MIS Quarterly

SPECIAL ISSUE: BUSINESS INTELLIGENCE RESEARCH

WEB 2.0 ENVIRONMENTAL SCANNING AND ADAPTIVE DECISION SUPPORT FOR BUSINESS MERGERS AND ACQUISITIONS[1]

Raymond Y. K. Lau and Stephen S. Y. Liao
Department of Information Systems, City University of Hong Kong, Tat Chee Avenue, Kowloon,
Hong Kong SAR, CHINA {raylau@cityu.edu.hk} {issliao@cityu.edu.hk}

K. F. Wong
Department of Systems Engineering and Engineering Management, The Chinese University of Hong Kong,
Shatin, N.T., Hong Kong SAR, CHINA {kfwong@se.cuhk.edu.hk}

Dickson K. W. Chiu
Department of Information Systems, City University of Hong Kong, Tat Chee Avenue, Kowloon,
Hong Kong SAR, CHINA {kwchiu@cityu.edu.hk}

Globalization has triggered a rapid increase in cross-border mergers and acquisitions (M&As). However, research shows that only 17 percent of cross-border M&As create shareholder value. One of the main reasons for this poor track record is top management's lack of attention to nonfinancial aspects (e.g., sociocultural aspects) of M&As. With the rapid growth of Web 2.0 applications, online environmental scanning provides top executives with unprecedented opportunities to tap into collective web intelligence to develop better insights about the sociocultural and political–economic factors that cross-border M&As face. Grounded in Porter's five forces model, one major contribution of our research is the design of a novel due diligence scorecard model that leverages collective web intelligence to enhance M&A decision making. Another important contribution of our work is the design and development of an adaptive business intelligence (BI) 2.0 system underpinned by an evolutionary learning approach, domain-specific sentiment analysis, and business relation mining to operationalize the aforementioned scorecard model for adaptive M&A decision support. With Chinese companies' cross-border M&As as the business context, our experimental results confirm that the proposed adaptive BI 2.0 system can significantly aid decision makers under different M&A scenarios. The managerial implication of our findings is that firms can apply the proposed BI 2.0 technology to enhance their strategic decision making, particularly when making cross-border investments in targeted markets for which private information may not be readily available.

Keywords: Domain-specific sentiment analysis, business relation mining, statistical learning, evolutionary learning, business intelligence, Web 2.0, mergers and acquisitions

[1]Hsinchun Chen, Roger Chiang, and Veda Storey were the accepting senior editors for this paper. Atish Sinha served as the associate editor.

The appendix for this paper is located in the "Online Supplements" section of the *MIS Quarterly*'s website (http://www.misq.org).

Introduction

Mergers and acquisitions (M&As) are a common strategy for firms to maintain sustainable growth and achieve a competitive advantage (Porter 1985). Globalization has led to a rapid increase in cross-border M&A initiatives, which involve acquirers and acquiring targets operating in different countries (Shimizu et al. 2004). However, research shows that more than 60 percent of M&As fail to meet their anticipated financial objectives (Andrade et al. 2001; Calipha et al. 2010; Cartwright and Cooper 1993; Marks and Mirvis 2001; Tetenbaum 1999), and only 17 percent of cross-border M&As can create shareholder value (Shimizu et al. 2004). It is believed that top management's over-attention to the financial factors and neglect of sociocultural and other nonfinancial aspects of M&As is one of the main reasons for the large number of M&A failures (Calipha et al. 2010; Galpin and Herndon 2007; Shimizu et al. 2004; Stahl and Voigt 2008; Weber et al. 1996).

El Sawy (1985) suggested that top executives' strategic decisions (e.g., M&As) would be enhanced by environmental scanning. Environmental scanning is the acquisition and utilization of information about events (e.g., sociocultural issues of M&As), trends, and relationships extracted from an organization's external environment to support top management's strategic planning and decision making (Albright 2004; Choo 1999; McEwen 2008; Yasai-Ardekani and Nystrom 1996). Continuous environmental scanning is particularly important to support top executives' adaptive decision making in turbulent business environments (Choo 1999; El Sawy 1985; McEwen 2008). With the rapid growth and the increased volume of different types of information on the Internet, online environmental scanning is becoming more popular (Choo 1999). However, it is not practical for top executives to manually perform environmental scanning over the Internet due to the problem of information overload (Farhoomand and Drury 2002; Lau, Bruza et al. 2008; Yan et al. 2011). Accordingly, online environmental scanning enhanced by business intelligence (BI) is desirable. In particular, it enables top executives and M&A consultants to identify the sociocultural and political–economic issues to help them address post-acquisition integration problems, thereby improving the success rates of M&As.

The rise of Web 2.0 applications (Oreilly 2007; Raman 2009) has led to a generation of rich environmental signals that have triggered the development of emerging BI technologies, collectively called BI 2.0 (Chen 2010). These technologies support automated sentiment and affect analysis (Abbasi, Chen, and Salem 2008; Abbasi et al. 2008; Archak et al. 2011; Bollen et al. 2011; Das 2010; Das and Chen 2007), the real-time visualization of unstructured information (Abbasi and Chen 2008; Bhagwan et al. 2009; Chung et al. 2005; Ong et al. 2005), and the discovery of hidden business relationships from financial texts (Bao et al. 2008; Ma, Sheng, and Pant 2009) to support real-time managerial decision making. In fact, Web 2.0 information has been shown to be valuable for supporting various financial investment tasks (Antweiler and Frank 2004; Bollen et al. 2011; Tetlock et al. 2008). This paper illustrates the design and development of a BI 2.0 system called Adaptive Business Intelligence for Mergers and Acquisitions (ABIMA) that enhances top executives' M&A decision making. In particular, we employ Chinese companies' cross-border M&As as an exemplar to demonstrate the functionality and effectiveness of the proposed design and its instantiation.

The Chinese government's "going outside" strategy has driven many cross-border M&As since 2002 (Tan and Ai 2010; He and Lyles 2008). Apart from the aim of improving China's global political presence, the desire to secure natural and strategic resources that sustain economic growth is another key driver behind the cross-border M&As of Chinese companies.[2] However, Chinese companies' cross-border M&As in some foreign countries, such as the United States, have encountered serious challenges. These have mainly been the result of Chinese companies' lack of knowledge about the sociocultural and political–economic characteristics of the targeted M&A environments (Tan and Ai 2010; He and Lyles 2008). One well-known example is CNOOC's attempt to acquire Unocal, which was stopped by the U.S. Congress in 2005 (Murphy 2005). Another example is the Aluminum Corporation of China's attempt to acquire an 18 percent stake in the Australian–British mining giant, the Rio Tinto group, which was rejected by Rio Tinto's shareholders in 2009 (Barboza and Wines 2009). If Aluminum Corporation had been able to apply a BI 2.0 system to identify and analyze the sociocultural comments posted online by Rio Tinto's shareholders early on, the company might have been able to employ appropriate public relations or a conflict resolution strategy (e.g., acquisition via a third party) to rescue its failing M&A attempt.

Grounded in Porter's (1980) five forces model, one major contribution of our research is the design of a novel due diligence scorecard model that leverages collective web intelligence to enhance M&A decision making. In addition, guided by the design science research methodology (Hevner et al. 2004; March and Storey 2008; Peffers et al. 2008), another

[2]"China Spreads Its Wings—Chinese Companies Go Global" (http://www.accenture.com/us-en/Pages/insight-china-spreads-wings-chinese-companies-go-global.aspx).

important contribution of our research is the design and development of an adaptive BI 2.0 system (i.e., an instantiation) that operationalizes the aforementioned scorecard model for adaptive M&A decision support. In particular, a computational algorithm (i.e., a method) called pseudo labeling (PL) is designed to support the automatic construction of domain-specific sentiment lexicons that facilitate the sentiment analysis of the events (e.g., the sociocultural issues) and trends related to M&As. A hybrid natural language processing (NLP) (Arazy and Woo 2007) and statistical learning-based algorithm for business relation mining is designed to support relationship scanning in M&A environments. Moreover, a hierarchical coevolution genetic algorithm (HCGA) has been developed to empower the BI system with learning and adaptation capabilities.

The managerial implication of our research is that top executives can apply the proposed BI 2.0 technology to effectively and efficiently scan the sociocultural and political–economic events from the Web 2.0 environment to enhance strategic decision making. Specifically, the proposed technology offers complementary decision support for cross-border corporate investment for which private information of the targeted markets may not be readily available. The societal implication is that the proposed BI technology may promote fair financial trading because enterprises that lack private market information may leverage collective web intelligence to improve their financial decision-making processes. The main research questions addressed by our study are summarized as follows:

- Is the proposed computational method for domain-specific sentiment analysis more effective than other existing methods?
- Is the proposed computational method for business relation mining more effective than other existing methods?
- Can the ABIMA system provide significant decision support to people conducting M&A activities?
- Can the ABIMA system offer adaptive decision support under different M&A scenarios?

The rest of this paper is organized as follows. The next section discusses research related to BI 2.0 applications, sentiment and affect analysis, and business relation mining. Following that, the system architecture and the implementation of the proposed ABIMA system are highlighted. The computational methods for domain-specific sentiment lexicon construction, business relation mining, M&A target scoring, and learning and adaptation are then illustrated. Next, an evaluation of the proposed ABIMA system in the context of Chinese companies' cross-border M&As is reported. Finally, we offer concluding remarks and suggest directions for future research.

Related Research

BI 2.0 Applications

Srivastava and Cooley (2003) developed a Web business intelligence (WBI) system to fetch information from diverse sources on the Web and present relevant information to users via the most suitable forms according to specific user profiles. Based on functional linguistic theory, Abbasi and Chen (2008) developed a BI 2.0 framework and an instantiation called CyberGate for the interactive analysis and visualization of unstructured textual data generated from computer-mediated communications (CMC). Chung et al. (2005) proposed the business intelligence explorer (BIE) framework, which incorporates Web mining methods, data clustering techniques, and visualization methods to effectively extract knowledge from the Web. Ong et al. (2005) designed a method based on self-organizing map (SOM) hierarchical knowledge generation to extract and visualize business intelligence based on Chinese news articles. The sound index system explored user-generated Web content, online communities, and social networks to dynamically compose music charts (Bhagwan et al. 2009). In contrast, our proposed ABIMA system utilizes an unsupervised rather than a supervised learning approach to extract both sentiments and business relations from financial news articles and investors' comments posted on the Web.

Sentiment and Affect Analysis

Sentiment and affect analysis facilitates the extraction of collective wisdom from opinionated expressions by scanning the Internet in general and the Web 2.0 information environment in particular. Pang et al. (2002) applied supervised machine learning techniques, such as SVM, to predict the sentiment polarity of movie reviews. Turney and Littman (2003) developed an inference-based opinion mining method called semantic orientation (SO) analysis to estimate the polarity of sentiments. The SO of an arbitrary word can be estimated based on the strength of its association with 14 seeding sentiment indicators measured in terms of point-wise mutual information (PMI). Context-sensitive sentiment analysis was conducted by first parsing and constructing a dependency tree from a sentence, at which point a set of linguistic features was used to train the supervised AdaBoost classifier to predict the sentiment polarity of a word (Wilson et al. 2006).

Based on the polarities of sentiments extracted from online message boards, a supervised machine learning approach was applied to predict the movement of the Dow Jones stock index (Das 2010; Das and Chen 2007). Abbasi, Chen, and Salem (2008) developed an entropy-weighted genetic algorithm (EWGA) to select the best syntactic and stylistic features for multilingual sentiment classification against various extremist online forums. Archak et al. (2007, 2011) applied sentiment analysis to the fine-grained product feature level, incorporating the polarities of sentiments at this level into a consumer preference model to predict product sales. Our proposed domain-specific sentiment analysis method also conducts fine-grained sentiment polarity classification at the aspect level (e.g., the sociocultural aspects). In addition, we propose and evaluate an unsupervised learning method for building a domain-specific sentiment lexicon to analyze financial texts.

