

# Constructing Thai Opinion Mining Resource: A Case Study on Hotel Reviews

Choochart Haruechaiyasak, Alisa Kongthon,  
Pornpimon Palingoon and Chatchawal Sangkeettrakarn

Human Language Technology Laboratory (HLT)  
National Electronics and Computer Technology Center (NECTEC)

choochart.har@nectec.or.th, alisa.kon@nectec.or.th,  
pornpimon.pal@nectec.or.th, chatchawal.san@nectec.or.th

## Abstract

Opinion mining and sentiment analysis has recently gained increasing attention among the NLP community. Opinion mining is considered a domain-dependent task. Constructing lexicons for different domains is labor intensive. In this paper, we propose a framework for constructing Thai language resource for *feature-based* opinion mining. The feature-based opinion mining essentially relies on the use of two main lexicons, *features* and *polar words*. Our approach for extracting features and polar words from opinionated texts is based on syntactic pattern analysis. The evaluation is performed with a case study on hotel reviews. The proposed method has shown to be very effective in most cases. However, in some cases, the extraction is not quite straightforward. The reasons are due to, firstly, the use of conversational language in written opinionated texts and, secondly, the language semantic. We provide discussion with possible solutions on pattern extraction for some of the challenging cases.

## 1 Introduction

With the popularity of Web 2.0 or social networking websites, the amount of user-generated contents has increased exponentially. One interesting type of these user-generated contents is texts which are written with some opinions and/or sentiments. An in-depth analysis of these

opinionated texts could reveal potentially useful information regarding the preferences of people towards many different topics including news events, social issues and commercial products. Opinion mining and sentiment analysis is such task for analyzing and summarizing what people think about a certain topic.

Due to its potential and useful applications, opinion mining has gained a lot of interest in text mining and NLP communities (Ding et al., 2008; Jin et al., 2009). Much work in this area focused on evaluating reviews as being positive or negative either at the document level (Turney, 2002; Pang et al., 2002; Dave et al., 2003; Beineke et al., 2004) or sentence level (Kim and Hovy, 2004; Wiebe and Riloff, 2005; Wilson et al., 2009; Yu and Hatzivassiloglou, 2003). For instance, given some reviews of a product, the system classifies them into positive or negative reviews. No specific details or features are identified about what customers like or dislike. To obtain such details, a *feature-based* opinion mining approach has been proposed (Hu and Liu, 2004; Popescu and Etzioni, 2005). This approach typically consists of two following steps.

1. Identifying and extracting features of an object, topic or event from each sentence upon which the reviewers expressed their opinion.
2. Determining whether the opinions regarding the features are positive or negative.

The feature-based opinion mining could provide users with some insightful information related to opinions on a particular topic. For example, for hotel reviews, the feature-based opinion

mining allows users to view positive or negative opinions on hotel-related features such as price, service, breakfast, room, facilities and activities. Breaking down opinions into feature level is very essential for decision making. Different customers could have different preferences when selecting hotels to stay for vacation. For example, some might prefer hotels which provide full facilities, however, some might prefer to have good room service.

The main drawback of the feature-based opinion mining is the preparation of different lexicons including *features* and *polar words*. To make things worse, these lexicons, especially the features, are domain-dependent. For a particular domain, a set of features and polar words must be prepared. The process for language resource construction is generally labor intensive and time consuming. Some previous works have proposed different approaches for automatically constructing the lexicons for the feature-based opinion mining (Qiu et al., 2009; Riloff and Wiebe, 2003; Sarmiento et al., 2009). Most approaches applied some machine learning algorithms for learning the rules from the corpus. The rules are used for extracting new features and polar words from untagged corpus. Reviews of different approaches are given in the related work section.

In this paper, we propose a framework for constructing Thai language resource for the feature-based opinion mining. Our approach is based on syntactic pattern analysis of two lexicon types: *domain-dependent* and *domain-independent*. The domain-dependent lexicons include features, sub-features and polar words. The domain-independent lexicons are particles, negative words, degree words, auxiliary verbs, prepositions and stop words. Using these lexicons, we could construct a set of syntactic rules based on the frequently occurred patterns. The rule set can be used for extracting more unseen sub-features and polar words from untagged corpus.

