(good|J)⟩ ⇒ ⟨(performance|N)⟩	⟨(bad|J)⟩ ⇒ ⟨(film|N)⟩
⟨(good|J)⟩ ⇒ ⟨(movie|N)⟩	⟨(bad|J)⟩ ⇒ ⟨(thing|N)⟩
⟨(good|J)⟩ ⇒ ⟨(friend|N)⟩	⟨(bad|J)⟩ ⇒ ⟨(year|N)⟩
⟨(great|J)⟩ ⇒ ⟨(film|N)⟩	⟨(bad|J)⟩ ⇒ ⟨(time|N)⟩
⟨(great|J)⟩ ⇒ ⟨(be|V)⟩	⟨(bad|J)⟩ ⇒ ⟨(dialogue|N)⟩
⟨(special|J)⟩ ⇒ ⟨(be|V)⟩	⟨(stupid|J)⟩ ⇒ ⟨(be|V)⟩
⟨(special|J)⟩ ⇒ ⟨(effect|N)⟩	⟨(poor|J)⟩ ⇒ ⟨(be|V)⟩

Fig. 4. The 10 most frequent sentiment rules.

Belief Base of Positive Sentiment	**Belief Base of Negative Sentiment**
[⟨(be|V)⟩ ; ⟨(good|J)⟩ ; ⟨(bad|J)⟩]	[⟨(not|R)(bad|J)⟩ ; ∅; ⟨(guy|N)⟩]
[⟨(be|V)⟩ ; ⟨(good|J)⟩ ; ⟨(not|R)(good|J)⟩]	[⟨(n't|R)(bad|J)⟩ ; ∅; ⟨(guy|N)⟩]
[⟨(be|V)⟩ ; ⟨(good|J)⟩ ; ⟨(n't|R)(good|J)⟩]	[⟨(bad|J)⟩ ; ⟨(be|V)⟩ ; ⟨(be|V)(not|R)⟩]
[⟨(bad|J)⟩ ; ∅; ⟨(film|N)⟩]	[⟨(bad|J)⟩ ; ⟨(be|V)⟩ ; ⟨(be|V)(n't|R)⟩]
[⟨(not|R)(good|J)⟩ ; ∅; ⟨(film|N)⟩]	[⟨(good|J)⟩ ; ∅; ⟨(film|N)⟩]
[⟨(n't|R)(good|J)⟩ ; ∅; ⟨(film|N)⟩]	[⟨(not|R)(bad|J)⟩ ; ∅; ⟨(film|N)⟩]
......	......

Fig. 5. The belief base for mining unexpected sentences.

5 Conclusion

In this paper we present a belief-driven approach that extracts contextual opposite sentiment as unexpected sentences from classified free text reviews. We adapt the sequence data model to text mining with Part-of-Speech tags, so that the extraction is associated with the semantic property of each word contained in the text reviews, thus the sequence mining techniques can be applied. Our experimental results show that the accuracy of the extracted opposite sentiments is in the acceptable range. Our future work includes to combine the adverbs and the conjunctions (like *however*, *but*) into the extraction process, and to integrate contextual opposite sentiments into document-sentence classifications.

References

1. R. Agrawal and R. Srikant. Mining sequential patterns. In *ICDE*, pages 3–14, 1995.
2. K. Dave, S. Lawrence, and D. M. Pennock. Mining the peanut gallery: Opinion extraction and semantic classification of product reviews. In *WWW*, pages 519–528, 2003.
3. A. Esuli and F. Sebastiani. PageRanking WordNet synsets: An application to opinion mining. In *ACL*, pages 424–431, 2007.
4. C. Fellbaum. *WordNet: An electronic lexical database*. MIT Press, 1998.
5. M. Hu and B. Liu. Mining and summarizing customer reviews. In *KDD*, pages 168–177, 2004.
6. J. Kamps, R. J. Mokken, M. Marx, and M. de Rijke. Using WordNet to measure semantic orientation of adjectives. In *LREC*, pages 1115–1118, 2004.
7. D. H. Li, A. Laurent, and P. Poncelet. Mining unexpected sequential patterns and rules. Technical Report RR-07027 (2007), LIRMM, 2007.
8. B. Pang and L. Lee. A sentimental education: Sentiment analysis using subjectivity summarization based on minimum cuts. In *ACL*, pages 271–278, 2004.
9. B. Pang, L. Lee, and S. Vaithyanathan. Thumbs up? Sentiment classification using machine learning techniques. In *EMNLP*, pages 79–86, 2002.
10. B. Santorini. Part-of-Speech tagging guidelines for the Penn Treebank project. Technical Report MS-CIS-90-47, Department of Computer and Information Science, University of Pennsylvania, 1990.
11. H. Schmid. Probabilistic Part-of-Speech tagging using decision trees. In *NeMLaP*, 1994.
12. D. D. Sleator and D. Temperley. Parsing English with a link grammar. In *3rd International Workshop on Parsing Technologies*, 1993.
13. P. D. Turney. Mining the Web for synonyms: PMI-IR versus LSA on TOEFL. In *ECML*, pages 491–502, 2001.
14. P. D. Turney. Thumbs up or thumbs down? Semantic orientation applied to unsupervised classification of reviews. In *ACL*, pages 417–424, 2002.
15. T. Wilson, J. Wiebe, and P. Hoffmann. Recognizing contextual polarity in phrase-level sentiment analysis. In *HLT/EMNLP*, 2005.
16. H. Yu and V. Hatzivassiloglou. Towards answering opinion questions: Separating facts from opinions and identifying the polarity of opinion sentences. In *EMNLP*, pages 129–136, 2003.