Abbasi et al. (2008) proposed an SVM regression correlation ensemble (SVRCE) method that utilizes an ensemble of classifiers to predict affect intensities of online texts. Calix et al. (2010) found that a PMI-based automatic emotion word feature selection method was as effective as a manual emotion feature selection approach. Based on the theory of kinesics, the CAO (emotiCon Analysis and decOding) system (Ptaszynski et al. 2010) classified a token into different semantic areas, such as "mouth" or "eyes," according to its frequency of appearance in a training corpus. Calvo and D'Mello (2010) conducted a comprehensive meta-analysis of existing approaches for the design and development of affect detection systems. More recently, supervised affect analysis has been successfully applied to predict the movement of the Dow Jones Industrial Average (Bollen et al. 2011). A lexicon-based affect feature extraction approach is applied to the proposed ABIMA system because labeled training examples are difficult to acquire for the finance domain.

Business Relation Mining

Bernstein et al. (2003) proposed a computational method to predict the associations among companies based on the co-occurrence statistics of the corresponding stock tickers found in financial news articles. The CoMiner system makes use of several predefined syntactic patterns (e.g., company A versus company B) to identify competitive companies based on a Web corpus (Bao et al. 2008). Each syntactic pattern was assigned a weight and the PMI measure was used to estimate the strength of competitiveness between two companies over the Web corpus. Ma, Pant, and Sheng (2009) developed a link-based weighted directed graph approach to extract company competitor relations from online financial news articles. For instance, if two companies co-occurred in a financial news article, a directed link between these companies would be established.

Ma, Sheng, and Pant (2009) also applied a link-based approach to predict company revenue relations based on online financial news, and found that both the decision tree classifier and the logistic regression classifier achieved comparable performance. Further, Pant and Sheng (2009) used both the hyperlinks and the content of home pages for a pair of companies to predict whether they were competitors. Given that labeled training examples are not readily available, our proposed computational method for business relation mining utilizes semi-automatically constructed relationship lexicons and generic NLP rules to identify competitor, collaborator, and supplier relationships among companies.

System Architecture and Implementation

A formal environmental scanning process consists of five subtasks: the identification of scanning needs, information gathering, information analysis, results communication, and informed decision making. Accordingly, the design of the ABIMA system architecture is driven by the formal environmental scanning model that supports the aforementioned subtasks. Since an effective environmental scanning system should take the scanning context into account (Yasai-Ardekani and Nystrom 1996), the design of the ABIMA system considers the pragmatic aspects of M&As (i.e., the scanning context) as well. M&As typically go through several stages, such as strategic planning, target selection (i.e., preliminary due diligence), legal preparation, target valuation, financing, structuring transaction, due diligence inquiry, negotiation and filing the letter of intent, closing, and post-merger integration (Emott 2011; Galpin and Herndon 2007; Reed et al. 2007; Wu and Xie 2010). The proposed ABIMA system supports top management's decision making at the stages of strategic planning, target selection, target valuation, due diligence inquiry, and post-merger integration. The focus of this paper is on the target selection stage.

Given that sociocultural elements have been identified as critical success factors for M&As (Birkinshaw et al. 2000; Harding and Rouse 2007; Shimizu et al. 2004), the ABIMA system supports automated qualitative analyses (e.g., sociocultural analysis, business network analysis, internal analysis,

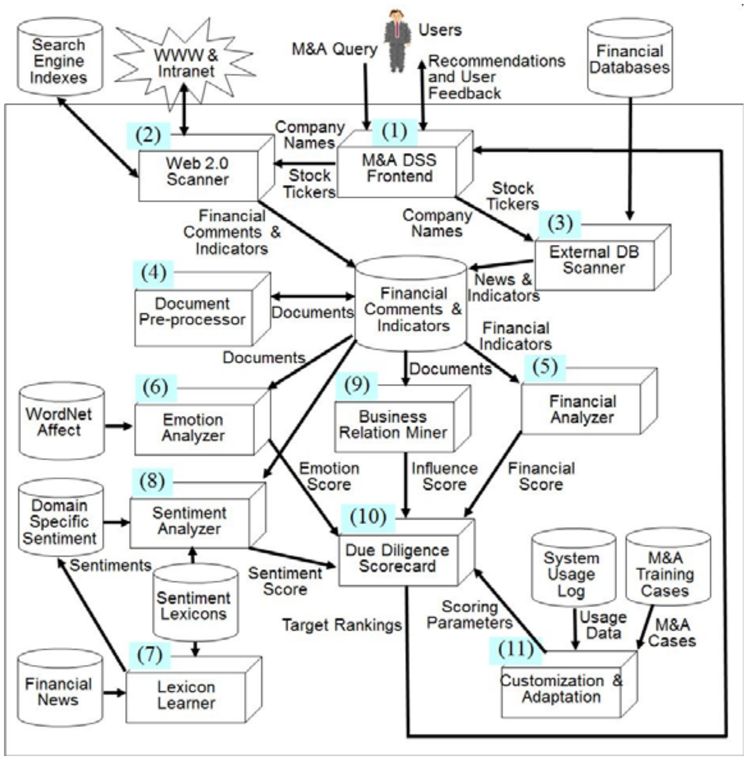

Figure 1. The General System Architecture of ABIMA

etc.) to augment traditional quantitative M&A analysis. By continuously scanning and analyzing the sociocultural and political–economic information retrieved from the Web 2.0 environment, the ABIMA system is able to leverage collective web intelligence to enhance top executives' M&A decision-making processes. Figure 1 presents the general system architecture of ABIMA. Directed arrow lines represent information flow among the system components. Module (1) supports both the identification of scanning needs and the resulting communication subtasks of the environmental scanning process, whereas modules (2) and (3) support the information-gathering subtask. In addition, modules (4), (5), (6), (7), (8), and (9) perform the information analysis subtask. Module (10) facilitates the informed decision-making subtask. Finally, module (11) enables an adaptive environmental scanning process. The ABIMA system was developed using Java (J2SE v 5.0), Java Server Pages (JSP) 2.1, and Servlet 2.5. The prototype system is hosted on a DELL 1950 III Server with Quad-Core Xeon 2.33GHz Processors, 16GB main memory, and 6TB secondary storage. The server operates under Windows Server 2003 R2 (x64 Edition).

The following paragraphs illustrate the various components of the ABIMA system:

1. A user (e.g., a top executive) initiates a query by entering the information of an M&A scenario through the decision support front end. In particular, s/he enters the name (or stock ticker) of the acquirer, and other optional information, such as the targeted M&A industrial sector, the range of market values of the targets, the targeted countries, the preference for targets providing complementary products or services, the time window of due diligence, and so forth into the ABIMA system. For commercial due diligence, the time window of data analysis usually ranges from the past 3 to 5 years (Emott 2011, p. 51; Howson 2006). If the user does not specify the time window for due diligence, a system default of three years will be applied (i.e., analyzing the historical data of the targets for the past three years). A screen-shot of the cross-border M&A recommendations for the SAIC Motor case is shown in Figure 2. For example, the user can click the "drill-down" button beside a recommended target, such as "Isuzu Motors" to display the quantitative (e.g., EBITDAR) and qualitative details (e.g., the results of sentiment analysis for the company and its products) about the company.

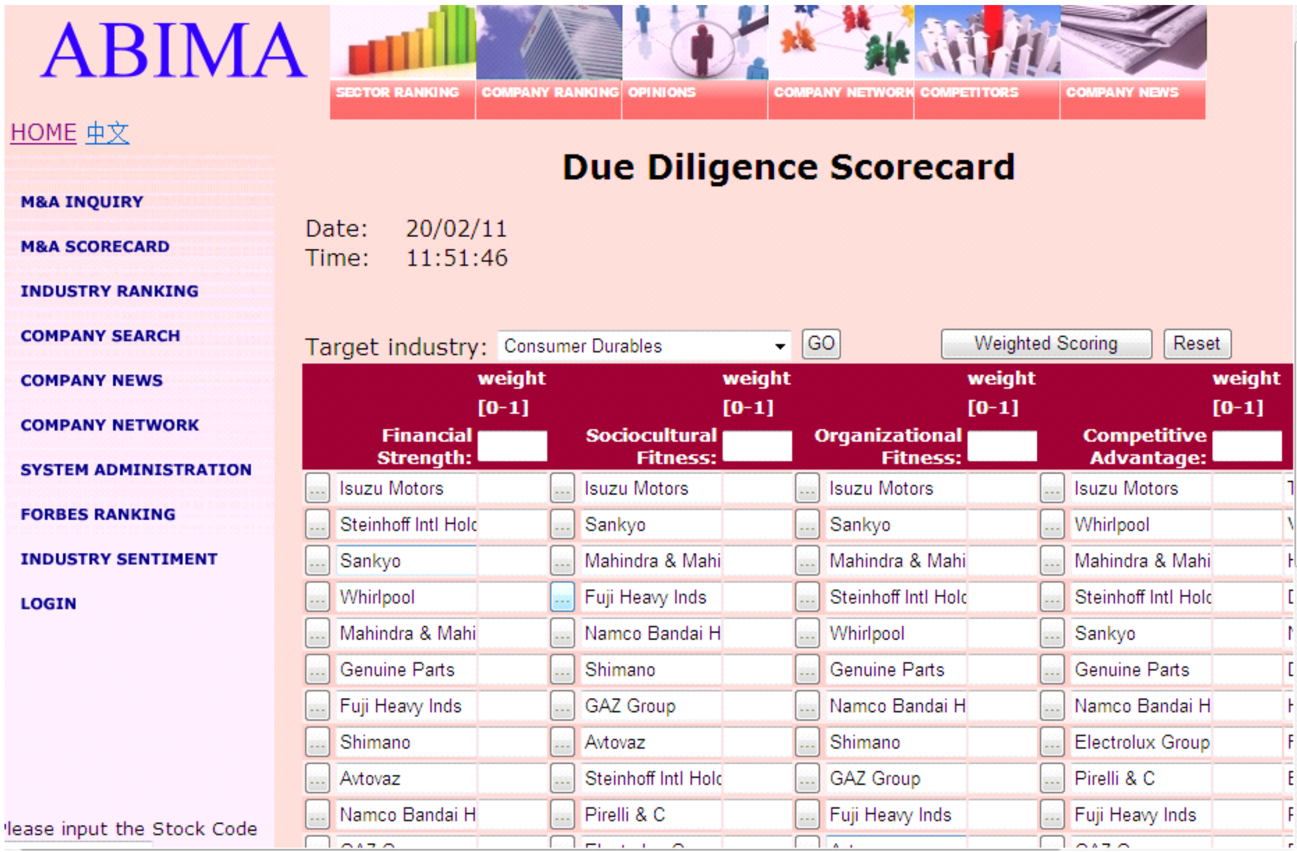

Figure 2. The Due Diligence Scorecard of ABIMA (February 2011)

2. Public information, such as online financial news, RSS feeds, investors' comments, company Web sites, company filings archived at stock exchange sites (e.g., EDGAR), and more are retrieved through the Web 2.0 scanner and stored in ABIMA's local database. The Web 2.0 scanner utilizes Web services, application programming interfaces (APIs),[3] RSS feed readers,[4] Internet search engines, and dedicated crawler programs to retrieve all of the relevant information pertaining to the acquirer and the potential targets of an M&A query. Our crawler programs periodically (e.g., daily) extract up-to-date financial information from a list of authoritative Web sites (e.g., Yahoo Finance, Google Finance, Forbes, Reuters, Hoovers, Fortune, etc.). In addition, if the M&A query is initiated by the acquirer, additional information can be retrieved from its intranet as well.