We evaluated the proposed framework on the domain of hotel reviews. The experimental results showed that our proposed method is very

effective in most cases, especially for extracting polar words. However, in some cases, the extraction is not quite straightforward due to the use of conversational language, idioms and hidden semantic. We provide some discussion on the challenging cases and suggest some solutions as the future work.

The remainder of this paper is organized as follows. In next section, we review some related works on different approaches for constructing language resources for opinion mining and sentiment analysis. In Section 3, we present the proposed framework for constructing Thai opinion mining resource by using the dual pattern extraction method. In Section 4, we apply the proposed framework with a case study of hotel reviews. The performance evaluation is given with the experiment results. Some difficult cases are discussed along with some possible solutions. Section 5 concludes the paper with the future work.

## 2 Related work

The problem of developing subjectivity lexicons for training and testing sentiment classifiers has recently attracted some attention. The Multi-perspective Question Answering (MPQA) opinion corpus is a well-known resource for sentiment analysis in English (Wiebe et al., 2005). It is a collection of news articles from a variety of news sources manually annotated at word and phrase levels for opinions and other private states (i.e., beliefs, emotions, sentiments, speculations, etc.). The annotation in this work also took into account the context, which is essential for resolving possible ambiguities and accurately determining polarity.

Although most of the reference corpora has been focused on English language, work on other languages is growing as well. Kanayama et al. (2006) proposed an unsupervised method to detect sentiment words in Japanese. In this work, they used clause level context coherency to identify candidate sentiment words from sentences that appear successively with sentences containing seed sentiment words. Their assumption is that unless the context is changed with adversative expressions, sentences appearing together

in that context tend to have the same polarities. Hence, if one of them contains sentiment words, the other successive sentences are likely to contain sentiment words as well. Ku and Chen (2007) proposed the bag-of-characters approach to determine sentiment words in Chinese. This approach calculates the observation probabilities of characters from a set of seed sentiment words first, then dynamically expands the set and adjusts their probabilities. Later in 2009, Ku et al. (2009), extended their bag-of-characters approach by including morphological structures and syntactic structures between sentence segment. Their experiments showed better performance of word polarity detection and opinion sentence extraction.

Some other methods to automatically generate resources for subjectivity analysis for a foreign language have leveraged the resources and tools available for English. For example, Benea et al. (2008) applied machine translation and standard Naive Bayes and SVM for subjectivity classification for Romanian. Their experiments showed promising results for applying automatic translation to construct resources and tools for opinion mining in a foreign language. Wan (2009) also leveraged an available English corpus for Chinese sentiment classification by using the co-training approach to make full use of both English and Chinese features in a unified framework. Jijkoun and Hofmann (2009) also described a method for creating a Dutch subjectivity lexicon based on an English lexicon. They applied a PageRank-like algorithm that bootstraps a subjectivity lexicon from the translation of the English lexicon and rank the words in the thesaurus by polarity using the network of lexical relations (e.g., synonymy, hyponymy) in Wordnet.

### 3 The proposed framework

The performance of the feature-based opinion mining relies on the design and completeness of related lexicons. Our lexicon design distinguishes lexicons into two types, domain-dependent and domain-independent. The design of domain-dependent lexicons is based on the

feature-based opinion mining framework proposed by Liu et al. (2005). The framework starts by setting the domain scope such as *digital camera*. The next step is to design a set of features associated with the given domain. For the domain of digital camera, features could be, for instance, “price”, “screen size” and “picture quality”. Features could contain sub-features. For example, the picture quality could have the sub-features as “macro mode”, “portrait mode” and “night mode”. Preparing multiple feature levels could be time-consuming, therefore, we limit the features into two levels: *main features* and *sub-features*.

Another domain-dependent lexicon is *polar words*. Polar words are sentiment words which represent either positive or negative views on features. Although some polar words are domain-independent and have explicit meanings such as “excellent”, “beautiful”, “expensive” and “terrible”. Some polar words are domain-dependent and have implicit meanings depending on the contexts. For example, the word “large” is generally considered positive for the *screen size* feature of digital camera domain. However, for the *dimension* feature of mobile phone domain, the word “large” could be considered as negative.

