Textual Demand Analysis: Detection of Users' Wants and Needs from Opinions

Hiroshi Kanayama Tetsuya Nasukawa

Tokyo Research Laboratory, IBM Japan, Ltd.

1623-14 Shimotsuruma, Yamato-shi, Kanagawa-ken, 242-8502 Japan

{hkana, nasukawa}@jp.ibm.com

Abstract

This paper tackles *textual demand analysis*, the task of capturing what people want or need, rather than identifying what they like or dislike, on which much conventional work has focused. It exploits syntactic patterns as clues to detect previously unknown demands, and requires domain-dependent knowledge to get high recall. To build such patterns we created an unsupervised pattern induction method relying on the hypothesis that there are commonly desired aspects throughout a domain corpus. Experimental results show that the proposed method detects twice to four times as many demand expressions in Japanese discussion forums compared to a baseline method.

1 Introduction

Increasingly we can access many opinions towards products, services, or companies through electronic documents including questionnaires, call logs, and other consumer-generated media (CGM) such as Internet discussion forums and blogs. It is very important for companies to get insights from their customers' opinions by analyzing such documents in large numbers.

The most popular way to utilize such data has involved *sentiment analysis* (SA) (Nasukawa and Yi, 2003; Yi et al., 2003), which is the task of recognizing the writers' feelings as expressed in positive or negative comments. Typically, a SA system focuses on expressions to identify the strong

or weak points of the subjects as in (1) or in the writers' evaluations as in (2).

- (1) I think the pictures are beautiful.
- (2) I don't like this camera very much.

Here we call them *polar expressions* because they convey positive or negative polarities. By counting the polar expressions related to products or services, one can quantitatively compare the goodness of competing services, find the drawbacks of specific products, and so on.

In addition to polar expressions, there are other types of expressions that provide valuable information, especially for the supplier side rather than the consumer side. Examples (3) and (4) express the demands of the writers.

- (3) I'd be happy if it is equipped with
a crisp LCD.
- (4) I'm waiting for a single-lens reflex less
than 30,000 yen to come on the market.

We call such expressions "*demand expressions*", and the underlined phrases "*demand targets*."

While sentiment analysis reveals evaluations of existing products or services, the task proposed here, textual demand analysis¹, gives more direct suggestions to companies: things they should do to attract customers. For example, by investigating demand targets, companies can add new functions to products on the market or plan new services to satisfy customers. These activities should lead to positive evaluations in the future.

Interestingly, demand expressions may be noise in sentiment analysis, because the demand expressions do not actually convey positive or neg-

© 2008. Licensed under the *Creative Commons Attribution-Noncommercial-Share Alike 3.0 Unported* license (<http://creativecommons.org/licenses/by-nc-sa/3.0/>). Some rights reserved.

¹Note that textual demand analysis is different from the demand analysis in the field of marketing or software engineering.

Figure 1: A demand analysis system and the flow of the pattern induction method.

ative evaluations of existing products or services, even though these demand expressions often contain positive or negative words, as in Example (3) which contains the positive expressions ‘happy’ and ‘crisp LCD’.

The detection of novel demand targets requires deep syntactic information because such demand targets themselves can not be predefined. For example, to regard the underlined parts of (3) and (4) as demand targets, the non-underlined parts of these sentences have to be properly interpreted as triggers. This is a major difference from sentiment analysis where the polar expressions can be defined in the lexicon.

The left parts of Figure 1 illustrate the concepts of a system that visualizes the users’ demands described in the input opinion data, where the main analysis component processes the documents and extracts the demand targets. The output of the system is created by a demand outliner, which the company uses to grasp the trends of consumers’ demands.

The syntactic patterns that can be used as clues to demand expressions depend on the topic domain or the writing style. To organize this linguistic knowledge we propose an unsupervised induction method for syntactic patterns. The right part of Figure 1 shows the flow of pattern induction.

In the next section, we review related work, and Section 3 defines our task more formally. In Section 4 we describe a naive approach to the task and Section 5 shows a form of unsupervised pattern induction used to cover more demand expressions. Section 6 gives the experimental results and we conclude in Section 7.

