

# Opinion Sentence Search Engine on Open-domain Blog

Osamu Furuse, Nobuaki Hiroshima, Setsuo Yamada, Ryoji Kataoka

NTT Cyber Solutions Laboratories,

NTT Corporation

1-1 Hikarinooka Yokosuka-Shi Kanagawa, 239-0847 Japan

furuse.osamu@lab.ntt.co.jp

## Abstract

We have introduced a search engine that can extract opinion sentences relevant to an open-domain query from Japanese blog pages. The engine identifies opinions based not only on positive or negative measurements but also on neutral opinions, requests, advice, and thoughts. To retrieve a number of opinion sentences that a user could reasonably be expected to read, we attempted to extract only explicitly stated writer's opinions at the sentence-level and to exclude quoted or implicational opinions. In our search engine, opinion sentences are identified based on features such as opinion clue expressions, and then, the relevance to the query of each identified opinion sentence is checked. The experimental results for various topics, obtained by comparing the output of the proposed opinion search engine with that of human judgments as to whether the sentences were opinions, showed that the proposed engine has promise as a practical application.

## 1 Introduction

An enormous number of blog pages are freely written and frequently updated as private articles about various topics, including very timely ones. As numbers of blog writers and readers rapidly increase, blog pages as a consumer-generated medium (CGM) become increasingly important information sources about people's personal ideas, beliefs, feelings, and sentiments (positive or negative measurement). Such subjective information in blog pages can often be useful for finding out what people think about various topics in making a decision.

Studies on automatically extracting and analyzing reviews about a specific subject on the web [Dave *et al.*, 2003; Morinaga *et al.*, 2002; Turney, 2002; Nasukawa and Yi, 2003] have been conducted. An attempt has also been made to develop a system to analyze sentiments about open-domain queries in blogspace [Nanno *et al.*, 2004]. These efforts have focused on positive or negative measurement.

Sentiments and different kinds of subjective information such as neutral opinions, requests, and judgments provide useful information. For instance, opinion sentences like "In my opinion this product should be priced around \$15," which

do not explicitly express sentiments, can also be informative for a user who wants to know others' opinions about a product.

The sentence-level subjectivity classification approaches [Cardie *et al.*, 2003; Riloff and Wiebe, 2003; Wiebe and Riloff, 2005] try to identify subjective information that is broader than sentiments and suggest a way of searching for opinion sentences in open-domain topics. In these efforts, the subjectivity/objectivity of a current sentence is judged based on the existence of subjective/objective clues in both the sentence itself and the neighboring sentences. The subjective clues, such as adjectives, nouns, verb phrases, and other collocations, are learned from corpora [Wiebe, 2000; Wiebe *et al.*, 2001].

Opinion sentence searches using sentence-level subjectivity classification often collect too many sentences for a user to read. According to a previous study [Wiebe *et al.*, 2001], 70% of sentences in opinion-expressing articles like editorials and 44% of sentences in non-opinion expressing articles like news reports were judged to be subjective. Sentence-level subjectivity can be used to analyze subjectivity in inputted documents [Wilson *et al.*, 2003; Wilson *et al.*, 2005]. However, in searching opinion sentences from web pages, it is necessary to limit the number of retrieved sentences so that a user can survey them without undue effort.

We introduce opinion clue expressions, which are more restrictive than sentence-level subjectivity in conventional methods, as a criterion for judging opinion sentences. We also propose a method for searching opinion sentences from web pages using these clue expressions. Using the proposed method, we created a prototype opinion sentence search system in blogspace. The search engine extracts opinion sentences relevant to a user's query phrase about open-domain topics on products, persons, events, and social phenomena. The search engine identifies opinion sentences based on sentiments, neutral opinions, requests, advice, and thoughts. To retrieve a number of opinion sentences that is reasonable and that a user will want to read, we attempted to extract only explicitly stated writer's opinions at the sentence-level and to exclude quoted or implicational opinions. Section 2 describes what sentences should be searched as opinions. Section 3 gives an overview of our prototype opinion search system for Japanese blog pages. Sections 4 and 5 explain the two major modules, opinion search ex-

traction and query-relevant sentence extraction, and Section 6 evaluates the opinion sentence search method of our prototype system.

