Mining comparative opinions from customer reviews for Competitive Intelligence

Kaiquan Xu [a,*], Stephen Shaoyi Liao [a], Jiexun Li [b], Yuxia Song [a]

[a] *Department of Information Systems, City University of Hong Kong, Kowloon, Hong Kong SAR, Hong Kong*
[b] *College of Information Science and Technology, Drexel University, Philadelphia, PA 19104, USA*

ARTICLE INFO

ABSTRACT

Competitive Intelligence is one of the key factors for enterprise risk management and decision support. However, the functions of Competitive Intelligence are often greatly restricted by the lack of sufficient information sources about the competitors. With the emergence of Web 2.0, the large numbers of customer-generated product reviews often contain information about competitors and have become a new source of mining Competitive Intelligence. In this study, we proposed a novel graphical model to extract and visualize comparative relations between products from customer reviews, with the interdependencies among relations taken into consideration, to help enterprises discover potential risks and further design new products and marketing strategies. Our experiments on a corpus of Amazon customer reviews show that our proposed method can extract comparative relations more accurately than the benchmark methods. Furthermore, this study opens a door to analyzing the rich consumer-generated data for enterprise risk management.

© 2010 Elsevier B.V. All rights reserved.

1. Introduction

Competitive Intelligence (CI) involves the early identification of potential risks and opportunities by gathering and analyzing information about the environment to support managers in making strategic decisions for an enterprise [33]. Most firms realize the importance of CI in enterprise risk management and decision support, and invest a large amount of money in CI. A survey from the American Futures Group consulting firm indicates that 82% of large enterprises and over 90% of the Forbes top 500 global firms adopt CI for risk management and decisions. By the end of the 20th century, the overall production value of CI industry had reached 70 billion U.S. dollars [23].

In order to identify potential risks, it is important for companies to collect and analyze information about their competitors' products and plans. Based on such information, a company can learn the relative weaknesses and strengths of its own products, and can then design new pointed products and campaigns to countervail those of its competitors. Traditionally, information about competitors has mainly come from press releases, such as analyst reports and trade journals, and recently also from competitors' websites and news sites. Unfortunately, such information is mostly generated by the company that produces the product. Therefore, the amount of available information is limited and its objectivity is questionable. The lack of

sufficient and reliable information sources about competitors greatly restricts the capability of CI.

With the emergence of Web 2.0, an increasing number of customers now have opportunities to directly express their opinions and sentiments regarding products through various channels, such as online shopping sites, blogs, social network sites, forums, and so forth. These opinion data, coming directly from customers, become a natural information source for CI. There are some existing studies on mining customer opinions [6,7,27,31,34]. However, these studies mainly focus on identifying customers' sentiment polarities toward products. The most important problem in CI—i.e., collecting and analyzing the competitors' information to identify potential risks as early as possible and plan appropriate strategies—has not been well studied.

Customer reviews are often a rich source of comparison opinions. Users usually prefer to compare several competitive products with similar functions, for example,

> *Nokia N95 has a stronger signal than iPhone.*
> *The iPhone has better looks, but a much higher price than the BB Curve.*
> *Compared with the v3, this V8 has a bigger body, and it has a much worse keyboard than Nokia E71.*

These comparison opinions are precious information sources for identifying the relative strengths and weaknesses of products, analyzing the enterprise risk and threats from competitors, and further designing new products and business strategies.

Mining such comparison opinions is a non-trivial task due to the large amount of customer reviews and their informal style. In this

* Corresponding author.
  E-mail address: kaiquan.xu@student.cityu.edu.hk (K. Xu).

paper, we propose a novel approach to extracting product comparative relations from customer reviews, and display the results as comparative relation maps for decision support in enterprise risk management.

The remainder of this paper is organized as follows: Section 2 reviews the related work in comparative opinion mining. Section 3 introduces our overall approach of comparative relation extraction. Section 4 introduces a novel graphical model we propose for comparative relation extraction. Section 5 presents our experiments that evaluate the proposed relation extraction approach. Section 6 concludes our study and discusses some future directions for research.