2 Related Work

Sentiment analysis (SA) and related topics have been extensively studied in recent years. The textual demand analysis proposed in this paper shares some properties with phrase-level SA, the detection of sentiments and evaluations expressed in phrases, rather than document-level SA, the classification of documents in terms of goodness of reputation. Yi et al. (2003) pointed out that the multiple sentiment aspects in a document should be extracted, and Nasukawa and Yi (2003) clarified the need for deep syntactic analysis for the phrase-level SA.

The acquisition of clues is a key technology in these research efforts, as seen in learning methods for document-level SA (Hatzivassiloglou and McKeown, 1997; Turney, 2002) and for phrase-level SA (Wilson et al., 2005; Kanayama and Nasukawa, 2006).

As well as the sentiment expressions leading to evaluations, there are many semantic aspects to be extracted from documents which contain writers’ opinions, such as subjectivity (Wiebe and Mihalcea, 2006), comparative sentences (Jindal and Liu, 2006), or predictive expressions (Kim and Hovy, 2007). However, the extraction of the contents of writers’ demands which this paper handles is less studied while this type of information is very valuable for commercial applications.

For the tasks of information extraction and relation extraction, bootstrapping approaches have been proven successful (Yangarber, 2003; Pantel and Pennacchiotti, 2006). The pattern induction method in this paper exploits their ideas, but their application to the demand detection is not trivial, because some instances of demands are previously unknown and do not appear frequently, so they have to be abstracted effectively.

The work by Inui et al. (2003) handles semantics of a type similar to ours. They aimed to detect the requests in the responses to open-ended questionnaires, seeking *direct requests* such as “... てほしい” (‘[I] would like you to ...’) and other forms which can be paraphrased as direct requests. They classified sentences into requests or non-requests, where their source documents were responses to questionnaires, and where more than 60% of the utterances could be regarded as requests of some sort. In contrast, our method detects the content of the demands in the form of noun phrases, and handles more general target documents including

CGM documents that contain much more diverse utterances.

3 Task Definition

As shown in Section 1, our goal is to create a system to enumerate in an easily understandable way the demand targets in the input text. This section describes the definition of a demand target and its representation format.

3.1 Demand targets

Demands or requests written in opinion texts can be represented by verb phrases as in *e.g.* ‘I want to V.’ and ‘I want you to V.’, or noun phrases as in ‘I want N.’² In this paper we focus on the last type, *i.e.* noun phrases which represent desired objects, because they are easier to aggregate and grasp than verb phrases. Another reason is that some demands represented with a verb phrase only describe the objects that are desired. For example, ‘I want to buy N’ and ‘I want you to provide N’ can be simply interpreted as meaning that N is what the writer wants. We call such a noun phrase a *demand target*, and these are the outputs of our system.

For applications, the demand targets to be detected by the system depend on the type of input documents. For example, from a consumers’ forum on digital cameras, the underlined parts in Examples (3) and (4) from Section 1 apparently describe the writer’s demands, so they are valuable information for such users of demand analysis such as the makers of digital camera. However, the request in Example (5) does not express the author’s demands about any digital camera, but rather it is written for other participants in the forum. This type should be excluded from the output.

- (5) Please give me a good advice.

In contrast, when the responses to a questionnaire about the activities of an organization are processed, statements such as Example (5) should be regarded as a demand target, since the writer wrote it as a request to the sponsor of the questionnaire and the ‘advice’ is indeed a thing that can be provided by the sponsoring organization.

3.2 Representation of demand targets

A demand target tends to be expressed by a noun phrase with modifiers, as seen in Examples (3) and

²‘V’ and ‘N’ indicate a verb phrase and a noun phrase, respectively.

Figure 2: Syntactic trees for the sentences (6) and (7). ‘*’ indicates the headword of the demand targets.

(4), rather than by a single noun. The headwords of such phrases (*e.g.* ‘LCD’ in (3) and ‘reflex’ in (4)) represent the main categories of the demanded objects, but they are not distinctive enough to recognize as knowledge of the authors’ demands.

Therefore the key task of this research was to find ways to markup the headword of a noun phrase that represents the content of a demand in the syntactic parse tree. Examples (6) and (7) are the original Japanese sentences corresponding to Examples (3) and (4).