## 2 Opinion Sentences to be Searched

We judge a sentence to be an opinion if it explicitly declares the writer’s idea or belief at a sentence level. We define as an “opinion clue”, the part of a sentence that contributes to explicitly conveying the writer’s idea or belief in the opinion sentence [Hiroshima *et al.*, 2006]. For example, “I am glad” in the sentence “I am glad to see you” can convey the writer’s pleasure to a reader, so we regard the sentence as an “opinion sentence” and “I am glad” as an “opinion clue”. Another example of an opinion clue is the exclamation mark in the sentence “We got a contract!” It conveys the writer’s emotion about the event to a reader.

The existence of word-level or phrase-level subjective parts does not assure that the sentence is an opinion. Some word-level or phrase-level subjective parts can make the sentence an opinion depending on where they occur in the sentence. Consider the following two sentences.

- (1) This house is beautiful.
- (2) We purchased a beautiful house.

Both (1) and (2) contain the word-level subjective part “beautiful”. Our criterion would lead us to say that sentence (1) is an opinion, because “beautiful” is placed in the predicate part and (1) is considered to declare the writer’s evaluation of the house to a reader. This is why “beautiful” in (1) contributes to make the sentence an opinion. By contrast, sentence (2) is not judged to be an opinion, because “beautiful” is placed in the object of the verb “purchase” and (2) is considered to report the event of the house purchase rather objectively to a reader. Sentence (2) contains subjective information about the beauty of the house; however this information is unlikely to be what a writer wants to emphasize. Thus, “beautiful” in (2) does not contribute to making the sentence an opinion.

These two sentences illustrate the fact that the presence of a subjective word (“beautiful”) does not unconditionally assure that the sentence is an opinion. Additionally, these examples do suggest that whether a sentence is an opinion can be judged depending on where such word-level or phrase-level subjective parts as evaluative adjectives are placed in the predicate part.

Some word-level or phrase-level subjective parts such as subjective sentential adverbs contribute to making the sentence an opinion depending on where they occur in the sentence. Sentence (3) is judged to be an opinion because its main clause contains a subjective sentential adverb “amazingly”, which expresses the writer’s feeling about the event.

- (3) Amazingly, few people came to my party.

The presence of idiomatic collocations in the main clause also affects our judgment as to what constitutes an opinion sentence. For example, sentence (4) can be judged as an opinion because it includes “my wish is”.

- (4) My wish is to go abroad.

Thus, depending on the type of opinion clue, it is necessary to consider where the expression occurs in the sentence to judge whether the sentence is an opinion.

## 3 Architecture of Opinion Sentence Search

Figure 1 shows the configuration of our prototype opinion sentence search system in blogspace.

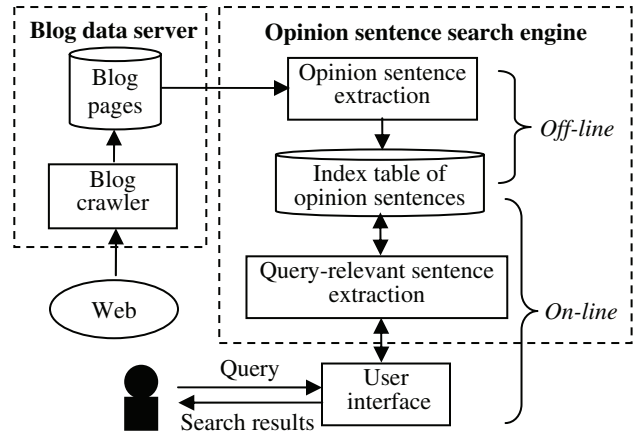


Figure 1: Configuration of opinion sentence search

The blog data server collects blog pages by periodically crawling the web. Our opinion sentence search engine, which receives blog pages from the blog data server, consists of two main modules: opinion sentence extraction and query-relevant sentence extraction.

The opinion sentence extraction module checks whether each sentence in the crawled blog pages can be considered an opinion. Opinion sentences are extracted and indexed as off-line processing, which, for a practical real-time search, should be as high a proportion of the entire processing as possible. The index table in the blog data server can accommodate more than 1,262,000 updated blog pages every month.