3. Proprietary financial information services can also be accessed through the external database scanner. This scanner utilizes the dedicated connection interfaces provided by service providers, such as Reuters and Bloomberg, to access proprietary information.

4. The document preprocessor applies traditional document preprocessing procedures (Salton and McGill 1983), such as stop word removal and part-of-speech (POS) tagging, to process the financial news articles or investor comments prepared in English. The ABIMA system employs the POS tagger provided by the GATE information engineering service (Cunningham 2002) to parse and tokenize financial news articles or comments written in English. The preprocessed financial information is stored in the system's local database. For financial news

[3]http://code.google.com/intl/en/apis/finance/.

[4]www.google.com/reader.

articles or investor comments written in Chinese, word segmentation (Foo and Li 2004) must first be conducted. A GNU-licensed word segmentation program developed by the Chinese Academy of Sciences, ICTCLAS,[5] is applied to extract meaningful word combinations and tag these combinations with appropriate POS tags.

5. The financial analyzer applies various financial appraisal methods and standard metrics, such as return on equity (ROE), return on asset (ROA), cash flow to debt ratio, net asset ratio, earnings before interest, taxes, depreciation, amortization, and rent (EBITDAR), debt to equity ratio, and others to quantitatively evaluate the potential M&A targets (Emott 2011; Madura et al. 1991; Reed et al. 2007). Because this paper focuses on applying innovative BI technologies to enhance M&A decision making, the details of financial analysis for M&As will not be discussed.

6. The emotion analyzer then extracts the affects embedded in texts (e.g., investors' comments about a potential M&A deal). The WordNet-affect lexicon (Valitutti et al. 2004), which is an extension of the standard WordNet lexicon (Miller et al. 1990), is used to assess the degree of affect implied in stakeholders' or investors' comments with reference to the classes of affect—anger, fear, happiness, and sadness (Calix et al. 2010). The computational details of the proposed affect analysis method are given in Appendix B.

7. The lexicon learner uses a novel unsupervised learning algorithm to build a domain-specific (e.g., finance) sentiment lexicon based on the general sentiment lexicons, such as OpinionFinder (Riloff et al. 2005). The proposed sentiment construction algorithm utilizes a statistical learning method that has been successfully applied to information retrieval (Lau, Bruza et al. 2008) and ontology learning (Lau, Song et al. 2009) in previous studies. Our design choice is to use unsupervised learning methods rather than supervised machine learning methods because it is difficult to acquire a large number of labeled training examples to train a supervised classifier in the finance domain.

8. The sentiment analyzer then applies the automatically learned domain-specific sentiment lexicon to analyze the sentiments presented in financial texts (e.g., news, conversations captured from financial message boards, blog posts, etc.). In particular, an aspect-oriented sentiment analysis method (Brody and Elhadad 2010; Lau, Lai et al. 2009) is applied to identify the sentiments related to sociocultural and political–economic aspects of M&As.

9. The business relation miner applies a hybrid NLP (Arazy and Woo 2007) and unsupervised statistical inference method to extract competitor, collaborator, and supplier relationships among the companies of a specific industrial sector based on financial news articles. Several network-based metrics underpinned by Porter's (1980) five forces model (Porter 1980) are developed to estimate the competitive advantage of a firm or an entire industrial sector.

10. The due diligence scorecard assesses a potential industrial sector or targeted firms by considering the most important M&A factors, such as financial merits, business fitness, sociocultural fitness, and competitiveness (Calipha et al. 2010; Garbuio et al. 2010; Harding and Rouse 2007; Porter 1980; Shimizu et al. 2004; Weber et al. 1996). The results of the multi-dimensional M&A analysis are presented to the user via the decision support front end. The idea of our due diligence scorecard is similar to that of the balanced scorecard (Kaplan and Norton 1996), although the assessment dimensions are slightly different. As Figure 2 shows, our design allows users to assign weights to various dimensions according to their specific preferences and the specific M&A scenarios. The ABIMA system can compute aggregated scores for the potential M&A targets and generate a single ranked list of firms if the user prefers such a presentation format.

11. The customization and adaptation module conducts learning and adaptation tasks at both the system level and the individual user level. First, it performs a system-wide adaptation to refine the M&A decision support mechanism by taking into account changing environmental signals such as training M&A cases and the corresponding financial documents in different periods. Specifically, the hierarchical coevolutionary approach (Delgado et al. 2004; Huang et al. 2009) is adopted to design the system-wide adaptation mechanism. Second, the customization and adaptation module observes a user's interactions with the system over time, and utilizes a genetic algorithm (GA) (Goldberg 1989; Holland 1992) to learn and infer the user's specific information scanning and results communication preferences to realize a personalized environmental scanning process (El Sawy 1985; Yasai-Ardekani and Nystrom 1996). Due to space constraints, this paper only focuses on the system-wide learning and adaptation mechanism.

[5] http://www.nlp.org.cn/.

The Computational Methods of ABIMA

Business Relation Mining

According to the theory of resource-based view (RBV) (Barney 1991), researchers have explained the motive of many M&As in terms of an acquirer's intention to secure complementary resources from an acquired company. The unique contribution of business relation mining in the context of M&As is that the potential of *synergy* between two companies can be estimated based on a supply chain network discovered via business relation mining. For instance, if two companies appear in a supply chain, complementary resources between them may exist, which creates the potential for an acquisition. The structural embeddedness model posits that a company's embeddedness in a business network impacts its competitive performance (Gnyawali and Madhavan 2001). Bernstein et al. (2003) also indicate that a firm's relationships with other firms could significantly influence its status in a specific industrial sector. Therefore, business relation mining contributes to the effective assessment of the competitiveness of a potential M&A target through the extraction and visualization of its business network.

The proposed computational methods for company name identification and company relation extraction are conducted at the fine-grained sentence level to improve overall accuracy. Identifying company names and their relationships from financial documents is quite challenging because natural languages are ambiguous and very flexible. For company name identification, there are at least two main challenges. First, a variety of names can be used to refer to the same company. Second, co-references (e.g., pronouns) may be used to refer to a company. For instance, the company General Electric can also be referred to as GE (an abbreviation) or co-referenced by "it" or "the company" in financial documents. To deal with the first problem, the ABIMA system utilizes an unsupervised learning method to automatically extract various abbreviations that are statistically associated with a company's name. For the second problem, the ABIMA system employs an NLP rule-based approach to resolve the co-references appearing in texts. The advantage of the proposed method is that manually labeled training examples are not required to resolve the co-references to a company.

The full names and stock tickers of companies stored in our local database are passed to the named entity recognition (NER) module of GATE (Maynard et al. 2001) to compose the company name dictionary for basic name detection. GATE is a general information engineering service, which also contains a text tokenizer, a sentence splitter, a part-of-speech tagger, a morphological analyzer, and a VP chunker (Cunningham 2002). The NER rules of GATE are extended to consider the specific requirements of company name identification. For instance, the tokens preceding "Inc," "Co," "Ltd," and so forth are likely company names. For business relation mining, the ABIMA system first utilizes GATE to identify valid lexical patterns, such as (organization, relation, organization) and (relation, organization, organization). A seeding relationship lexicon is built using the common relationship keywords often used to describe competitors, collaborators, and suppliers. This is the only human intervention involved when applying this method across business domains.

For instance, words such as *cooperate*, *ally*, *collaborate*, *joint*, *together*, *partner*, and their various forms are used to build the collaborative relationships set of the relationship lexicon. In contrast, words such as *compete*, *challenge*, *versus*, *enemy*,= *against*, *vie*, and their various forms are used to build the competitive set. Supplier relationship mining utilizes keywords, such as *supply*, *purchase*, *procure*, *lend*, and so on to build the supplier set. Based on the proposed relationship lexicon, each valid lexical tuple returned by GATE is classified as a competitive, collaborative, or supply chain-related business relation by comparing the relationship keywords appearing in both the lexicon and the lexical tuple. However, such an approach is expected to have a low recall because some relationship indicators in the financial domain may not be captured by the seeding relationship lexicon. To make the proposed computational method more robust and easily applicable to other business domains, a statistical learning method is applied to extract additional relationship indicators, which are statistically correlated to the seeding relationship keywords based on a domain-specific corpus (e.g., financial documents).

The skeleton of a novel computational algorithm called business relation mining by expansion (BRME) is shown in Figure 3. The first step of the BRME algorithm is to expand the full names of companies (e.g., the company full names extracted from Yahoo Finance) by using a variant of PMI called the balanced mutual information (BMI) method, which has been successfully applied to find statistically correlated terms for concept discovery (Lau 2003; Lau, Song et al. 2009). Compared to the standard PMI often used to infer the statistical associations among terms (Turney and Littman 2003), the distinct advantage of the BMI method is that it can take into account both positive and negative correlation information to infer the strength of an association .

Algorithm

RelationMiningbyExpansion(T^{ORG}, L^R, D, ϖ_{Abbr}, $\varpi_{Keyword}$, ϖ_{BMI}, ϖ^{Abbr}_{window}, $\varpi^{Keyword}_{window}$, ϖ^{CoR}_{window}, ϖ^{Rel}_{window}, ϖ_{frq})

Inputs:
- L^R /* a starting relationship lexicon with predefined keywords
- D /* a collection of documents of a specific domain
- T^{ORG} /* a table of company full names
- ϖ_{Abbr} /* a dynamic pruning threshold for abbreviations extraction
- $\varpi_{Keyword}$ /* a dynamic pruning threshold for relational keyword extraction
- ϖ_{BMI} /* a parameter for BMI value computation
- ϖ^{Abbr}_{window} /* the size of a virtual text window for abbreviation mining
- $\varpi^{Keyword}_{window}$ /* the size of a virtual text window for relational keyword mining
- ϖ^{CoR}_{window} /* the size of a virtual text window for co-reference extraction
- ϖ^{Rel}_{window} /* the size of a virtual text window for company relation mining
- ϖ_{frq} /* a frequency threshold for pruning mined relationships

Output: T^{REL} /* tuples of recognized organization relations

Main Procedure:

1. expand the organization name table T^{ORG} with abbreviations:
   1.1 for each document $d \in D$, utilize the T^{ORG} table and invoke GATE's NER module to find organizations names and tokens co-located within a window of size ϖ^{Abbr}_{window};
   1.2 apply the BMI formula and parameter ϖ_{BMI} to compute an organization's association scores with all candidate abbreviations;
   1.3 for each candidate abbreviation a, perform pruning using abbreviation entropy (AE);
   1.4 if $AE(a) > \varpi_{Abbr}$, discard the candidate abbreviation;
2. expand the relationship lexicon L^R:
   2.1 for each document $d \in D$ and for each seeding relational keywords $k \in L^R$;
      2.1.1 find candidate relational keywords with the same part-of-speech as k and co-located within a virtual text window of size $\varpi^{Keyword}_{window}$;
      2.1.2 apply the BMI measure to compute the association scores of all the co-located candidate relational keywords with k;
   2.2 for each candidate relational keyword k', perform pruning using keyword entropy (KE);
   2.3 if $KE(k') > \varpi_{Keyword}$, discard the candidate relational keyword because of its high entropy;
3. resolve the co-references for each document $d \in D$ using distance threshold ϖ^{CoR}_{window};
4. extract company relationships in D through:
   4.1 filter out documents without 2 or more company names;
   4.2 for each document $d \in D$, identify all organizations and relational keywords with close proximity defined by ϖ^{Rel}_{window}, add the relation to T^{REL};
   4.3 prune relationships by type: for each pair of organizations, retain only one type of relationship with the highest frequency;
   4.4 prune relationships by frequency: delete all relationships with frequency < ϖ_{frq};
5. return tuples of recognized organization relations T^{REL};

Figure 3. The Business Relation Mining by Expansion (BRME) Algorithm

Because statistical learning may introduce noisy abbreviations, the second step of the BRME algorithm is to utilize an entropy-based method called abbreviation entropy to prune invalid abbreviations. Abbreviation entropy is defined by

$$AE(a) = \frac{\sum_{org \in T^{ORG}} Ass(a, org)}{|T^{ORG}|}$$

where *AE(a)* is the abbreviation entropy of a candidate abbreviation *a* of a company. *Ass(a, org)* returns 1 if *a* is associated with a company name defined in the table T^{ORG}. $|T^{ORG}|$ is the cardinality of the name table. Hence, if a candidate abbreviation *a* is associated with many company names at the same time, its entropy is high and it should be pruned. Similarly, the seeding lexicon of relationship keywords is expanded by using the same statistical learning and pruning method.