- (6) 私としては 綺麗な液晶 があると嬉しい。
‘I’d be happy if it is equipped with a crisp LCD.’
(7) 3万円で買える一眼レフ を出してほしい。
‘[I’m] waiting for a single-lens reflex
less than 30,000 yen to come on the market.’

Figure 2 represents the parse trees corresponding to sentences (6) and (7), where the demand targets are identified by the mark ‘*’.

This simple representation is advantageous for both the collection of and the deeper investigation of the demand targets. One can easily grasp the content of a demand if the application shows the whole surface structure of the subtree under ‘*’ in the tree, *e.g.* the underlined parts of Examples (6) and (7). At the same time the tree structure supports the further analysis of the trends of the demands by picking up the headwords or modifiers prominent in the subtrees that were detected as demand targets.

4 Baseline Method of Textual Demand Analysis

This section describes an algorithm to extract demand targets with high precision and describes a preliminary experiment to measure the performance.

Figure 3: (a) is a demand pattern which indicates that the noun in N^* is detected as a demand target. (b) and (c) are auxiliary patterns, where V indicates the node matches any verb.

4.1 Syntactic patterns and top-down matching

A major purpose of textual demand analysis is to discover novel demands embedded in the text, thus the triggers of their detection should not be a predefined set of demand targets but should be their surrounding syntactic information. We use two types of syntactic patterns shown in Figure 3. Those patterns are compared with the syntactic tree as the parsing result of the input sentence.

The pattern (a) in Figure 3 is a *demand pattern*, which is used to search for demand targets. The node with the ‘*’ mark indicates the corresponding node will be the headword of a demand target. Hence we write the pattern (a) as “ N^* -が-欲しい” for simplicity. The patterns are applied in a top-down manner, that is, initially the top node of the input tree is examined to see whether or not the node and its combination of children nodes match with one of the patterns in the pattern repository. This method supports higher precision in the detection than the surface pattern matching. For example, the expression “不良品が欲しい人はいない” (‘There is no one who wants low quality goods’) should not be misunderstood to express a demand.

The patterns (b) and (c) in Figure 3 are *auxiliary patterns*. These are used to apply the demand patterns to nodes other than the root of the syntactic tree. For example, by applying the patterns (b) and (c), the pattern (a) can then be applied to the expressions “Nが欲しいと思っています” (‘I think that I want N’) and “Nが欲しいものの、お金が足りない” (‘Though I want N, I don’t have enough money’), respectively, even though “ N^* -が-欲しい” doesn’t appear at the top of the trees. In other words, the auxiliary patterns contribute to generate variations of the demand patterns.

In addition, simple rules can be applied to filter out certain meaningless outputs. When a noun phrase that matched to the ‘*’ part of the demand

Table 1: The result on the small gold standard with DP_1 . PM signifies surface pattern matching, TM signifies tree matching. ‘+AP’ means that auxiliary patterns are used.

Method	Precision		Recall	
PM	39%	(14/36)	25%	(14/56)
TM	92%	(11/12)	20%	(11/56)
TM+AP	94%	(17/18)	30%	(17/56)

pattern was a pronoun or very common noun (*e.g.* ‘camera’ in the camera domain) without any modifier, it is not output as a demand target.

4.2 Preliminary experiment

We conducted a preliminary experiment to assess the feasibility of our approach.

We prepared a small gold-standard dataset which consists of 1,152 sentences from a discussion forum on digital cameras, for which two human annotators attached marks to the demand targets. There were 56 demand targets that at least one of the annotators detected, and the sentence-level agreement value³ was $\kappa = 0.73$, which is regarded as a good level of agreement. There was no sentence in which the two annotators attached marks to different nouns.

First, we made a minimum set of demand patterns DP_1 , which contained only one basic pattern “ N^* -が-欲しい⁴” (‘I want N^* ’).

To see the effect of the top-down matching and the auxiliary patterns described in Section 4.1, demand targets in the gold-standard corpus were automatically detected using three methods: pattern matching with surface strings like “が欲しい” (PM), tree matching without the auxiliary patterns (TM), and tree matching with the auxiliary patterns⁵ (TM+AP).

Table 1 shows the results. The top-down matching on the syntactic tree resulted in much higher precision than the surface pattern matching, and the auxiliary patterns improved the recall. The only misdetection in the tree matching method was due to an error in the sentence segmentation.