The query-relevant sentence extraction module retrieves opinion sentences relevant to the user’s query phrases from the index table of opinion sentences in the blog page server. Since users’ queries cannot be predicted, query-relevant sentence extraction has to include on-line processing.

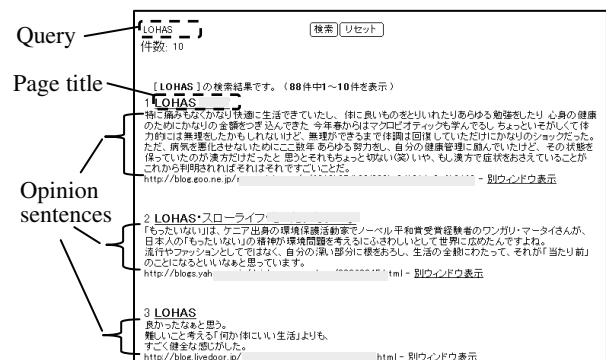


Figure 2: User interface by open-domain query

Figure 2 shows the user interfaces we provide now. A user inputs open-domain keyword phrases in the query box and then clicks the search button. The opinion sentences resulting from the search are presented in a blog page unit. The result pages can be ranked according to the number of opinion sentences, the ratio of opinion sentences to total sentences, or total strength of the opinion sentences.

## 4 Opinion Sentence Extraction

It is difficult to enumerate the opinion-judgment rules describing diversified features under some conditions in a rule-based method. To avoid the poor performance caused by data sparseness and the daunting task of writing rules, we adopted a learning method that binarily classifies sentences using opinion clues and their positions in sentences [Hiroshima *et al.*, 2006] as feature parameters of a Support Vector Machine (SVM). An SVM can efficiently learn the model for classifying sentences as opinion and non-opinion, based on the combinations of multiple feature parameters. Following are the feature parameters of our method.

- 2,936 opinion clue expressions
- 2,715 semantic categories
- 150 frequent words
- 13 parts of speech

Opinion clue expressions and semantic categories are crucial feature parameters. The semantic categories we adopted have a hierarchical structure and are from a Japanese thesaurus [Ikehara *et al.*, 1997].

### 4.1 Clue Expression Collection

Whether expressions have opinion clues is a basic criterion for judging whether a sentence expresses an opinion. To collect opinion clue expressions for an open-domain opinion sentence search, we extracted opinion sentences from the top twenty Japanese web pages retrieved with forty queries on various kinds of topics. The queries correspond to possible situations in which a user wants to retrieve opinions from web pages about a particular topic. The retrieved pages were unrestricted to blog pages because we target the opinion sentence search engine applicable not only to blog pages but also to other CGM pages or all web pages, and we hypothesize that opinion clues do not differ between blog pages and other web pages.

Out of 75,575 sentences in the total 800 retrieved pages, the 13,363 sentences judged unanimously by three evaluators to be opinions were extracted. Then, of these 13,363 sentences considered very likely to be opinions, 8,425 were used to extract opinion clues by the human analysts, while the remaining 4,938 were reserved for future assessment for other CGM pages or general web pages.

The total number of opinion clues obtained was 2,936. These clue expressions were classified into two groups, as shown in the example sentences below. The underlined expressions in example sentences are extracted as opinion clues. There were 2,514 clues and 422 clues in each group. The example sentences are translations of Japanese opinion sentences extracted by human analysts.

- **2,514 clues appearing in the predicate part**

**Thought:** I think this book is his.

**Intensifier:** They played extremely well.

**Impression:** This terminology is confusing.

**Emotion:** I am glad to see you.

**Positive/negative judgment:**

Your audio system is terrific.

**Modality about propositional attitude:**

You should go to the movie.

**Value judgment:** This sentence makes no sense.

**Utterance-specific sentence form:**

However, it's literally just a dream now.

**Symbol:** We got a contract!

**Uncertainty:** I am wondering what I should eat for lunch.

**Imperative:** Don't do that.

- **422 clues not appearing in the predicate part**

**Declarative adverb:**

I will possibly go to Europe next year.