2. Related work

2.1. Sentiment analysis of user opinions

Much research exists on sentiment analysis of user opinion data [6,7,27,31,34], which mainly judges the polarities of user reviews. In these studies, sentiment analysis is often conducted at one of three levels: the document level, sentence level, or attribute level. Sentiment analysis at the document level classifies reviews into the types of polarities—positive, negative, or neutral—based on the overall sentiments in the reviews. A number of machine learning techniques have been adopted to classify the reviews [32]. Abbasi and Chen et al. propose the sentiment analysis methodologies for classification of Web forum opinions in multiple languages [1]. Sentiment analysis at the sentence level mainly focuses on identifying subjective sentences and judging their polarities. Most of these studies adopted the machine learning methods [42,47]. Sentiment analysis at both the document level and sentence level has been too coarse to determine precisely what users like or dislike. In order to address this problem, sentiment analysis at the attribute level is aimed at extracting opinions on products' specific attributes from reviews. In [17], Part Of Speech (POS) tag sequence rules were used to extract product attributes, and then the polarities of opinion phrases on the attributes were judged based on the context information. For the sentiment analysis, various features can be used. Term presence has been more effective than term frequency in classifying the polarities of documents, and the positions of terms also have had an important influence on sentiment analysis [32]. The POS tags of words, such as adjectives and adverbs, have been good indicators for the subjectivity detection and sentiment polarity classification [2,40]. The syntax features (such as dependency tree) outperform the bag-of-words for the sentiment polarity classification in some situations [24]. The interactions between topic and sentiment play an important role in sentiment analysis [13]. The techniques for sentiment analysis are mainly classified into two categories: unsupervised approaches and supervised approaches. The unsupervised approaches usually create a sentiment lexicon and determine the polarities by counting the positive and negative phrases [15,40]. The supervised approaches use the labeled data to train some classifiers (such as Naive Bayes, Maximum Entropy, or Support Vector Machine) to predict the unlabeled data [32,42,47]. Other important research on sentiment analysis includes identifying the sentiment target/topic and the opinion holder. The purpose of identifying the sentiment target/ topic is to discover subjectivity and sentiment sentences [43,45,46], and the existing research proposes various linguistic clues for this task. For the task of identifying the opinion holders associated with particular opinions [3,22], the semantic parsing techniques are proposed.

In addition, several systems [28,45,46] have been developed to automatically analyze customer reviews, mine opinions toward a product or attribute, and visually show the mined information for aiding users' decision making. Some systems [28] can aggregate opinions about certain attributes of several competing products and help users compare their pros and cons. However, these existing systems only mine general user opinions, which can be biased in acquiring competitors' information and identifying the potential operational risks. Unlike these systems, our study focuses on mining users' comparative opinions from product reviews. Because these comparative opinions can better reflect customers' preferences on competitive products, they should be more effective in tracing information regarding competitors and supporting enterprise risk management.

As noted, these studies mainly focused on judging customers' sentiment polarities toward products. However, few studies have focused on extracting the sentiment polarities of user opinions on comparisons of competitive products. Usually, this type of sentiment polarity is more important for enterprises to learn, so as to discover their products' weaknesses and design pointed products.

2.2. Relation extraction

Another closely linked research area is relation extraction, which detects if there exists a specific relation between entities, such as *work_in (Tom, IBM)* (meaning *Tom works in IBM*). A number of methods currently exist for relation extraction. Some of these methods are based on rules/templates, and others formulate the relation extraction as a classification problem and use various classification techniques to resolve it. In the rule-based methods, the extraction rules are defined manually [9] or learned from large annotated training corpora [14]. The classification-based methods can be divided into two categories: feature-based methods [21] and kernel-based methods [4,5]. Feature-based methods define the feature set (such as words, POS tags, entity types, path in parse tree, etc.), and represent the examples using these features to train classifiers. Usually, the computing complexity of this kind of method is relatively low, while the choice of features is intuitive and difficult. The kernel-based methods structurally represent examples and define kernels to compute the similarities in a high-dimension space implicitly. For example, in [8] and [48], the shallow parse tree kernels and the dependency tree kernel were used separately. In [26], a trace kernel was incorporated with the tree kernel to capture richer contextual information for biomedical relation extraction. This kind of methods does not need to define a feature set, but the computation complexity is relatively high.

Comparative relation extraction differs from the existing relation extraction in two important ways: 1) the comparative relation is a higher-order relation. That is, one comparative relation contains four entities/arguments, while the existing relation extraction methods mainly address relations with two entities. 2) Comparative relation extraction not only detects if the comparative relations occur or not, but also recognizes their directions. That is, the extraction task needs to indicate that product *A* is better than product *B* on a particular attribute, or the inverse. In contrast, the existing relation extraction methods mainly detect merely the occurrences of relations. These two characteristics bring particular challenges to comparative relation extraction: Involvement of multiple entities means that the long-range features need to be captured for better extraction; recognizing the direction of the relation makes the comparative relation extraction become a multi-class classification problem, unlike the existing relation extraction tasks, which were typically binary-class classification problems.

The only work up to now on extracting the comparative relations from customer opinions is [18]. Their work only identified the comparative sentences and extracted the relation items. In their research, the comparative relation included two product entities, one attribute entity, and keywords expressing the comparative relation. The comparative relations were classified