The third step of the BRME algorithm is to resolve co-references by using the proposed backward searching method. According to the expanded company name table, all the company names within a document are first identified. By invoking GATE's part-of-speech recognizer, all of the pronouns within the document are also identified. Starting with the first ambiguous pronoun in a document, the proposed co-reference resolution method finds the nearest company name by moving a virtual pointer backward. If the distance between the pronoun and the nearest company name does not exceed the co-reference distance threshold ϖ^{CoR}_{window}, the co-reference is resolved by using the nearest company name to replace the pronoun. This procedure continues to resolve the ambiguous pronouns that follow in the document.

After co-reference resolution, the BRME algorithm identifies pairs of company names by using the relationship keywords defined in the expanded relationship lexicon. If the three elements of a valid relation pattern are colocated in a close proximity defined by the parameter , a candidate business relation is extracted. This process is repeated for the entire corpus. Finally, the candidate business relations are pruned using relation type and frequency. For example, if both competitive and collaborative relations of a pair of companies are found in the corpus, th ϖ^{Rel}_{window} e BRME algorithm retains the type of relation based on which one has a higher occurrence frequency.

The BMI measure is applied to find candidate company abbreviations or candidate relationship keywords from a corpus, expanding the company name table or the relationship lexicon. The *BMI* measure is defined by the following:

$$BMI(t_i, t_j) = \varpi_{BMI} \times \left(\Pr(t_i, t_j) \log_2 \left(\frac{\Pr(t_i, t_j) + 1}{\Pr(t_i)\Pr(t_j)} \right) + \Pr(\neg t_i, \neg t_j) \log_2 \left(\frac{\Pr(\neg t_i, \neg t_j) + 1}{\Pr(\neg t_i)\Pr(\neg t_j)} \right) \right) -$$

$$(1 - \varpi_{BMI}) \times \left(\Pr(t_i, \neg t_j) kig_2 \left(\frac{\Pr(t_i, \neg t_j) + 1}{\Pr(t_i)\Pr(\neg t_j)} \right) + \Pr(\neg t_i, t_j) \log_2 \left(\frac{\Pr(\neg t_i, t_j) + 1}{\Pr(\neg t_i)\Pr(t_j)} \right) \right)$$

where $BMI(t_i, t_j)$ is a function to estimate the statistical correlation between two terms t_i and t_j. For example, one of the terms is the recognized full company name. The parameter $\varpi_{BMI} \in [0,1]$ is used to adjust the relative weight of positive and negative evidence, respectively (Lau 2003; Lau, Song et al. 2009). $\Pr(t_i, t_j)$ is the joint probability that both terms appear in a text window, and $\Pr(t_i)$ is the probability that a term t_i appears in a text window. The probability $\Pr(t_i)$ is estimated based on $|\frac{w_t}{w}|$ where $|w_i|$ is the number of virtual text windows containing the term *t* and $|w|$ is the total number of windows constructed from a training corpus. Similarly, $\Pr(t_i, t_j)$ is the fraction of the number of windows containing both terms out of the total number of text windows. Negation such as $\Pr(t_i, \neg t_j)$ is interpreted in the way that t_i but not t_j appears in a text window. After computing the *BMI* scores, the tokens with the top scores and high string similarity (Islam and Inkpen 2008) are selected as the candidate abbreviations for company names.

Each pair of companies with a recognized relation is used to build a business network for a particular industrial sector. For both competitive and collaborative relations, the directions of the relations are not extracted because it is too challenging to predict these based on shallow text parsing and light-weight NLP techniques. The Pajek Java library (Batagelj and Mrvar 1998) has been applied to develop our business network visualization module. The business network diagram for the banking industry in the Forbes 2000 list is depicted in Figure 4. In the diagram, dashed lines represent competitive relations, solid lines represent collaborative relations, and directed solid lines indicate supplier relations pointing from suppliers to consumers. It is easy to observe that some nodes (i.e., companies) are well connected to many other nodes. One example is Citigroup, one of the hubs in the business network. The bank's structural embeddedness in the banking network (e.g., connecting many nodes via the collaborative links) has implications for its competitiveness. Indeed, Citigroup is one of the leading commercial banks in the world.

Automated Construction of Domain-Specific Sentiment Lexicons

The identification of sociocultural and political–economic issues is essential at both the target-selection and due diligence stages because these issues can cause serious problems

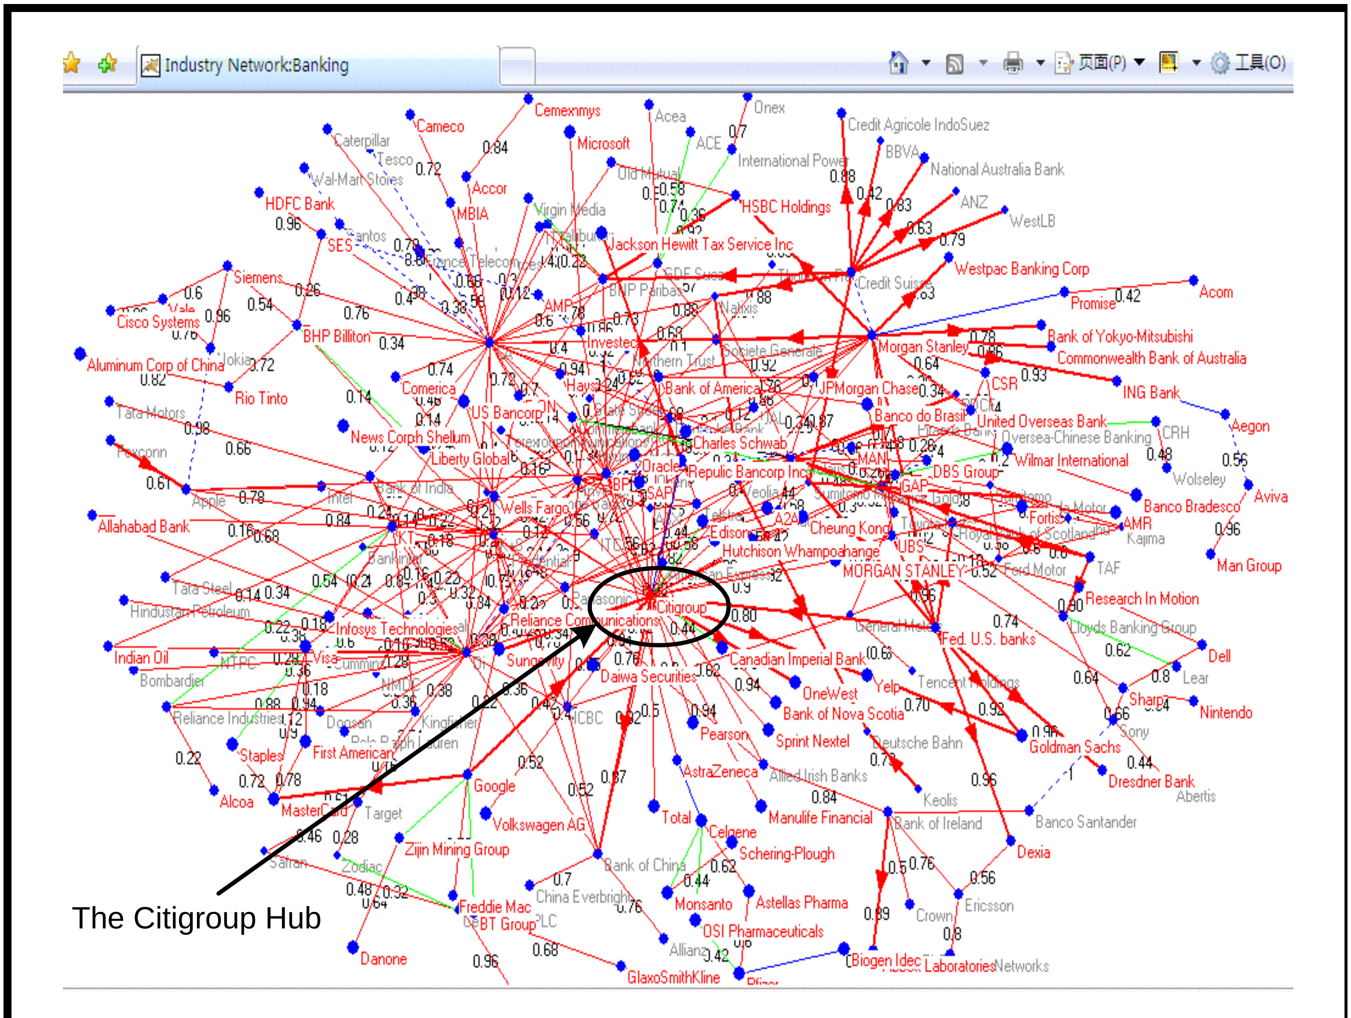

Figure 4. The Business Relation Network of the Forbes Banking Industry

at the post-acquisition integration stage, leading to the total failure of an acquisition (Birkinshaw et al. 2000; Calipha et al. 2010; Deresky 2011; Harding and Rouse 2007; Shimizu et al. 2004). At the same time, it has been indicated that analyzing consumers' attitudes toward a targeted firm is important when the social scanning of the firm is conducted (Albright 2004). Apart from external environmental scanning, an internal analysis to understand the strengths and weaknesses of both the acquirer and the target is also important in M&As (Deresky 2011; Garbuio et al. 2010; Reed et al. 2007). As Figure 5 shows, sentiment analysis helps top executives quickly identify potential sociocultural issues early on, enabling them to make a more informed decision as early as at the target-selection stage. The <neg-sent> tag indicates a negative sentiment, and the <aspect> tag refers to a sociocultural or political–economic aspect in Figure 5. Even if top management decides to acquire a target involving some sociocultural concerns, they can be better prepared to handle post-acquisition integration problems if they are alerted about the issues in the beginning stages. For instance, if a Chinese acquirer realizes that there is skepticism about its cross-border acquisitions in Australia, it may attempt an alternative route to the approach that was used with Rio Tinto in 2009.

Although lexicon-based methods have been widely used for sentiment analysis (Esuli and Sebastiani 2005; Riloff et al. 2005), a low recall of sentiment identification may occur. The reason is that domain-specific sentiment indicators may not be defined in a generic sentiment lexicon. One of the main contributions of our research is the design of an automated method for domain-specific sentiment lexicon construction that enhances sentiment analysis. An unsupervised


About the Company | Display Suppliers | Display Competitors | Display Collaborators

Title:

Rio Tinto Appointment Fails to Calm Merger Fears in Australia

Date:

March 17, 2009 00:00 PM GMT

Content:

Rio Tinto PLC (ADR:RTP) appointed a new chairman yesterday (Tuesday) in what was seen as an effort to calm political turmoil in Australia over its $19.5 billion tie-up with China's Aluminum Corp. of China Ltd. (ADR: ACH), also known as Chinalco.