**Interjection:** Oh, wonderful.

**Idiomatic collocation:** It's hard to say.

The opinion clues in the Japanese examples are placed in the last part of sentences in the first group. This reflects the heuristic rule that Japanese predicates are in principle placed in the last part of a sentence.

### 4.2 Augmentation by Semantic Categories

Opinion clue expressions can be augmented by the semantic categories of the words in the expressions. The feature parameters for a semantic category have two roles: one is to compensate for the insufficient amount of opinion clue expressions, and the other is to consider the relations between clue expressions and co-occurring words in the opinion sentences. Consider the following two sentence patterns:

(5) X is beautiful.

(6) X is pretty.

The words “beautiful” and “pretty” are adjectives in the common semantic category, “appearance”, and the degree of sentence-level subjectivity of these sentences is almost the same regardless of what X is. Therefore, even if “beautiful” is learned as an opinion clue but “pretty” is not, the semantic category “appearance” to which the learned word “beautiful” belongs, enables (6) to be judged an opinion as well as (5).

Many of the opinion clue expressions have co-occurring words in the opinion sentence. Consider the following two sentences.

(7) The sky is high.

(8) The quality of this product is high.

Both (7) and (8) contain the word “high” in the predicate part. Sentence (7) is considered to be less of an opinion than (8) because (7) might be judged to be the objective truth, while (8) is likely to be judged an opinion. The adjective “high” in the predicate part can be validated as an opinion clue depending on co-occurring words. However, providing all possible co-occurring words with each opinion clue expression is not a realistic option. The co-occurrence infor-

mation about each opinion clue expression can be generalized using semantic categories.

### 4.3 Training and Test Set

The training set was chosen from blog pages different from the web pages used for opinion clue collection. This was done in order to conduct training specific to blog searches and in order to conduct training and testing independently of opinion clue collection.

We used the same procedure as we did to collect opinion clues, to prepare training and test sets that are both specific to blog search. We first retrieved Japanese blog pages with ninety queries covering a wide range of topics:

- Culture:** movies, books, music
- Entertainment:** sports, TV drama, games
- Facilities:** museums, zoos, amusement parks
- Food:** beer, noodles, ice cream
- Health:** medicine, syndromes
- Local information:** restaurants, hotels, hot springs
- Person:** comedians, idols
- Phenomena:** lifestyle, environment
- Politics, Economy:** elections, gasoline prices
- Products:** cell phones, cars, beer, cosmetics, software

Opinion sentences were extracted from the top ten retrieved blog pages for each query, leaving 900 pages and 29,486 sentences in total. Three evaluators judged whether each sentence was an opinion or not. Out of 29,486 sentences, 2,868 were judged to be opinions by all three evaluators, 3,725 by two evaluators, 3,248 by one evaluator, and 19,645 were judged to be non-opinions by all three evaluators.

Number	Training set	Test set
Query	72	18
Total sentences	23,800	5,686
Sentences all three judged opinions	2,416	452
Sentences the two judged opinions	3,003	722
Sentences the one judged opinions	2,631	617
Sentences none judged opinions	15,750	3,895

Table 1: Training and test set

Eighteen queries, one-fifth of the total, were randomly selected, and the sentences for the queries were used for testing. The sentences of the other seventy-two queries were used for training. The breakdown of training and test sets is shown in Table 1.

We set the sentences with at least one judged opinion as a cut-off point. Thus, 8,050 were then used to learn positive examples in the SVM, and 1,791 were used to assess the performance of the opinion sentence search system (Section 6). 15,750, non-opinion sentences were used to learn negative examples, and 3,895 were used for assessment.

## 5 Query-relevant Sentence Extraction

The three evaluators also judged whether each opinion sentence in a training and test set in Section 4.3 was query-relevant. Of the 9,841 sentences that at least one

evaluator judged to be an opinion, 2,544 were judged to be relevant to the queries by at least one evaluator, and 7,297 were judged by all three evaluators to be unrelated to the queries. The high percent of the latter figure, 7,297, which is 74.1% of the 9,841 opinion sentences, shows that it is inappropriate to accept as search results all opinion sentences in the pages retrieved by a user's query.