However, all of them show low recall values,

³The agreement on whether or not the sentence has a demand target.

⁴Apparent character variations like ‘欲しい’ and ‘ほしい’, and alternative forms of particles were aggregated in the parsing process.

⁵A total of 95 auxiliary patterns which Kanayama et al. (2004) used for the sentiment analysis.

Table 2: The list of augmented demand patterns DP_q .

N^* -が-欲しい (I want N^*), N^* -を-望む (I hope N^*), N^* -を-お願い-する (Please [give] N^*), N^* -を-願う (I wish N^*), N^* -を-して-下さい (Please do N^*), N^* -を-頼む (I ask [you] N^*), N^* -が-ある-べきだ (N^* should be), N^* -を-して-頂きたい (Please do N^*)

Table 3: The result with the minimum set of demand patterns DP_1 and the larger set DP_q .

Patterns	Precision	Recall
DP_1	94% (17/18)	30% (17/56)
DP_q	78% (18/22)	32% (18/56)

since only one demand pattern was used. To make the recall higher, we created another set of demand patterns DP_q listed in Table 2, which are generally used as clues for the request expressions in the analysis of responses to open-ended questionnaires. They are derived from the previous work on request classification (Yamamoto et al., 2007).

The result in Table 3 shows that the patterns newly added in DP_q do not perform well. This is because these patterns appear in responses to questionnaires but are not suitable for the writing styles used in discussion forums, as mentioned in Section 3.1.

Therefore we used the TM+AP method with the minimum pattern set DP_1 as the baseline in this paper rather than the pattern set DP_q .

5 Automatic Pattern Induction

The preliminary experiment in Section 4.2 showed that high precision can be obtained by the top-down matching method, and at the same time, revealed the difficulty in building demand patterns to achieve high recall. To overcome this problem, we devised an automatic pattern induction algorithm.

5.1 Frequent fragments of demand targets

Here we make an assumption that there are commonly desired aspects or things throughout a domain corpus. Based on this assumption, we extract the syntactic fragments which appear relatively frequently as the elements of demand targets from the training corpus in a specific domain.

First we obtain demand targets from the domain corpus, in this case from the discussion forum on digital cameras, by using the baseline method with

Figure 4: Sample extraction of demand instances from a demand target detected by the system. ‘N’ denotes the wildcard for any nouns.

the pattern set DP_1 . Next, *demand instances* are extracted from each demand target. A demand instance is a one-to-three-node subtree of a demand target, and shares the root node with the demand target. The root node that is modified by one or more nodes can be replaced with a wildcard. Figure 4 shows a sample extraction of demand instances from a demand target “3万円で購入できる一眼レフ” (‘single-lens reflex less than 30,000 yen’), where nine demand instances are extracted, and four of them have a wildcard at the root node.

The demand instances which appear more than θ_f times in the corpus are selected as *frequent demand instances* (FDIs), and each FDI is assigned the following reliability value r_i :

$$r_i = \frac{\text{freq}_{DT}(i)}{\text{freq}(i)} \quad (8)$$

where ‘ $\text{freq}(i)$ ’ denotes the frequency of the instance subtree in the whole corpus and ‘ $\text{freq}_{DT}(i)$ ’ means the i ’s frequency in the demand targets. The notion of an instance’s reliability is inspired by the method of relation extraction (Pantel and Pennacchiotti, 2006), but our usage is different from theirs. Here we use the reliability value only for the relative comparison among demand instances, so normalization of the values is not needed. The intrinsic difference from the instance of relation extraction is that the demand instances are not the final outputs of the extraction, but are just triggers for new demand patterns.

When θ_f was set to 5, a total of 42 FDIs which had reliabilities above 0.01 were picked from 150,000 postings in the discussion forum on digital cameras. Table 4 shows examples of FDIs.

5.2 Frequent patterns as the clue

The FDIs with high reliability correspond to aspects which are likely to be demanded, therefore the syntactic structures which often govern such

Table 4: Examples of frequent demand instances (FDIs) with $r_i > 0.01$ in the digital camera domain. ‘-’ denotes the split of nodes.