On the other hand, the domain-independent lexicons are regular words which provide different parts of speech (POS) and functions in the sentence. For opinion mining task, we design six different domain-independent lexicons as follows (some examples are shown in Table 1).

- **Particles (PAR):** In Thai language, these words refer to the sentence endings which are normally used to add politeness of the speakers (Cooke, 1992).
- **Negative words (NEG):** Like English, these words are used to invert the opinion polarity. Examples are “not”, “unlikely” and “never”.
- **Degree words (DEG):** These words are used as an intensifier to the polar words. Examples are “large”, “very”, “enormous”.

- **Auxiliary verbs (AUX):** These words are used to modify verbs. Examples are “should”, “must” and “then”.
- **Prepositions (PRE):** Like English, Thai prepositions are used to mark the relations between two words.
- **Stop words (STO):** These words are used for grammaticalization. Thai language is considered an isolating language, to form a noun the words “karn” and “kwam” are normally placed in front of a verb or a noun, respectively. Therefore, these words could be neglected when analyzing opinionated texts.

Lexicons	Examples
<b>Particles (PAR)</b>	เลย, หน่อย, นึก, ด้วย, เท่าไหร่, ลี, ซิ, นะ, ครับ, ค่ะ, นะคะ, อ่ะ, จ๊ะ, ครับ
<b>Negative words (NEG)</b>	ไม่ (not), ไม่น่า (unlikely), ไม่เคย (never)
<b>Degree words (DEG)</b>	มาก (large), มากมาก (very), มากๆ (very) มากมาย (enormous), ค่อนข้าง (most likely) พอควร (borderline), เกินไป (exceed), สุดยอด (awesome)
<b>Auxiliary verbs (AUX)</b>	ก็ (then), ควร (should), ควรจะ (should), ค่อนข้าง (likely), ต้อง (must), น่าจะ (should), ยัง (yet)
<b>Preposition (PRE)</b>	และ (and), กับ (with), ของ (of), ใน (in) รวมไปถึง (including), บน (on), ที่ (at)
<b>Stopword (STO)</b>	การ, ความ

Table 1: Domain-independent lexicons

Although some of the above lexicons are similar to English, however, some words are placed in different position in a sentence. For example, in Thai, a degree word is usually placed after a polar word. For example, “very good” would be written as “good very” in Thai.

Figure 1 shows all processes and work flow under the proposed framework. The process starts with a corpus which is tagged based on two lexicon types. From the tagged corpus, we construct patterns and lexicons. The pattern construction is performed by collecting text segments which contain both features and polar

words. All patterns are sorted by the frequency of occurrence. The lexicon construction is performed by simply collecting words which are already tagged with the lexicon types. The lexicons are used for performing the feature-based opinion mining task such as classifying and summarizing the reviews as positive and negative based on different features. The completeness of lexicons is very important for the feature-based opinion mining. To collect more lexicons, patterns are used in the dual pattern extraction process to extract more features and polar words from the untagged corpus.

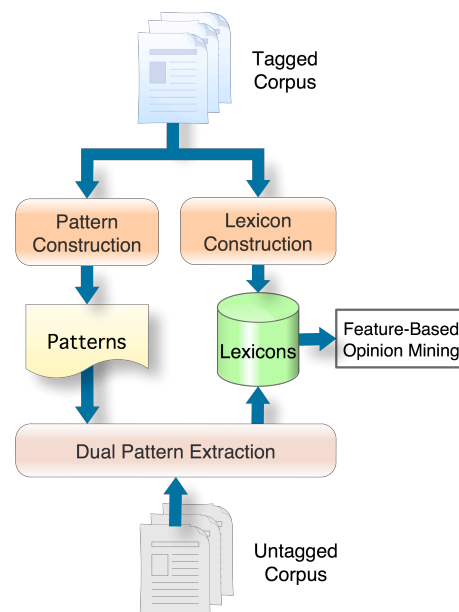


Figure 1: The proposed opinion resource construction framework based on the dual pattern extraction.