### 5.1 Permissible Scope of Query Relevance

Not all of the retrieved opinion sentences are closely related to the query because some of the pages describe miscellaneous topics. The permissible scope between individual users for query relevance of a sentence differs. The following are opinion sentences from the retrieved pages queried with "product name of a game console". The number of evaluators who judged the sentence to be query-relevant is shown in parentheses.

- (9) I was impressed with the compactness. (all three)
- (10) An adult also can enjoy this. (two of the three)
- (11) The manufacturer is marvelous. (one of the three)
- (12) Technological advancement is very rapid. (none)

The above sentences show that individual judgments differ when a sentence tends to be indirectly or weakly relevant to the query. We take a stand on accepting weak query relevance because it is more advantageous in a real-time search to pursue possible query relevance heuristically or eliminatively rather than to verify query relevance precisely. Thus, we considered the sentences judged by at least one evaluator query-relevant to be a correct answer.

### 5.2 Strategies about Query Relevance

Query-relevant sentence extraction in the prototype system has the following two heuristic and simple strategy options.

- (a) A sentence is relevant to the query only when a query phrase occurs in the sentence or within some number of sentences before the sentence.
- (b) A sentence is relevant to the query only when a query phrase occurs in the sentence or within the chunk that the sentence belongs to and only opinion sentences consecutively appear in.

The strategy adopted affects the number of opinion sentences an index table can accommodate. From this viewpoint, Strategy (b) is better because an index table needs information only about opinion sentences. In contrast, with Strategy (a), an index table has to accommodate non-opinion sentences immediately before opinion sentences in addition to the opinion sentences themselves.

## 6 Experiments

We conducted experiments in the prototype opinion sentence search system in blogspace to assess opinion sentence extraction, query-relevant sentence extraction, and a combination of the two. All experiments used the Japanese sentences described in Sections 4.3. The numbers of sentences used for training and testing are shown in Table 1.



## 6.1 Evaluation of Opinion Sentence Extraction

The experiments on opinion sentence extraction were designed to ascertain the effect of feature parameters on opinion sentence learning and the effect of position information on opinion clues. Answers where at least one of the three evaluators judged the sentence to be an opinion were defined as correct, and answers where no evaluator judged the sentence to be an opinion were defined as wrong.

Method	Precision	Recall	Accuracy
Baseline	67.5%	40.3%	75.1%
Proposed (without semantic categories)	75.0%	47.6%	78.5%
Proposed (with semantic categories)	72.5%	54.8%	79.2%

Table 2: Comparison with baseline

The main feature parameters for the SVM learner are clue expressions and semantic categories, as explained in Section 4. We prepared a baseline method that regards a sentence as an opinion if it contains a number of opinion clues that does not dip below a certain threshold. The best threshold was set through trial and error at four occurrences. The experimental results in Table 2 show that our method performs better than the baseline method. Precision is defined as the correctness ratio of the sentences extracted as opinions. Recall is defined as the ratio of opinion sentences correctly extracted over the total number of test opinion sentences. Accuracy is defined as the correct judgment ratio of all the test (both opinion and non-opinion) sentences. The two bottom rows show the results of our opinion sentence extraction method. The second bottom row concerns methods that do not use semantic categories, and the bottom row concerns those that do. The results in these two cases show that clue expressions are effective and that semantic categories improve performance.

We also evaluated the effect of position information of 2,936 opinion clues based on the heuristic rule that a Japanese predicate part almost always appears in the last ten words in a sentence. Instead of more precisely identifying predicate position from parsing information, we employed this heuristic rule as a feature parameter in the SVM learner and classifier for practical reasons.

Position	Precision	Recall	Accuracy
All words	71.2%	48.6%	77.6%
Quasi predicate part	72.5%	54.8%	79.2%

Table 3: Effect of opinion-clue position restriction

Table 3 lists the experimental results for position restriction of opinion clues. “All words” indicates that all feature parameters are permitted at any position in the sentence. “Quasi predicate part” indicates that all feature parameters are permitted only if they occur within the last ten words in the sentence. Although we narrowed the scope to consider the feature parameters and adopted an expedient method to locate the predicate part, feature parameters within the last ten words perform better in all evaluations than those without position restriction. The fact that the equal position restriction on all opinion clues improved performance suggests that

assigning the individual position condition to each opinion clue or locating the predicate part more precisely significantly improves performance.