れる-デジカメ (‘digital camera - which can’)
mm-レンズ (‘mm - lens’), 後継機 (‘newer model’),
良い-もの (‘good - thing’), 明るい-レンズ (‘sharp lens’),
綺麗な-N (‘beautiful - N’), マクロ-N (‘macro - N’), ...

Figure 5: An example of extraction of a candidate demand pattern. From the sentence “小さくて接写ができるカメラを早く出して欲しい。” (‘I want a small camera with the close-up function sold earlier’), the CDP “N* -を-出して-欲しい” (the oval) is extracted, triggered by the FDI “できる-カメラ” (the square).

FDIs are expected to be clues for detecting additional demands.

Following this expectation, the *candidate demand patterns* (hence CDPs) are extracted. A CDP is a subtree that connects the head of an FDI and the root of the syntactic tree, or auxiliary patterns which cover the root of the tree. A CDP forms a node sequence without a branch, and corresponds to a sentence-final expression in Japanese that usually conveys modality information. Figure 5 illustrates the extraction of a CDP from a syntactic tree triggered by an FDI.

For each CDP extracted in this way, the reliability is determined by Equation (9):

$$r_p = \sum_{i \in FDI} \frac{\text{freq}(i, p) \cdot r_i}{\text{freq}(p)} \quad (9)$$

where $\text{freq}(i, p)$ denotes the frequency of the collocation of the instance i and the pattern p , and $\text{freq}(p)$ is the frequency of p in the entire corpus. These ratios are summed up over all of the FDIs, weighted by the reliability of the instance. Also r_p is used only for the relative comparison among CDPs, so it is not normalized to be in the range [0,1].

Table 5 shows the CDPs which appeared 10 times or more and their reliability values. Some

Table 5: The extracted candidate demand patterns (CDPs) sorted by their reliability.

Candidate Demand Pattern	Reliability
N*-が-あれば-良い (be good if it includes N*)	1.70×10^{-2}
N*-を-買って-下さい (please buy N*)	1.48×10^{-2}
N*-が-あると-良い (be good if it includes N*)	1.12×10^{-2}
N*-が-あると-便利だ (be convenient if it includes N*)	4.81×10^{-3}
N*-を-買おうと-思っ-ている ([I’m] going to buy N*)	3.32×10^{-3}
N*-が-あれば-便利だ (be convenient if it includes N*)	3.00×10^{-3}
N*-を-出して-欲しい ([I] want N* to be sold)	1.88×10^{-3}
N*-が-希望-だ (N* is longed for)	1.34×10^{-3}
N*-を-薦める ([I] recommend N*)	1.06×10^{-3}
N*-の-購入-を-検討して-いる ([I’m] thinking about buying N*)	8.92×10^{-4}
N*-が-不足する (N* is lacking)	3.53×10^{-4}
⋮	⋮
N*-を-使う (to use N*)	5.28×10^{-5}
N*-を-購入する (to purchase N*)	4.31×10^{-5}
N*-で-撮る (to take [pictures] with N*)	6.06×10^{-6}
⋮	⋮

of them apparently reflect the writing style of the discussion forum and the digital camera domain. The effect of these patterns will be verified in Section 6.

6 Evaluation

We conducted experiments to assess the contributions of the candidate demand patterns acquired in Section 5.

6.1 Experimental setup

In Section 4.2 we created a gold-standard dataset with human annotations, however, the number of annotation is not enough to fairly compare the several methods due to the sparseness of the demand targets in the original corpus, and it would be very laborious to prepare a larger annotated dataset.

Therefore we used an unannotated corpus for the evaluation in this section. A total of 50,000 postings in the digital camera forum⁶ were processed by each method, and 100 demand targets were randomly selected from the system output for each trial and a human evaluator decided for each demand target whether or not it referred to any demanded object related to the domain.

⁶They are separate from the 150,000 postings used for the training.

Table 6: The evaluations when CDPs with reliability more than θ were used.