Rio entered the agreement with Chinalco in February to help cut its $39 billion debt burden, but the deal has **<neg-sent>ruffled</neg-sent> <aspect>political feathers</aspect>** in Australia on concerns key assets are **<neg-sent>falling</neg-sent>** into **<aspect>Chinese hands</aspect>**.

After a previous **<aspect>attempts</aspect> <neg-sent>failed</neg-sent>**, Rio appointed Jan du Plessis, chairman of British American Tobacco PLC (BAT), as its new chairman. Rio chairman-elect Jim Leng quit the board last month, after **<neg-sent>objecting</neg-sent>** to the **<aspect>deal</aspect>**.

But that did little to deter Senator Barnaby Joyce, an Australian politician who took out television ads on Tuesday urging that the **<aspect>Chinalco deal</aspect>** be **<neg-sent>blocked</neg-sent>**, Reuters reported.

Figure 5. Sociocultural Issues of Chinese Companies' Cross-Border M&As

learning method is preferred because manually labeling training examples to train a supervised classifier for each application domain is extremely labor intensive. As Figure 1 (module 7) shows, the inputs for the proposed sentiment lexicon learning module are a large, domain-specific corpus of documents (e.g., financial news articles) and generic lexicons, such as OpinionFinder (Riloff et al. 2005) and SentiWordNet (Esuli and Sebastiani 2005). The proposed learning method is called *pseudo labeling* since it explores both term-to-term and term-to-document relations of a domain-specific corpus to automatically infer the polarity label of each document, thereby extracting the domain-specific sentiments from within each document. The proposed PL algorithm is depicted in Figure 6.

Initially, a generic sentiment lexicon (e.g., OpinionFinder) is used to identify domain-independent sentiments. For example, strong subjective sentiments with one single polarity defined in OpinionFinder are treated as domain-independent sentiments. According to the general linguistic rule of contextual coherence (Kanayama and Nasukawa 2006), some of the initial domain-specific sentiments can be extracted from a corpus *D*. The contextual coherence rule states that sentiments colocated within the same textual unit (e.g., a clause) should have the same polarity unless a negation word is found. In addition, the BMI measure is applied to extract domain-specific sentiments associated with the seeding sentiments. According to the set of domain-specific sentiments extracted so far, a pseudo polarity label can be inferred for each document in the corpus *D*. According to the distributional characteristic, a positive document tends to contain more positive terms than negative terms. In addition, a strong positive term tends to appear repeatedly in many of the positive documents of a corpus (Kindo et al. 1997; Lau, Bruza et al. 2008). Therefore, the distributional characteristic can be exploited to discover positive and negative sentiments based on a pseudo labeled corpus. Finally, the domain-specific sentiment lexicon consists of sentiments extracted by using the distributional characteristic, contextual coherence, the statistical term association estimated according to the BMI measure, and the domain-independent sentiments.

To extract domain-specific sentiments using the distributional characteristic of terms exhibited in pseudo labeled documents, a variant of the keyword classifier (KC) (Kindo et al. 1997) is applied. The KC method has been successfully adapted to extract positive or negative keywords representing an information seeker's specific information retrieval (IR) preferences (Lau, Bruza 2008); it has also been applied to extraction of sentiment indicators for opinion mining (Lau, Lai et al. 2009; Lau, Lai, and Li 2009). We calibrate the KC formulation to develop a sentiment extraction formula that estimates the

Algorithm PseudoLabeling(L^G,D,POS, ϖ_{bmi}, ϖ_{label}, ϖ_{KC})

Inputs: L^G /* a generic sentiment lexicon
D /* a collection of documents of a specific domain
POS /* preferred part-of-speech for domain-specific sentiment extraction
ϖ_{bmi} /* a threshold for extracting sentiments using the BMI measure
ϖ_{label} /* a threshold for assigning pseudo labels to documents of a corpus
ϖ_{KC} /* a threshold for extracting sentiments using the polarity score formula

Output: L^D /* a domain-specific sentiment lexicon

Main Procedure:

1. extract domain-independent sentiments S^G from a generic sentiment lexicon L^G e.g., only the strong subjective sentiments in OpinionFinder;
2. apply the principle of contextual coherence to extract domain-specific sentiments S^C based on the set S^G and the pre-defined POS;
3. construct an initial domain-specific sentiment lexicon: $L^D := S^G \cup S^C$;
4. expand L^D by examining each term $s' \in D$ with specific part-of-speech defined in POS; examine its statistical association with each sentiment $s \in L^D$ using the balance mutual formation formula *BMI(s,s')*;
5. let the set of domain-specific sentiments extracted by BMI method be $S^{BMI} = \{s' \in D: BMI(s,s') > \varpi_{bmi}\}$, where ϖ_{bmi} is a sentiment extraction threshold;
6. let $L^D := S^{BMI} \cup S^G \cup S^C$;
7. assign pseudo label to each document $d \in D$ based on the initial domain-specific sentiment lexicon L^D;
8. if the document label score $label(d) = \frac{|PosSent| - |NegSent|}{|PosSent| + |NegSent|} > \varpi_{label}$, the document $d \in D$ is positive, where $PosSent$ ($NegSent$) represents the set of positive (negative) sentiments found in d;
9. if $label(d) < -\varpi_{label}$, the document $d \in D$ is negative; otherwise it is neutral;
10. extract additional domain-specific sentiments using the keyword classifier formula *polarity(t)* and the pre-defined POS;
11. let the set of additional domain-specific sentiments be $S^{KC} = \{t \in D: polarity(t) > \varpi_{KC}$ *or* $polarity(t) < -\varpi_{KC}\}$, where ϖ_{KC} is a sentiment extraction threshold;
12. let the final domain-specific sentiment lexicon be $L^D := S^{KC} \cup S^{BMI} \cup S^G \cup S^C$;
13. return the domain-specific sentiment lexicon L^D;

Figure 6. Pseudo Labeling Algorithm for Domain-Specific Sentiment Lexicon Learning

degree of polarity for each potential token with specific POS, such as adjective and adverb (Subrahmanian and Reforgiato 2008):

$$Polarity(t) = \tanh\left(\frac{df(t)}{\varpi_{pos}} \times \Pr(pos|t) \times \log_2 \frac{\Pr(pos|t\}}{\Pr(pos)} - \frac{df(t)}{\varpi_{neg}} \times \Pr(neg|t) \times \log_2 \frac{\Pr(neg|t)}{\Pr(neg)}\right)$$

The parameters ϖ_{pos} and ϖ_{neg} control the learning rates for positive and negative evidences, respectively. These parameters are established and refined by using the ABIMA system's customization and adaptation module. The hyperbolic tangent function *tanh* ensures that the induced polarity score of a sentiment falls in the unit interval. The term

$$\Pr(pos|t) = \frac{df(t_{pos})}{df(t)}$$

is the estimated conditional probability that a training document is positive given that it contains the particular token t. The term $\Pr(pos|t)$ is estimated based on the fraction of the number of positive documents that contain the token t over the total number of documents containing t.

Similarly,

$$\Pr(neg|t) = \frac{df(t_{neg})}{df(t)}$$

is the estimated conditional probability that a document is negative given that it contains the token t. The document frequency $df(t_{neg})$ represents the number of negative documents that contain the token t. In addition,

$$\Pr(pos) = \frac{|D^+|}{|D^+|+|D^-|}$$
$$\Pr(neg) = \frac{|D^-|}{|D^+|+|D^-|}$$

is *a priori* probability that a document is positive (negative), respectively; $D^+(D^-)$ is the set of positive (negative) documents pseudo labeled using the initial sentiment lexicon. If the *Polarity*(t) of a token t is greater than a threshold ϖ_{KC}, the token is treated as a positive sentiment; if *Polarity*(t) $\leq \varpi_{KC}$ is true, the token is treated as a negative sentiment; otherwise, it is neutral. Only positive or negative sentiments are included in our domain-specific sentiment lexicon. Table 1 shows some of the top sentiments automatically discovered by applying the proposed PL algorithm to our financial corpus.

Sentiment analysis for a financial document is conducted by first consulting the domain-specific sentiment lexicon. If a token is not found in the domain-specific sentiment lexicon, the generic sentiment lexicons are consulted. To analyze the sentiments embedded in financial documents, both the sentiments and the associated "aspects" (Jo and Oh 2011; Thet et al. 2010) are identified by our computational model. Aspects are also called features for certain problem domains (Archak et al. 2007, 2011; Lau, Lai et al. 2009). Aspects are the financial, political, technological, and social factors that can largely influence the success of M&As (Birkinshaw et al. 2000; Deresky 2011; Shimizu et al. 2004). We apply the GATE general information engineering service as the basis for conducting the extraction of aspects (noun phrases). For each sentiment identified in a financial document, if a noun phrase appears nearby (defined by a proximity factor ϖ_{Asp} of words within the same sentence), the noun phrase is considered a relevant aspect associated with that sentiment.

For the proposed computational model, each sentiment indicator si in a document is a triple $si = (n_i, s_i, f_i)$,where n_i, s_i, and f_i represent the negation indicator, the sentiment, and the aspect, respectively. The negation indicator is optional, and any ordering of these three elements in a sentiment indicator is allowed. Figure 5 shows some examples of ABIMA's markups (with XML-like syntax) for M&A aspects and sentiments. To improve readability, an un-stemmed version is shown. For instance, "political feathers" is an aspect to which the negative sentiment "ruffled" is applied. The polarity score of a financial document is computed according to the average polarity score of the sentiment indicators found in that document. The polarity score of a sentiment indicator is determined based on the score of the sentiment defined in our domain-specific sentiment lexicon. If a negation indicator is included in the sentiment indicator, the sign of the polarity score is reversed. The sentiment score of an M&A deal, that is, *Sent*(*deal*), is a weighted average polarity score of all the financial documents (news or comments) concerning that deal computed with respect to a predefined due diligence time window.

M&A Target Scoring

Howson (2006, pp. 81-97) proposed applying Porter's five forces model to identify targeted M&A sectors, and then selected target companies. Accordingly, the ABIMA system applies an extension of the five forces model (Brandenburger and Nalebuff 1996) to support target selection and due diligence inquiries regarding M&As. In particular, the proposed system adopts an M&A scorecard approach (Galpin and Herndon 2007, pp. 175-203) to highlight the most important assessment dimensions (i.e., evaluation criteria) of M&As. Due to limited space, this paper only discusses the proposed due diligence scorecard at the targeted company level rather than the targeted industry level. Apart from assessing the financial factors, the sociocultural fitness of a targeted company is considered a critical success factor in M&As in general and in cross-border M&As in particular (Calipha et al. 2010; Harding and Rouse 2007; Shimizu et al. 2004; Weber et al. 1996).

In addition, assessing the internal strengths and weaknesses of the involved companies (i.e., business fitness) is believed to be an important step for M&As (Deresky 2011; Garbuio et al. 2010; Reed et al. 2007). Therefore, the proposed due diligence scorecard takes the following main dimensions of the potential M&A targets into account: financial strength, sociocultural fitness, business fitness, and competitive advantage. We stress that the proposed due diligence scorecard for ABIMA is not intended to replace the traditional due diligence methods (Emott 2011; Howson 2006). Instead, it provides M&A practitioners with valuable complementary information based on the collective intelligence extracted from the Web 2.0 information environment.