The ratios of sentences the system judged opinion were, 74.3% to the opinion sentences three evaluators judged to be opinions, 62.0% to those two judged to be opinions, 44.4% to those one judged to be opinions, and 11.4% to those three judged to be non-opinions. Even though all sentences judged by at least one evaluator to be opinions were equally trained as correct answers, the higher the number of evaluators judging a sentence to be an opinion, the more likely our method was to judge it an opinion. This result shows that our method is congruent with human judgment.

## 6.2 Evaluation of Query Relevance

We investigated the performance of the query-relevant sentence extraction strategies described in Section 5.2, using all 1,791 opinion sentences in a test set in Table 1. The performance values were computed based on the correct answers being the 429 sentences that at least one of the three evaluators had judged to be query-relevant and the wrong answers being 1,362 sentences that all three evaluators had judged to be query-irrelevant. We modified Strategy (a) in Section 5.2, as follows.

- (a) A sentence is relevant to the query only when a query phrase exists in the sentence or in those right before the sentence.

Strategy (b) was not modified for the evaluation. We prepared a baseline method that regards a sentence as query relevant if it contains a query phrase.

Method	Precision	Recall	Accuracy
Baseline	74.0%	16.6%	78.6%
Strategy (a)'	65.0%	33.3%	79.7%
Strategy (b)	53.2%	41.3%	77.2%

Table 4: Evaluation of query relevance strategy

Table 4 shows the experimental results of query-relevance extraction from 2,868 opinion sentences in the baseline, Strategy (a)', and Strategy (b). These results show that our strategies performed with much better recall and slightly worse precision than the baseline method. Although the above results show that our strategies need improvement, Strategy (a)' and Strategy (b) seem to amount to a practical solution at present. Strategy (b), which our system is currently using is advantageous from the viewpoint of the amount of opinion sentences in an index table but is somewhat inferior to Strategy (a)' in precision.

## 6.3 Evaluation of Total Performance

The total performance of the opinion sentence search is obtained by multiplying performance of the two modules, opinion sentence extraction, and query-relevant sentence extraction. The performance values were computed based on the correct answers being the 429 sentences that were judged by at least one of the three evaluators to be query-relevant opinions out of all 5,686 test sentences in Table 1. The ratio of opinion query-relevant sentences in test sentences, 7.5%,

which is 429 out of 5,686, suggests that the number of sentences for which the system pursues retrieval is a reasonable amount for a user to read.

Table 5 shows the experimental results for total performance obtained by combining the two modules. For opinion sentence extraction, all feature parameters described in Section 4 and the opinion-clue position restriction described in Section 6.1 were used. In query-relevant sentence extraction, two trials, one using Strategy (a)', and the other using Strategy (b) were attempted.

Query-relevance	Precision	Recall	Accuracy
Strategy (a)'	55.2%	18.6%	92.7%
Strategy (b)	52.2%	14.0%	92.5%

Table 5: Evaluation of total performance

The results show that although total performance must be improved, the precision values, which were not low, suggest that system's output is reasonably reliable. The precision of total performance was higher than the multiplication products of the two modules. This is thought to be because opinion sentences tend to be more query-relevant than non-opinion sentences.

## 7 Conclusion and Future Work

We proposed an opinion sentence search method for Japanese open-domain blog pages. The experiments suggested that the performance of the prototype system has promise as a practical application. While the performance of opinion sentence extraction was good, it is necessary to improve the query-relevant sentence extraction strategy while storing as many opinion sentences as possible in the space available in the index table in the blog data server. Another avenue of future work is to develop a richer user interface where extracted opinion sentences can be classified in terms of emotion, sentiment, requirement, and suggestion, so that a user can retrieve relevant opinions on demand, and where extracted sentences are summarized so that the user can quickly learn what the writer wanted to say.