#### 4 A case study of hotel reviews

To evaluate the proposed framework, we perform some experiments with a case study of hotel reviews. In Thailand, tourism is ranked as one of the top industries. From the statistics provided by the Office of Tourism Development<sup>1</sup>, the number of international tourists visiting Thailand in 2009 is approximately 14 mil-

<sup>1</sup>The Office of Tourism Development, <http://www.tourism.go.th>

lions. The number of registered hotels in all regions of Thailand is approximately 5,000. Providing an opinion mining system on hotel reviews could be very useful for tourists to make decision on hotel choice when planning a trip.

#### 4.1 Corpus preparation

We collected customer reviews from the *Agoda* website<sup>2</sup>. The total number of reviews in the corpus is 8,436. Each review contains the name of the hotel as the title and comments in free-form text format. We designed a set of 13 main features: service, cleanliness, hotel condition, location, food, breakfast, room, facilities, price, comfort, quality, activities and security. The set of main features is designed based on the features obtained from the *Agoda* website. Some additional features, such as activities and security, are added to provide users with more dimensions.

In this paper, we focus on two main features: *breakfast* and *service*. Table 2 shows the domain-dependent lexicons related to the breakfast feature. For breakfast main feature (FEA), we include all synonyms which could be used to describe breakfast in Thai. These include English terms with their synonyms, transliterated terms and abbreviations.

The breakfast sub-features (FEA\*) are specific concepts of breakfast. Examples include “menu”, “taste”, “service” and “coffee”. It can be observed that some of the sub-features could also act as a main feature. For example, the sub-feature “service” of breakfast is also used as the main feature “service”. Providing sub-feature level could help revealing more insightful dimension for the users. However, designing multiple feature levels could be time-consuming, therefore, we limit the features into two levels, i.e., main feature and sub-feature. The polar words (POL) are also shown in the table. We denote the positive and negative polar words by placing [+] and [-] after each word. It can be observed that some polar words are dependent on sub-features. For example, the polar word “long line” can only be used for the sub-feature “restaurant”.

<sup>2</sup>Agoda website: <http://www.agoda.com>

Lexicons	Examples
Features (FEA)	ABF, Breakfast, เบรคฟาสต์ (Breakfast), อาหารเช้า (Breakfast)
Sub-features (FEA*)	เมนู (menu), รสชาติ (taste), ห้องอาหาร (restaurant), คุณภาพ (quality), บริการ (service), ปริมาณ (quantity), ขนมปัง (bread), กาแฟ (coffee), ที่นั่ง (seat), พนักงาน (waiter)
Polar words (POL)	ดี (good)[+], หลากหลาย (various)[+], สด (fresh)[+], มีคุณภาพ (with quality)[+], สะอาด (clean)[+], ประทับใจ (impressive)[+], แย่ (terrible)[-], ไม่ได้เรื่อง (awful)[-], จำเจ (repeating)[-], ต้องรอคิว (long line)[-], น้อย (little)[-], น่าเบื่อ (boring)[-], คับแคบ (confined)[-]

Table 2: Domain-dependent lexicons for the *breakfast* feature.

Table 3 shows the domain-dependent lexicons related to the service feature. The main features include synonyms, transliterated and English terms which describe the concept service. The service sub-features are, for example, “reception”, “security guard”, “maid”, “waiter” and “conciierge”. Unlike the breakfast feature, the polar words for the service feature are quite general and could mostly be applied for all sub-features. Another observation is that some of the polar words are based on Thai idiom. For example, the phrase “having rigid hands” in Thai means “impolite”. In Thai culture, people show politeness by doing the “wai” gesture.