Common approaches to evaluating the financial strength of targeted firms include a market-based approach (e.g., capital intensity as a percentage of sales) and an income-based approach (e.g., discounted cash flow) (Emott 2011; Reed et al. 2007). Moreover, common financial measures applied to the due diligence of M&As include ROE, ROA, cash flow to debt ratio, net asset ratio, EBITDAR, debt to equity ratio, and others. Because well-established financial appraisal methods for M&As already exist (Emott 2011; Howson 2006; Reed et al. 2007), we will not discuss them in detail in this paper. To allow ABIMA users to view information at different levels of

Table 1. Top-15 Items in Automatically Generated Sentiment Lexicon

Positive Sentiment	Strength	Negative Sentiment	Strength
lifted	0.831	weakened	0.843
supported	0.829	red	0.843
backed	0.825	suspended	0.842
aggressive	0.824	alleged	0.842
earned	0.824	incurable	0.835
completed	0.821	overgrown	0.834
tony	0.821	blocked	0.832
winning	0.819	fighting	0.829
cheaper	0.817	slashed	0.828
encourage	0.816	stuck	0.825
topped	0.814	harder	0.821
advancing	0.811	tighter	0.819
pushy	0.808	dropping	0.817
grown	0.807	demanding	0.816
eased	0.791	conservative	0.805

granularity (Yan et al. 2011), the various financial indicators can be combined using weighted linear or nonlinear aggregation methods, tuned according to a GA-based learning and adaptation method (module 11 shown in Figure 1).

For the various financial scores, they are first converted to unit intervals via linear normalization and then combined using one of the following two methods:

$$val_{aggregated} = \sum_{m \in M} w(m) \times val(m)$$

$$val_{aggregated} = \frac{1}{1 + \exp\left(-\sum_{m \in M} w(m) \times val(m)\right)}$$

where $val_{aggregated}$ is an aggregated financial score based on the individual normalized financial score $val(m)$ of a specific financial measure m among the set of standard financial measures M. The weight of each financial measure $w(m)$ is subject to $\sum_{m \in M} w(m)$. A weighted combination of various financial metrics for target ranking is a common practice during due diligence inquiry (Emott 2011, p. 51).

The sociocultural fitness dimension is estimated using three methods: cultural clustering (Gupta et al. 2002), affect analysis, and sentiment analysis. If the acquirer *acq* and target *tar* belong to the same cultural cluster, then the cultural fitness function $fit_{cluster}(acq, tar)$ is 1; otherwise, it is 0. In addition, if an acquirer has announced an M&A attempt for which preliminary due diligence has been conducted, then both the sentiment and affect of the potential deal are analyzed. Every public comment about an M&A deal is treated as a document, with the average sentiment scores of all such documents computed. In addition, to provide more in-depth analysis of public emotion about a potential M&A deal, an affect score is calculated on the basis of all affect indicators appearing in the same document. If the ABIMA user conducts a preliminary target-selection exercise (i.e., the potential M&A deal is not disclosed), sociocultural fitness is estimated based on public feeling in the targeted nation concerning the acquiring nation and its companies. Under such circumstances, all news and comments about the two nations or companies concerned are collected, and sentiment and affect analysis is then applied. In addition, any major events (e.g., natural disasters) that take place in the targeted nation are analyzed via sentiment analysis. Figure 7 presents the results of sentiment analysis of a news article concerning the financial performance of Fuji Heavy Industries in the wake of the Tohoku earthquake and tsunami that occurred in March 2011. Such events affect the sociocultural fitness of potential targets in the area in question.

In general, recent observations are assigned greater weight when they are applied to predict an investment's returns (Huang et al. 2009). For M&A target scoring, recent events generally have a greater impact on the current fitness of a potential target compared to events that occurred a long time

Figure 7. Societal Event Affecting Fuji Heavy Industries

ago. Accordingly, an exponential decay function is applied to aggregate the sentiment or affect scores of a potential M&A target over a predefined time window of due diligence (e.g., 36 months). An exponential decay function is often used to model the time-varying phenomena in financial markets (Barari and Mitra 2008; Jo et al. 1997). Specifically, the following exponential decay function is used to weight a sentiment or affect score:

$$score(d,t) = score(d) \times e^{-\left(\frac{t_{current}-t}{\tau/2}\right)}$$

where $score(d)$ is the raw sentiment or affect score without weighting, and $score(d, t)$ is the weighted score of document d posted at time point t. The term τ is the due diligence time window specified in terms of months, and the term $(t_{current} - t)$ represents the elapsed time (in months) between the time t when a financial document containing sentiments or affects is posted and the time $t_{current}$ when M&A scoring is conducted. For a sentiment or affect score computed based on a document posted in the same month that M&A target scoring is performed, the elapsed time of $(t_{current} - t)$ is zero. Finally, the resulting scores from cultural clustering and the decayed scores from sentiment and affect analysis are combined via a weighted linear or nonlinear aggregation function similar to that used to estimate a target's financial strength.

For the dimension of business fitness, sentiment analysis is first applied to assess the strengths and weaknesses of typical aspects (e.g., products and services, production process, distribution channels, sales-force, etc.) of each M&A target (Deresky 2011; Emott 2011; Reed et al. 2007). Following this, an aggregated sentiment score is computed for each potential target. Moreover, the synergy between an acquirer and a potential target is estimated in terms of their complementary resources (Calipha et al. 2010; Shimizu et al. 2004). For this purpose, the result of business relation mining is utilized. Specifically, the ABIMA system uses a coarse-grained method to estimate the synergy score. If both the acquirer and the target appear in a supply chain relation, their synergy score is assigned a 1; otherwise, it is given a 0. Both the aggregated sentiment score and the synergy score of a target are combined via a weighted linear or nonlinear formula, which is similar to the one applied to the aforementioned financial strength scoring. A weighted combination of the respective scores pertaining to the strengths and weaknesses of an M&A target, and the synergy between an acquirer and a target, has been shown to be a practical approach for a due diligence inquiry (Emott 2011, p. 25).

The competitive advantage dimension is assessed based on Porter's five forces model, augmented by the complementary force (Brandenburger and Nalebuff 1996). In particular, competitive forces such as the bargaining power of customers, the bargaining power of suppliers, the threat of substitutes, the intensity of rivalry, the threat of new entrants, and the effect of complementors (the sixth force) are taken into account by the proposed computational method. Among these forces, the bargaining power of suppliers $bs(tar)$, the intensity of rivalry $ir(tar)$, and the complementary force $cp(tar)$ of a targeted company tar are estimated based on the results of business relation mining. For instance, the number of suppliers connected to a targeted company determines the bargaining power of the target; the more suppliers are connected to the target, the less bargaining power each individual supplier has with respect to the target (i.e., the target enjoys a greater competitive advantage). Similarly, the intensity of rivalry of a target *ir*(*tar*) is estimated based on the number of competitors connected to the target. The more competitors are connected to the target, the less competitive advantage the target may have. However, the complementors (e.g., the strategic partners) of a company *cp*(*tar*) are considered beneficial (Brandenburger and Nalebuff 1996). Therefore, the more collaborators connected to the target, the more competitive power the targeted firm may have.

Formally, the bargaining power of suppliers with respect to a target *bs*(*tar*), the intensity of rivalry of a target *ir*(*tar*), and the complementary force of a target *cp*(*tar*) are defined by the following:

$$bs(tar) = \frac{1}{\sum_{l \in Supp(tar)} w(1)}$$

$$ir(tar) = \frac{\sum_{l \in Comp(tar)} w(l)}{\sum_{tar' \in Sector} \sum_{l \in Comp(tar')} w(l)}$$

$$cp(tar) = \frac{\sum_{l \in Coll(tar)} w(l)}{\sum_{tar' \in Sector} \sum_{l \in Coll(tar')} w(l)}$$

where the function *Supp*(*tar*) returns the set of immediate suppliers connected to the node representing the targeted company. The term *w*(*l*) represents the weight of a link *l* and is computed based on the occurrence frequency of the specific type of business relation observed in the training corpus. The function *Comp*(*tar*) represents the set of immediate competitors connected to the target node. The set *Sector* represents all of the companies in the specific industrial sector under evaluation. The function *Coll*(*tar*) returns the set of immediate collaborators connected to the target node.

A targeted company's bargaining power of customers and the threat of new entrants *te*(*tar*) are estimated based on market research data and financial data made available on the Internet. If the competitive analysis is conducted at the post-acquisition stage, additional information can be collected from the intranets of the acquirer and the acquired company, respectively. For instance, the threat of new entrants for a targeted company can be estimated based on its production plant and equipment costs (Howson 2006). The ABIMA system returns a default value of zero for these two factors if the relevant quantitative data is not available from the external information sources. To estimate the force of threat for substitutes *ts*(*tar*) of a targeted company *tar*, the ABIMA system crawls its product-related Web pages to extract the noun phrases describing the targeted company's main products.

The product descriptions (noun phrases) of a targeted company are then combined into a long document. At this point, term frequency, inverse document frequency (TFIDF) document representation and the cosine score measure (Salton and McGill 1983) are used to estimate the similarity of product lines between the target and those of its competitors. If the average cosine score is high, it suggests that the threat of substitutes may be high. Finally, the competitive advantage of a targeted company *ca*(*tar*) is estimated according to the weighted linear or nonlinear aggregation formula similar to that applied to financial strength scoring. For the weighted sum of competitive advantage, *bc*(*tar*), *ir*(*tar*), *bce*(*tar*), *te*(*tar*), and *ts*(*tar*) take a negative sign (i.e., reducing a target's competitive advantage), whereas *cp*(*tar*) takes a positive sign (i.e., increasing a target's competitive advantage).

Learning and Adaptation

It is not realistic to assume that a single M&A target scoring function is optimal for a wide variety of M&A situations. Accordingly, a learning and adaptation process is needed to adjust the proposed scoring function with respect to different M&A scenarios. The ABIMA system supports two types of learning and adaptation processes, namely system-wide adaptation for M&A decision support and user personalization. This paper focuses on the discussion of the former type. For system-wide learning and adaptation, a set of system and application parameter values and some low-level NLP features are continuously refined to drive the effective operations of sentiment analysis, business relation mining, and M&A target scoring under different M&A scenarios. This kind of parameter learning and customization process is considered crucial in supporting adaptive investment decision making (Armano et al. 2005; Chen 2002; Huang et al. 2009; Kuo et al. 2001; Ruiz-Torrubiano 2010). The system-wide learning and adaptation process is invoked periodically (e.g., quarterly), but it can also be launched on demand (e.g., when applying the system to different kinds of M&As). It utilizes a genetic algorithm (Goldberg 1989; Holland 1992) to conduct a heuristic search over a large solution space (e.g., the set of feasible parameter values) by taking into account environmental signals, such as a set of training M&A cases and the corresponding financial document corpus. For the training M&A cases, they are successful M&A deals in the real-world or some representative targets selected by M&A experts.

GAs have been successfully applied in ducting heuristic searches over large search spaces to obtain near-optimal solutions (Goldberg 1989; Holland 1992; Lau, Tang. 2006). Recently, a more advanced hierarchical coevolutionary approach has been successfully applied to adaptive stock trend prediction and financial trading (Huang et al. 2009). Accordingly, we have designed a hierarchical coevolution genetic algorithm (HCGA) to continuously refine the system's M&A target scoring function (i.e., a decision-support mechanism) according to environmental signals. Coevolution refers to the simultaneous evolution of two or more species (i.e., subspaces of a large solution space) with the fitness of each individual determined by other individuals of the same species as well as the individuals of other species via interspecies interactions (Delgado et al. 2004; Olsson 2001). One advantage of the coevolutionary approach is that a large solution space can be explored more effectively and efficiently through a parallel and diversified search over the subspaces (Olsson 2001).

For system-wide learning and adaptation, there are high-level species (i.e., a population of application parameters) and low-level species (i.e., two populations of system parameters), which coevolve at the same time. Each individual of a species (i.e., a population) is represented by a chromosome, which carries a fixed number of genes (Goldberg 1989; Holland 1992). For our HCGA algorithm, each gene encodes the feasible value of an application or system parameter. According to the principle of natural selection (Darwin 1936), a population of individuals gradually evolves to produce fitter individuals representing near-optimal solutions for a problem domain. For the first generation of a population, each chromosome is randomly assigned feasible gene values. The chromosomes of the current population are selected to produce offspring according to their fitness, as evaluated with respect to a fitness function that is defined according to the precision of ABIMA's M&A recommendations. In other words, the system's M&A target scoring module (i.e., the decision support model) is invoked whenever the fitness of a chromosome is assessed.