θ	Precision	Recall	F1
∞	93%	30%	0.45
10^{-2}	91%	31%	0.46
10^{-3}	87%	37%	0.52
10^{-4}	68%	59%	0.63
10^{-5}	33%	57%	0.42

In this way, the precision can be computed directly, and the recall can be estimated as follows:

$$\text{Rec}(T) \simeq \frac{\text{Num}(T)\text{Prec}(T)\text{Rec}(B)}{\text{Num}(B)\text{Prec}(B)} \quad (10)$$

where $\text{Rec}()$, $\text{Prec}()$, and $\text{Num}()$ denote the recall, the precision, and the number of demand targets detected by the system in the entire test corpus, respectively, and T and B denote the tested method and the baseline method, respectively. $\text{Prec}(B)$ is assumed to be 30% as observed in the preliminary experiment.

6.2 Effect of new demand patterns

The CDFs for the digital camera domain that were acquired with the method in Section 5 are tested by varying the threshold θ . The CDFs which have reliability value more than θ were added to the demand pattern set. Table 6 shows the results. The baseline method was without any newly acquired demand patterns (*i.e.* $\theta = \infty$), thus it is the same condition as the DP_1 in the preliminary experiment in Section 4.2.

When θ was set to 10^{-3} , the recall increased drastically with little harm to the precision. The value of $\theta = 10^{-5}$ did not work well because the precision was very low and the increase of the recall was limited. The value $\theta = 10^{-4}$ performed best in terms of the F1 value.

We observed the demand targets derived from the new demand patterns. In most cases desirable functions and features of the digital cameras were successfully obtained from conditional positive expressions such as “ N^* -が-あると-便利だ” (‘be convenient if it includes N^* ’). Also, preferred machines were clarified by the expression “ N^* -を-買って-下さい” (‘please buy N^* ’) which is in a postings to recommend something to other users. On the other hand, the expression “ N^* -が-不足する” (‘ N^* is lacking’) tend to result in noisy demand targets.

Table 7: The extracted CDPs and their reliability for the domain of company’s questionnaire.

Candidate Demand Pattern	Reliability
N^* -を-望む ([I] hope N^*)	3.10×10^{-2}
N^* -を-願いたい ([I] want to ask for N^*)	1.37×10^{-2}
N^* -を-願う ([I] wish N^*)	4.92×10^{-3}
N^* -が-望ま-れる (N^* is hoped for)	1.45×10^{-3}
N^* -を-提供して-欲しい ([I] want N^* to be provided)	1.04×10^{-3}
N^* -が-必要-と-考える ([I] think N^* is necessary)	6.72×10^{-4}
⋮	⋮
N^* -が-不足して-いる (N^* is lacking)	1.02×10^{-4}
N^* -が-悪い (N^* is bad)	7.22×10^{-5}
⋮	⋮

We also tried the iterative acquisition using the newly acquired patterns, but the useful patterns were rarely acquired in the second step. This is because FDIs cannot be definitive triggers, and the first seed pattern “ N^* -が-欲しい” (‘I want N^* ’) was a prominently reliable pattern compared with the other demand patterns.

6.3 Applicability to other demand targets

The pattern induction method was expected to have advantage that domain-dependent patterns can be acquired, and indeed some of the patterns were specific for the digital camera domain as shown in Table 5. To see the applicability of our algorithm to other domains or other types of text, the pattern induction was tested on another corpus.

The responses to a questionnaire about collaboration process in a company were used as the corpus. Unlike the discussion forum on digital cameras, the writing style of direct request such as ‘Please provide N^* ’ was observed frequently, and the demand targets are much more dense in the corpus. Table 7 shows the CDPs acquired in this domain, and Table 8 shows the evaluation where 25,000 and 5,000 sentences were used for the training and the test, respectively.

As a result, higher precision was achieved in this domain than in the digital camera domain, because the demands are stated more explicitly in the responses to the questionnaires. Unlike in the digital camera domain, the pattern “ N^* -が-不足して-いる” (‘ N^* is lacking’) worked well because in many cases of this domain what are lacking equal to what are needed. For example, “効率の良い話し合いが不足している” (‘effective discussion is lack-

Table 8: The evaluations for the company questionnaire domain. The initial recall was estimated as 15%. DC10⁻⁴ means that the CDPs for the digital camera domain ($\theta = 10^4$) were used.