#### 4.2 Experiments and results

Using the tagged corpus and the extracted lexicons, we construct the most frequently occurred patterns. For two main features, breakfast and service, the numbers of tagged reviews for each feature are 301 and 831, respectively. We randomly split the corpus into 80% as training set and 20% as test set. We only consider the patterns which contain both features (either main features or sub-features) and polar words. For the breakfast feature, the total number of extracted patterns is 86. For the service feature, the total number of extracted patterns is 192. Table 4 and 5 show some examples of most frequently

Lexicons	Examples
<b>Features (FEA)</b>	บริการ (service), service, เซอร์วิส (service)
<b>Sub-features (FEA*)</b>	bell boy, reception, receptionist, คนขับรถ (driver), พนักงานขับรถ (driver), เจ้าหน้าที่รักษาความปลอดภัย (security guard), รปภ. (security guard), บริการ (waiter), แผนกต้อนรับ (reception), แม่บ้าน (maid) พนักงานต้อนรับ (receptionist), พนักงานขนกระเป๋า (concierge)
<b>Polar words (POL)</b>	มีน้ำใจ (considerate)[+], สะอาด (clean)[+], ใจดี (kind)[+], ตูแผลตลอด (courteous)[+], กระตือรือร้น (eagerly)[+], อบอุ่น (warm)[+], ช่วยเหลือ (helpful)[+], จริงใจ (sincere)[+], เอาใจใส่ (courteous)[+], น่ารัก (lovely)[+], เป็นกันเอง (friendly)[+], นิสัยดี (nice)[+], ช้า (slow)[-], คอยไม่ให้อวด (arrogant)[-] ขาดคุณธรรม (deceitful)[-], จู้จี้จ้าน (nosy)[-] เฉยเมย (inattentive)[-], หงุดหงิด (grumpy)[-] มือไม่แข็ง (impolite)[-]

Table 3: Domain-dependent lexicons for the *service* feature.

occurred patterns extracted from the corpus. The symbols of the tag set are as shown in Table 1 and 2 with the tag <OTH> denoting any other words.

From the tables, two patterns which occur frequently for both features are <FEA><POL> and <FEA\*><POL>. These two patterns are very simple and show that the opinionated texts in Thai are mostly very simple. Users just simply write a word describing the feature followed by a polar word (either positive or negative) without using any verb in between. Some examples for the pattern <FEA\*><POL> are <coffee cup><dirty> and <employee><friendly>. In English, a verb “to be” (is/am/are) is usually required between <FEA\*> and <POL>.

Using the extracted patterns, we perform the dual pattern extraction process to collect the sub-features and polar words from the test data set. Table 6 shows the evaluation results of sub-features and polar words extraction for both breakfast and service features. It can be observed that the set of patterns could extract polar words (POL) with higher accuracy than sub-features (FEA\*). This could be due to the patterns used to

No.	Top-ranked "breakfast" patterns
1	<FEA><POL> <อาหารเช้า><อร่อย> <breakfast><delicious>
2	<FEA><AUX><POL> <อาหารเช้า><ก็><ดีเยี่ยม> <breakfast><"kor"><excellent>
3	<FEA><POL><DEG> <อาหารเช้า><แย่มาก> <breakfast><terrible><very>
4	<FEA*><POL> <ถ้วยกาแฟ><สกปรก> <coffee cup><dirty>
5	<FEA><POL><POL> <อาหารเช้า><สะอาด><มีคุณภาพ> <breakfast><clean><with quality>

Table 4: Top-ranked *breakfast* patterns with examples

describe the polar words are straightforward and not complicated. This is especially true for the case of breakfast feature in which the accuracy is approximately 95%.

No.	Top-ranked "service" patterns
1	<FEA*><POL> <พนักงาน><เป็นมิตร> <employee><friendly>
2	<FEA><POL> <บริการ><ประทับใจ> <service><impressive>
3	<FEA*><FEA><POL> <พนักงานขับรถ><บริการ><ดี> <driver><service><good>
4	<FEA*><FEA><POL><DEG> <พนักงาน><บริการ><สุภาพ><มาก> <employee><service><polite><very>
5	<FEA*><OTH><POL> <พนักงาน><ทุกคน><ยิ้มแย้มแจ่มใส> <employee><everyone><smiling>

Table 5: Top-ranked *service* patterns with examples

### 4.3 Discussion

Table 7 and 8 show some examples of challenging cases for breakfast and service features, respectively. The polar words shown in both tables are very difficult to extract since the patterns can