## References

- [Cardie *et al.*, 2003] Claire Cardie, Janyce Wiebe, Theresa Wilson, and Diane J. Litman. Combining Low-Level and Summary Representations of Opinions for Multi-Perspective Question Answering. *Working Notes - New Directions in Question Answering (AAAI Spring Symposium Series)*, 2003.
- [Dave *et al.*, 2003] Kushal Dave, Steve Lawrence, and David M. Pennock. Mining the peanut gallery: Opinion extraction and semantic classification of product reviews. *Proceedings of the 12th International World Wide Web Conference*, 519-528, 2003.
- [Hiroshima *et al.*, 2006] Nobuaki Hiroshima, Setsuo Yamada, Osamu Furuse, and Ryoji Kataoka. Searching for Sentences Expressing Opinions by using Declaratively Subjective Clues. *Proceedings of COLING-ACL 2006 Workshop "Sentiment and Subjectivity in Text"*, 39-46, 2006.
- [Ikehara *et al.*, 1997] Satoru Ikehara, Masahiro Miyazaki, Akio Yokoo, Satoshi Shirai, Hiromi Nakaiwa, Kentaro Ogura, Yoshifumi Ooyama, and Yoshihiko Hayashi. *Nihongo Goi Taikai – A Japanese Lexicon. Iwanami Shoten. 5 volumes.* (In Japanese), 1997.
- [Morinaga *et al.*, 2002] Satoshi Morinaga, Kenji Yamanishi, and Kenji Tateishi. Mining Product Reputations on the Web. *Proceedings of the eighth ACM SIGKDD International Conference on Knowledge Discovery and Data Mining (KDD 2002)*, 2002.
- [Nanno *et al.*, 2004] Tomoyuki Nanno, Toshiaki Fujiki, Yasuhiro Suzuki, and Manabu Okumura, Automatically Collecting, Monitoring, and Mining Japanese Weblogs, *13th International World Wide Web Conference*, 2004.
- [Nasukawa and Yi, 2003] Tetsuya Nasukawa and Jeonghee Yi. Sentiment Analysis: Capturing Favorability Using Natural Language Processing. *Proceedings of the 2nd International Conference on Knowledge Capture (K-CAP 2003)*, 2003.
- [Riloff and Wiebe, 2003] Ellen Riloff and Janyce Wiebe. Learning Extraction Patterns for Subjective Expressions. *Proceedings of the Conference on Empirical Methods in Natural Language Processing (EMNLP-03)*, 105-112, 2003.
- [Turney, 2002] Peter Turney. Thumbs Up or Thumbs Down? Semantic Orientation Applied to Unsupervised Classification of Reviews. *Proceedings of the 40th Annual Meeting of the Association for Computational Linguistics (ACL-2002)*, 417-424, 2002.
- [Wiebe, 2000] Janyce Wiebe. Learning Subjective Adjectives from Corpora. *Proceedings of the 17th National Conference on Artificial Intelligence (AAAI-2000)*. 2000.
- [Wiebe *et al.*, 2001] Janyce Wiebe, Theresa Wilson, and Matthew Bell. Identifying Collocations for Recognizing Opinions. *Proceedings of ACL/EACL 2001 Workshop on Collocation*, 2001.
- [Wiebe and Riloff, 2005] Janyce Wiebe and Ellen Riloff. Creating Subjective and Objective Sentence Classifiers from Unannotated Texts. *Proceedings of the Sixth International Conference on Intelligent Text Processing and Computational Linguistics (CICLing-2005)*, 486-497, 2005.
- [Wilson *et al.*, 2003] Theresa Wilson, David R. Pierce, and Janyce Wiebe, Identifying Opinionated Sentences. *Proceedings of the 2003 Human Language Technology Conference of the North American Chapter of the Association for Computational Linguistics on Human Language Technology (HLT-NAACL2003): Demonstrations*, 33-34. 2003.
- [Wilson *et al.*, 2005] Theresa Wilson, Paul Hoffmann, Swapna Somasundaran, Jason Kessler, Janyce Wiebe, Yejin Choi, Claire Cardie, Ellen Riloff, and Siddharth Patwardhan, OpinionFinder: A System for Subjectivity Analysis. *Proceedings of Human Language Technologies Conference/Conference on Empirical Methods in Natural Language Processing (HLT-EMNLP2005)*, 34-35. 2005.