θ	Precision	Recall	F1
∞	98%	15%	0.26
10 ⁻²	96%	24%	0.39
10 ⁻³	92%	30%	0.45
10 ⁻⁴	85%	71%	0.77
10 ⁻⁵	41%	73%	0.53
DC10 ⁻⁴	72%	40%	0.51

ing’) implies that the effective discussion is a demand.

When the demand patterns acquired in the digital camera domain (DC10⁻⁴) were used, the increase of the recall was limited. These results support the value of the unsupervised pattern induction method which works for any domain when only a raw domain corpus is provided.

7 Conclusion

We formalized the task *textual demand analysis* and proposed a pattern induction method to increase the coverage of the automatic detection of demand targets. The pattern induction proposed here allows for the discovery of novel demands that can be represented by various forms of noun phrases, though they were triggered by frequently appeared syntactic fragments. Beyond sentiment analysis, textual demand analysis provides valuable knowledge for industries, clarifying not only the favorable aspects in the current products, but also the essential features in the future.

References

- Hatzivassiloglou, Vasileios and Kathleen R. McKeown. 1997. Predicting the semantic orientation of adjectives. In *Proceedings of the 35th ACL and the 8th EACL*, pages 174–181.
- Inui, Hiroko, Masao Utiyama, and Hitoshi Isahara. 2003. Criterion for judging request intention in response texts of open-ended questionnaires. In *Proceedings of the second international workshop on Paraphrasing*, pages 49–56.
- Jindal, Nitin and Bing Liu. 2006. Mining comparative sentences and relations. In *Proceedings of the 21st National Conference on Artificial Intelligence (AAAI2006)*.
- Kanayama, Hiroshi and Tetsuya Nasukawa. 2006. Fully automatic lexicon extraction for domain-oriented sentiment analysis. In *Proceedings of the 2006 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 355–363.
- Kanayama, Hiroshi, Tetsuya Nasukawa, and Hideo Watanabe. 2004. Deeper sentiment analysis using machine translation technology. In *Proceedings of 20th International Conference on Computational Linguistics (COLING)*, pages 494–500.
- Kim, Soo-Min and Eduard Hovy. 2007. Crystal: Analyzing predictive opinions on the Web. In *Proceedings of the 2007 Joint Conference on Empirical Methods in Natural Language Processing and Computational Natural Language Learning (EMNLP-CoNLL)*, pages 1056–1064.
- Nasukawa, Tetsuya and Jeonghee Yi. 2003. Sentiment analysis: Capturing favorability using natural language processing. In *Proceedings of the Second International Conferences on Knowledge Capture*, pages 70–77.
- Pantel, Patrick and Marco Pennacchiotti. 2006. Espresso: leveraging generic patterns for automatically harvesting semantic relations. In *ACL ’06: Proceedings of the 21st International Conference on Computational Linguistics and the 44th annual meeting of the ACL*, pages 113–120.
- Turney, Peter D. 2002. Thumbs up or thumbs down? Semantic orientation applied to unsupervised classification of reviews. In *Proc. of the 40th ACL Conf.*, pages 417–424.
- Wiebe, Janyce and Rada Mihalcea. 2006. Word sense and subjectivity. In *ACL ’06: Proceedings of the 21st International Conference on Computational Linguistics and the 44th annual meeting of the ACL*, pages 1065–1072.
- Wilson, Theresa, Janyce Wiebe, and Paul Hoffmann. 2005. Recognizing contextual polarity in phrase-level sentiment analysis. In *Proceedings of HLT Conference and Conference on EMNLP*, pages 347–354, October.
- Yamamoto, Mizuki, Takashi Inui, Hiroya Takamura, Satoko Marumoto, Hiroko Otsuka, and Manabu Okumura. 2007. Extracting demands and their reasons in answers to open-ended questionnaires. In *The 13th Annual Meeting of The Association for Natural Language Processing*. (in Japanese).
- Yangarber, Roman. 2003. Counter-training in the discovery of semantic patterns. In *Proceedings of the 41st annual meeting of the ACL*, pages 343–350.
- Yi, Jeonghee, Tetsuya Nasukawa, Razvan Bunescu, and Wayne Niblack. 2003. Sentiment analyzer: Extracting sentiments about a given topic using natural language processing techniques. In *Proceedings of the Third IEEE International Conference on Data Mining*, pages 427–434.