Using a coevolutionary approach, selected individuals (e.g., those who are relatively fit) are exchanged among the coevolved populations in order to evaluate the fitness of each individual of each participating population (Delgado et al. 2004; Huang et al. 2009; Olsson 2001). Following this, standard genetic operators such as selection, cross-over, and mutation (Goldberg 1989; Huang et al. 2009; Lau, Tang et al. 2006) are applied to evolve the chromosomes of the current population to produce the next generation. Such an evolutionary process is repeated until certain termination conditions (e.g., the average fitness of the whole population exceeds a predefined threshold) are met. The fittest chromosomes (i.e., representing the near-optimal parameter values and NLP features) from each population of the final generation are then selected and applied to M&A target scoring. The computational details of the proposed HCGA algorithm are illustrated in Appendix A.

System Evaluation

General Evaluation Procedures and Data Sets

Evaluation of the ABIMA system was conducted under the business context of Chinese enterprises' cross-border M&As. In particular, our experiments were performed based on a set of core companies tabulated in the Forbes 2000 list.[6] Since the financial information regarding some Forbes companies is not available in English or Chinese, these companies were excluded from our test. As a result, a total of 1,568 core companies were included in our test set. The ABIMA system collected public information (both quantitative and qualitative) about the core companies from the Internet using our crawler programs and external APIs. A total of 765,103 financial documents (financial news, investors' comments, and blog messages) covering the period from January 1, 2008, to November 30, 2011, were downloaded to ABIMA's local database. In addition, the key financial indicators (e.g., ROE, PE, cash flow to debt ratio, EBITDAR, etc.) of the respective companies were also retrieved for quantitative financial analysis by the ABIMA system. A subset of our financial corpus was applied to evaluate the proposed sentiment analysis method and the business relation mining method.

To evaluate the proposed sentiment analysis method, 325 financial news articles were manually annotated by two human experts. There were 1,265 sentences containing sentiment indicators (812 positive and 679 negative) and 1,151 sentences not containing sentiment indicators. To evaluate the proposed business relation mining method, 766 financial news articles were manually annotated by the same human experts. This evaluation data set consisted of 516 collaborative relationships, 497 competitive relationships, and 406 sup-

[6]http://www.forbes.com/2009/04/08/worlds-largest-companies-business-global-09-global_land.html.

Table 2. Confusion Matrix for System Evaluation

System's Judgment		Gold Standard – Human's Judgment	
		Yes	No
	Yes	*TruePositive*	*FalsePositive*
	No	*FalseNegative*	*TrueNegative*

plier relationships. The ABIMA system was then applied to conduct sentiment analysis and business relation mining. The effectiveness of the proposed computational methods was assessed with respect to the set of manually annotated sentences. Since each M&A deal is conducted on a case-by-case basis in practice, we selected some companies from specific industrial sectors (e.g., the consumer durables and the materials industries according to the Forbes 2000 classification) to develop test cases for evaluating the effectiveness of the proposed due diligence scorecard.

On two different occasions (in February and November, 2011), we invited five Chinese M&A experts who specialize in cross-border M&As to examine the chosen M&A target-selection scenarios and develop their "gold standard." A majority vote was first applied to identify an appropriate M&A target for each given business scenario of a particular occasion. If there was a disagreement among the experts, they would discuss to resolve the conflict. Only if all of the experts agreed on a chosen M&A target would it be included in the gold standard set. Both the M&A experts and the ABIMA system referred to the same test set of 1,568 companies for M&A target selection. A subset of these established test cases was applied to the experiments reported in this paper, all of which were conducted using a DELL 1950 III Server with Quad-Core Xeon 2.33GHz Processors, 16GB main memory, and 6TB secondary storage. For the experiments reported in this paper, only public financial information collected from the Internet was utilized.

Performance Measures

The performance measures commonly adopted in information retrieval and opinion mining research were utilized in our experiments (Ounis et al. 2008; Salton and McGill 1983). With reference to the confusion matrix (Table 2), the various performance measures are defined by the following:

$$Recall = \frac{TruePositive}{TruePositive + FalseNegative}$$

$$Precision = \frac{TruePositive}{TruePositive = FalseNegative}$$

$$F_\beta = \frac{(1+\beta^2)\,Precision \times Recall}{\beta^2\,Precision + Recall}$$

$$Accuracy = \frac{TruePositive + TrueNegative}{TruePositive + FalsePositive + FalseNegative + TrueNegative}$$

where *TruePositive*, *FalsePositive*, *FalseNegative*, and *TrueNegative* refer to the number of predictions by the system that fall into each corresponding category in the confusion matrix (Table 2). *Recall* is the fraction of correct positive predictions made by the system from all of the positive predictions judged by human means. *Precision* is the fraction of correct positive predictions from all of the positive predictions produced by the system. F_β (F-measure) is the harmonic mean of precision and recall. We adopted $\beta = 1$ for the experiments reported in this paper.

For the TREC Blog Track (i.e., the opinion retrieval benchmark test), the measure P@10 (precision for top 10 documents) has been used to evaluate the performance of opinion retrieval systems (Ounis et al. 2008). We applied the P@10 measure to evaluate the quality of our system's recommendations via the due diligence scorecard. Since information users rarely review information items beyond the first page of a result set (Granka et al. 2004), it is considerably more important to evaluate the precision of the M&A recommendations ranked in the top 10 positions. Due to the complex dynamics of M&A activities, our preliminary experimentation showed that even M&A experts might not be completely certain about the *TruePositive* set given a particular business context. Specifically, after viewing the ABIMA system's due diligence scorecard, our M&A experts might feel a need to modify the *TruePositive* set (i.e., the gold standard). Accordingly, we propose the notion of an adjusted P@10 to capture such a phenomenon in system testing. Basically, the adjusted P@10 measure and the P@10 use the same formula. However, the *TruePositive* set of the adjusted P@10 measure is based on the revised *TruePositive* set developed after an expert has reviewed ABIMA's recommendations via the due diligence scorecard.

Experimental Results and Discussions

The First Controlled Experiment

For the first controlled laboratory experiment, we examined the effectiveness of the proposed PL algorithm by comparing the sentiment indicators identified by the system with those annotated by humans. Among the 2,416 annotated sentences, 1,265 sentences contained sentiment indicators (812 positive and 679 negative), and the remaining 1,151 sentences did not contain sentiment indicators. Considering that there were three classes involved in this classification task (positive sentiment indicators, negative sentiment indicators, and no sentiment indicator), the evaluation of the system's classification performance for the positive and negative classes was conducted separately. As the proposed sentiment analysis approach was empowered by the unsupervised method of domain-specific sentiment lexicon construction, we also indirectly evaluated the effectiveness of the pseudo labeling algorithm. Two other baseline methods were also examined in this experiment. The first one involved using the general sentiment lexicon SentiWordNet (Esuli and Sebastiani 2005) to identify sentiment indicators in financial documents. The second baseline method employed OpinionFinder (Riloff et al. 2005) to do the same.

The third baseline method employed a state-of-the-art supervised classifier, the conditional random field (CRF) classifier (Finkel et al. 2008; Sarawagi 2006), to learn and assign the label sequences of sentiment indicators, given various financial text segments. CRF is based on a discriminative probabilistic model (an undirected graph model), with each vertex representing a random variable and the edge representing the dependency between two random variables. The ultimate goal was to estimate the probability distribution of each vertex based on training data. We used the CRF package[7] developed by Sarawagi (2006) and adopted the default parameters for this experiment. For each class, 80 percent of the sentences were used as the training set, and the remaining 20 percent were held out as the test set. This process was repeated 10 times to produce 10 test sets. Each computational method was then evaluated 10 times, based on the test sets, to derive the average performance score. For the PL algorithm, labeled sentences were not required to train a classifier. Instead, 42,000 financial documents (excluding the documents used to build the evaluation dataset) were used as the corpus to generate the domain-specific sentiment lexicon for subsequent sentiment analysis. The results of this experiment are shown in Table 3.

[7]http://crf.sourceforge.net/.

For positive sentiment indicator identification, it is clear that the proposed PL method outperforms the SentiWordNet method, the OpinionFinder method, and the CRF method in terms of $F_{\beta=1}$ by 5.5 percent, 6.1 percent, and 5.8 percent, respectively. Moreover, for negative sentiment indicator identification, the proposed PL method outperforms the SentiWordNet method, the OpinionFinder method, and the CRF method by 4.5 percent, 8.6 percent, and 8.1 percent, respectively. The experimental results show that the PL method can greatly improve recall while maintaining a fair amount of precision compared to other baseline methods. Such a performance improvement is achieved by the effective unsupervised learning (extraction) of domain-specific sentiment indicators from the training financial text segments. For instance, the sentiment indicator "red" is not included in the OpinionFinder lexicon. However, the PL method can automatically discover this sentiment indicator and correctly predict its domain-specific polarity (i.e., a negative sentiment) with respect to the financial domain. Subsequently, the ABIMA system can make use of these domain-specific sentiment indicators to improve the accuracy of sentiment analysis.

Surprisingly, the SentiWordNet method also performs slightly better than the CRF method in both classification tasks. The reason for this may be that it is quite a challenging task for a CRF classifier to automatically learn the sequence labeling knowledge when given a small training set. Therefore, its performance is not as strong as that of the PL method. This experiment demonstrates that the concept of designing an unsupervised statistical learning method to extract domain-specific sentiment indicators from a large unlabeled corpus of financial documents to enhance sentiment analysis is promising. The merits of the proposed unsupervised sentiment analysis method stem not only from its better classification performance, but also from its great potential to be applied to different business domains with minimal human intervention.

The Second Controlled Experiment

For the second controlled experiment, we examined the effectiveness of the proposed company relation mining method (i.e., the BRME algorithm) by comparing the business relations identified by the system to those annotated by human means. Since some sentences contained no business relation at all from among the sentences containing the 516 collaborative relations, 497 competitive relations, and 406 supplier relations in the evaluation corpus, there were four classes for this classification problem. We evaluated the performance of the proposed business relation mining method separately under each classification task. The proposed BRME algorithm utilized an unsupervised statistical learning method to

Table 3. Comparative Performance of Sentiment Indicator Identification

Method	Recall	Precision	F-measure	Accuracy
Identification of Positive Sentiment Indicators				
PL	0.716	0.686	0.701	0.834
CRF	0.654	0.671	0.663	0.762
SentiWordNet	0.648	0.682	0.665	0.764
OpinionFinder	0.642	0.679	0.660	0.766
Identification of Negative Sentiment Indicators				
PL	0.719	0.634	0.674	0.818
CRF	0.637	0.609	0.623	0.772
SentiWordNet	0.652	0.638	0.645	0.785
OpinionFinder	0.629	0.612	0.620	0.776

Table 4. Comparative Performance of Business Relation Mining

Method	Recall	Precision	F-measure	Accuracy
Identification of Collaborative Business Relations				
BRME	0.660	0.618	0.639	0.770
CRF	0.583	0.594	0.588	0.738
BASIC	0.592	0.587	0.589	0.735
Identification of Competitive Business Relations				
BRME	0.636	0.583	0.609	0.783
CRF	0.566	0.560	0.563	0.735
BASIC	0.576	0.594	0.585	0.740
Identification of Supplier Relations				
BRME	0.616	0.604	0.611	0.763
CRF	0.575	0.571	0.572	0.726
BASIC	0.565	0.583	0.574	0.716

identify various abbreviations of company names and expanded the seeding company relation lexicon to enhance the performance of business relation mining.