Feature	Accuracy (%)	
	FEA*	POL
Breakfast	80.00	95.74
Service	82.56	89.29

Table 6: Evaluation results of features and polar words extraction.

not be easily captured. The difficulties are due to many reasons including the language semantic and the need of world knowledge. For example, in case #5 of service feature, the whole phrase can be interpreted as “attentive”. It is difficult for the system to generate the pattern based on this phrase. Another example is case #4 of both tables, the customers express their opinions by comparing to other hotels. To analyze the sentiment correctly would require the knowledge of a particular hotel or hotels in specific locations.

No.	Some difficult cases of "breakfast" feature
1	มีแต่ซีเรียลนมกับชากาแฟ (only provide cereals, milk, tea and coffee)
2	ปิด 10 โมงเร็วไป (close at 10AM, too soon)
3	แบบว่าเต็มสิบเอาไปเกือบเต็ม (almost 10 out of 10)
4	เหมือนไปพักที่โรงแรมในกรุงเทพ (just like staying at hotels in Bangkok)
5	ให้แค่ห้องละ 1 คน (only limit 1 person per room)

Table 7: Examples of difficult cases of *breakfast* feature

## 5 Conclusion and future work

We proposed a framework for constructing Thai opinion mining resource with a case study on hotel reviews. Two sets of lexicons, domain-dependent and domain-independent, are designed to support the pattern extraction process. The proposed method first constructs a set of patterns from a tagged corpus. The extracted patterns are then used to automatically extract and collect more sub-features and polar words from an untagged corpus. The performance evaluation

No.	Some difficult cases of "service" feature
1	ลืมให้เสื้อคลุมอาบน้ำ (forget to provide a bathrobe)
2	welcome drink ไป 3 ท่าน เซิร์ฟเพียง 1 ท่าน (for welcome drink, 3 persons, but only serve 1)
3	ควรลดเสียงในการลากหรือใช้อุปกรณ์ที่ส่งเสียงดัง (should reduce noise from dragging and using tools)
4	ระดับดุสิตธานี (like Dusit Thani hotel)
5	บริการจำได้ว่าแขกชอบ/ไม่ชอบอะไร (waiter could remember guests' preferences)

Table 8: Examples of difficult cases of *service* feature

was done with a collection of hotel reviews obtained from a hotel reservation website. From the experimental results, polar words could be extracted more easily than sub-features. This is due to the polar words often appear in specific positions with repeated contexts in the opinionated texts. In some cases, extraction of sub-features and polar words are not straightforward due to the difficulties in generalizing patterns. For example, some subjectivity requires complete phrases to describe the polarity. In some cases, the sub-features are not explicitly shown in the sentence. For future work, we plan to complete the construction of the corpus by considering the rest of main features. Another plan is to include the semantic analysis into the pattern extraction process. For example, the phrase “forget something” could imply negative polarity for the service feature.

## References

- Banea, Carmen, Rada Mihalcea, Janyce Wiebe, and Samer Hassan. 2008. Multilingual subjectivity analysis using machine translation. *Proc. of the 2008 empirical methods in natural language processing*, 127–135.
- Beineke, Philip, Trevor Hastie and Shivakumar Vaithyanathan. 2004. The sentimental factor: improving review classification via human-provided information. *Proc. of the 42nd Annual Meeting on Association for Computational Linguistics*, 263–270.

- Cooke, J.R. 1992. Thai sentence particles: putting the puzzle together. *Proc. of the The Third International Symposium on Language and Linguistics*, 1105–1119.
- Dave, Kushal, Steve Lawrence and David M. Pennock. 2003. Mining the peanut gallery: opinion extraction and semantic classification of product reviews. *Proc. of the 12th international conference on World Wide Web*, 519–528.
- Ding, Xiaowen, Bing Liu and Philip S. Yu. 2008. A holistic lexicon-based approach to opinion mining. *Proc. of the int. conf. on web search and web data mining*, 231–240.
- Hu, Minqing and Bing Liu. 2004. Mining and summarizing customer reviews. *Proc. of the 10th ACM SIGKDD international conference on Knowledge discovery and data mining*, 168–177.
- Jin, Wei, Hung Hay Ho and Rohini K. Srihari. 2009. OpinionMiner: a novel machine learning system for web opinion mining and extraction. *Proc. of the 15th ACM SIGKDD*, 1195–1204.
- Kim, Soo-Min and Eduard Hovy. 2004. Determining the sentiment of opinions. *Proc. of the 20th international conference on Computational Linguistics*, 1367–1373.
- Qiu, Guang, Bing Liu, Jiajun Bu, and Chun Chen. 2009. Expanding domain sentiment lexicon through double propagation. *Proc. of the 21st International Joint Conferences on Artificial Intelligence*, 1199–1204.
- Jijkoun, Valentin and Katja Hofmann. 2009. Generating a non-English subjectivity lexicon: relations that matter. *Proc. of the 12th Conference of the European Chapter of the Association for Computational Linguistics*, 398–405.
- Kanayama, Hiroshi and Tetsuya Nasukawa. 2006. Fully automatic lexicon expansion for domain-oriented sentiment analysis. *Proc. of the 2006 Conference on Empirical Methods in Natural Language Processing*, 355–363.
- Ku, Lun-Wei and Hsin-Hsi Chen. 2007. Mining opinions from the Web: Beyond relevance retrieval. *Journal of American Society for Information Science and Technology*, 58(12):1838–1850.
- Ku, Lun-Wei, Ting-Hao Huang and Hsin-Hsi Chen. 2009. Using morphological and syntactic structures for Chinese opinion analysis. *Proc. of the 2009 empirical methods in natural language processing*, 1260–1269.
- Liu, Bing, Minqing Hu and Junsheng Cheng. 2005. Opinion observer: analyzing and comparing opinions on the Web. *Proc. of the 14th World Wide Web*, 342–351.
- Pang, Bo, Lillian Lee and Shivakumar Vaithyanathan. 2002. Thumbs up?: sentiment classification using machine learning techniques. *Proc. of the ACL-02 conf. on empirical methods in natural language processing*, 79–86.
- Popescu, Ana-Maria and Oren Etzioni. 2005. Extracting product features and opinions from reviews. *Proc. of the conf. on human language technology and empirical methods in natural language processing*, 339–346.
- Riloff, Ellen and Janyce Wiebe. 2003. Learning extraction patterns for subjective expressions. *Proc. of the 2003 conference on empirical methods in natural language processing*, 105–112.
- Sarmiento, Luís, Paula Carvalho, Mário J. Silva, and Eugénio de Oliveira. 2009. Automatic creation of a reference corpus for political opinion mining in user-generated content. *Proc. of the 1st CIKM workshop on topic-sentiment analysis for mass opinion*, 29–36.
- Turney, Peter D. 2002. Thumbs up or thumbs down?: semantic orientation applied to unsupervised classification of reviews. *Proc. of the 40th ACL*, 417–424.
- Wan, Xiaojun. 2009. Co-training for cross-lingual sentiment classification. *Proc. of the joint conf. of ACL and IJCNLP*, 235–243.
- Wiebe, Janyce and Ellen Riloff. 2005. Creating subjective and objective sentence classifiers from unannotated texts. *Proc. of Conference on Intelligent Text Processing and Computational Linguistics*, 486–497.
- Wiebe, Janyce, Theresa Wilson and Claire Cardie. 2005. Annotating expressions of opinions and emotions in language. *Language Resources and Evaluation*, 39(2-3):165–210.
- Wilson, Theresa, Janyce Wiebe and Paul Hoffmann. 2009. Recognizing contextual polarity: An exploration of features for phrase-level sentiment analysis. *Comput. Linguist.*, 35(3):399–433.
- Yu, Hong and Vasileios Hatzivassiloglou. 2003. Towards answering opinion questions: separating facts from opinions and identifying the polarity of opinion sentences. *Proc. of the Conference on Empirical Methods in Natural Language Processing*, 129–136.