One of the baseline systems (BASIC) applied the standard company names, downloaded from Yahoo Finance and Forbes to identify companies. For business relation identification, this baseline system employed only the seeding relation lexicon. Both the BRME and the BASIC methods employed the same co-reference resolution method proposed in this paper. The second baseline method employed the CRF classifier with individual words as features to label the sequences of business relations. For each type of business relation classification, 80 percent of the relations of the annotated corpus were used as the training set and the remaining 20 percent were held out as the test set. Splitting of the training and test sets was repeated 10 times to produce 10 test sets. All of the systems being tested were evaluated using 10 test sets, and their average performance scores were then computed. The results of our experiment are shown in Table 4.

For collaborative business relation mining, it is clear that the proposed BRME algorithm outperforms the BASIC and the CRF methods in terms of $F_{\beta=1}$ by 8.3 percent and 8.5 percent, respectively. Moreover, for competitive business relation mining, the proposed method outperforms the BASIC and the CRF methods by 4.1 percent and 8.2 percent, respectively. For supplier relation mining, the proposed method outperforms the BASIC and the CRF methods by 6.2 percent and

6.5 percent, respectively. The BRME method performs better than the BASIC method because it can utilize unsupervised learning to discover various abbreviations of a company name and additional relation indicators based on a large financial corpus. For instance, the BRME method discovers the token "spx" as an abbreviation for the company Spirax-Sarco Engineering PLC. Accordingly, it can correctly identify a business relation whenever "spx" instead of Spirax-Sarco Engineering PLC is used to describe its relationship with another company in a testing sentence. Moreover, additional collaborative relation indicators such as "advisor" are automatically extracted by the BRME method to correctly classify collaborative business relations when a board member of one company serves as the advisor to another company. In contrast, the BASIC method cannot utilize company abbreviations nor automatically discovered relation indicators to classify business relationships. As a result, the recall of the BASIC method is considerably lower than that of the BRME method.

Interestingly, the BASIC method even performs slightly better than the CRF method for three types of business relation mining tasks in terms of the F-measure. This may be caused by the fact that the BASIC method has also adopted the proposed co-reference resolution method to effectively identify companies that appear in a business relation. In contrast, the CRF method mistakenly classifies many nonbusiness relations as one of the three business relations and vice versa, given a relatively small labeled training set. For instance, it classifies many sentences containing company names and the word "shares" as collaborative relations. Unfortunately, these sentences only indicate that the shares of both companies go up or down. This experiment demonstrates that the concept of designing an unsupervised learning method to automatically extract company abbreviations and relation indicators from a large unlabeled corpus of financial documents to improve business relation mining is promising. The main advantage of the BRME algorithm is that labeled training examples are not required to perform business relation mining from financial documents. This increases the chance of deploying the proposed method to real-world mining tasks.

The Third Controlled Experiment

The objective of the third controlled experiment was to evaluate whether the ABIMA system could provide significant decision aid to people engaging in M&A target-selection activities. All of the subjects involved were MBA students who had attended introductory lessons about M&As. These subjects voluntarily participated in this experiment and attended a 45 minute briefing session about the four cross-border M&A target-selection tasks in November 2011. They were randomly assigned to an experimental group (29 subjects) or a control group (28 subjects). The experimental group was offered an extra 30-minute tutorial session to learn the basic features of the ABIMA system after the common briefing session. After the briefing session, the subjects were given one week to submit a paper reporting their M&A recommendations. For each submission, the subjects were asked to rank the top 10 targeted companies and provide justifications for why they chose these targets. In addition, each subject agreed to fill out a questionnaire to provide additional feedback about these M&A tasks. Each subject was free to access the Internet or other data sources to obtain the information needed to identify the potential targets. The experimental group was provided with access accounts to the ABIMA system, and the subjects could refer to the ABIMA due diligence scorecard and the financial and sociocultural information about each potential target listed on it. This was the only treatment introduced to the experimental group.

For the first two M&A target-selection tasks, SAIC Motor and Dongfeng Motor Group were chosen as the Chinese acquirers, and the foreign companies belonging to the consumer durables industry from the Forbes 2000 list were assumed to be the potential targets. In addition, China Shenhua Energy and Baoshan Iron & Steel were selected as the Chinese acquirers for the remaining two M&A target-selection tasks, and the foreign firms belonging to the materials industry from the Forbes 2000 list were the potential targets. The gold standard developed by our M&A experts in November 2011 was used as the basis to assess the performance of the subjects' M&A target-selection tasks. The P@10 measure was used to quantify the subjects' performance. In addition, qualitative comments received from the subjects were also analyzed. Two-tail *t*-tests were applied to compare the M&A target-selection performance of the respective groups.

The mean and standard deviation of the P@10 scores achieved by the respective groups appear in Table 5, which shows that the experimental group consistently outperforms the control group in these M&A target-selection tasks. The differences between the two groups are statistically significant. Therefore, we conclude that the ABIMA system can provide considerable decision aid to people engaging in M&As. The qualitative comments received from the experimental group also indicated that the ABIMA system provided useful information to the subjects who were only M&A task novices. In particular, the majority of the subjects in the experimental group reflected that ABIMA's due diligence scorecard not only ranked the potential targets with respect to the most important M&A selection criteria, but also allowed easy access to both financial and sociocultural information

Table 5. Comparative Performance on the M&A Target-Selection Tasks

	Experimental Group		Control Group			
Industry: Consumer Durables	**Mean (P@10)**	**SD (P@10)**	**Mean (P@10)**	**SD (P@10)**	***t* statistic**	***p***
SAIC Motor	0.410	0.067	0.332	0.082	$t(55) = 3.944$	= .000
Dongfeng Motor Group	0.414	0.064	0.329	0.081	$t(55) = 4.417$	= .000
Industry: Materials	Mean (P@10)	SD (P@10)	Mean (P@10)	SD (P@10)		
China Shenhua Energy	0.407	0.075	0.321	0.096	$t(55) = 3.756$	= .000
Baoshan Iron & Steel	0.397	0.078	0.336	0.112	$t(55) = 2.375$	= .021

about the targets. By simply clicking the drill-down button beside a potential target listed on the due diligence scorecard, the key financial indicators and the relevant financial documents with highlighted opinion aspects were displayed. This information facilitated the M&A target-selection tasks considerably. In contrast, some of the subjects in the control group reported that there was an overwhelming amount of information on the Internet that they needed to filter out before it could be applied to the given M&A tasks.

The Field Tests

The aims of our field tests were threefold. First, we wanted to evaluate the usage of the ABIMA system by M&A experts in real-world settings. Second, we intended to assess the adaptive decision support capability of the ABIMA system. Third, we wanted to test whether the ABIMA system can provide complementary decision support to M&A experts. Two Chinese M&A experts who had previously participated in developing our M&A test cases were invited to take part in a series of field tests held at their offices. Four M&A target-selection tasks examined in the third controlled experiment were applied to these field tests. The first two field test sessions were conducted in the middle of November 2011, and the other two sessions took place one week later. For each field test session, online access to the Web-based ABIMA system was provided to an M&A expert, and a member of our research team guided them through the system.

For the first two field test sessions, the ABIMA system was configured to simulate the business environment in February 2011 (i.e., the ABIMA system only archived information up to the end of February 2011). As part of the system configuration process, the system's learning and adaptation process was invoked before the field tests began. With the assumption that it was the period of February 2011, each expert was guided to use the ABIMA system to conduct M&A target-selection tasks for four chosen Chinese companies (e.g., SAIC Motor). After entering an M&A query into the system, the expert examined the output of the due diligence scorecard. The top 10 recommendations from each of the four assessment dimensions (e.g., sociocultural fitness) on the scorecard were recorded by our research team member to calculate the P@10 score in the system. The P@10 score for each dimension of the due diligence scorecard was computed based on the gold standard developed by our M&A experts in February 2011.

A macro average of the P@10 scores for the four assessment dimensions was computed to assess the system's M&A target-selection performance for each M&A query in that period. To assess whether the ABIMA system can provide complementary decision support to M&A practitioners, each expert was invited to examine the top 10 recommendations under each of the four scorecard dimensions. If an expert decided that the recommendations provided by the ABIMA system were appropriate after careful examination of both the financial and nonfinancial details (e.g., sentiment analysis results) for the recommended targets, these *TruePositive* cases were recorded by the research team member to compute the adjusted P@10 score. This procedure was repeated for each of the four M&A target-selection tasks. In addition, each expert's qualitative feedback about the usability of the ABIMA system was recorded. Figure 2 highlights ABIMA's due diligence scorecard for the SAIC Motor case for the simulated period of February 2011.

For the remaining two field test sessions, the system was configured for the period of November 2011 (i.e., the same period when the field tests were performed). Again, as part of the configuration process, the system's learning and adaptation process was initiated before the field test sessions began. Basically, the two M&A experts carried out tasks similar to those they performed in the respective test sessions one week earlier. However, all of the M&A target-selection tasks were conducted with respect to the current period (i.e., November 2011). After an M&A query was submitted, the

Table 6. Performance of ABIMA in Two Different Time Periods

	Field Tests (February 2011)			Field Tests (November 2011)		
M&A Cases	**P@10**	**Adj. P@10**	**Adj. P@10 (All Experts)**	**P@10**	**Adj. P@10**	**Adj. P@10 (All Experts)**
SAIC Motor	0.425	0.613	0.575	0.450	0.600	0.575
Dongfeng Motor Group	0.400	0.625	0.600	0.425	0.638	0.600
China Shenhua Energy	0.425	0.663	0.650	0.400	0.650	0.625
Baoshan Iron & Steel	0.400	0.600	0.575	0.400	0.588	0.575
Mean	0.413	0.625	0.600	0.419	0.619	0.594

ABIMA system's recommendations were recorded to calculate the average P@10 score for that query. In addition, each expert examined whether the system's top 10 recommendations from each assessment dimension of the due diligence scorecard were indeed good alternatives.

According to an expert's judgment, the adjusted P@10 score of each query was calculated. The average P@10 and adjusted P@10 scores for the field test sessions conducted in each period are summarized in columns 2, 3, 5, and 6 of Table 6. According to two-tail *t*-test, significant differences in the P@10 scores ($t(6) = -.45$, $p = .67$) achieved by the system were not found on these two occasions. Nevertheless, obvious changes in the system's top 10 recommendations along the four scorecard dimensions were observed on these different occasions. For example, 22.5 percent of the top 10 recommended targets in the February 2011 period did not appear again as top recommendations in the November 2011 period for the SAIC Motor case. For all four M&A test cases, the average difference in the top 10 recommended targets for these two occasions was 25.3 percent.

These field test results suggest that the ABIMA system can adaptively make different M&A recommendations with respect to changing business environment in different time periods, and yet the quality of its recommendations is maintained. Figure 8 shows ABIMA's due diligence scorecard for the SAIC Motor case in the November 2011 period. As is evident in the figure, the Japanese motor company Fuji Heavy Inds. was removed from the top 10 recommendations of the sociocultural fitness dimension in November 2011. A closer look at the system's sentiment analysis reveals that a striking societal event, the Tohoku earthquake and tsunami, occurred in March 2011. Figure 7 shows the sentiment analysis results for one of the news articles revealing the poor business performance of Fuji Heavy Inds. due to the Tohoku earthquake and tsunami. From this news article, it is evident that the company had been an M&A target of a Chinese enterprise although the deal was not approved by the Chinese government.

In addition, the third and sixth columns of Table 6 show the higher adjusted P@10 score for each M&A target-selection case on two different occasions. This illustrates that the ABIMA system may provide complementary decision support to M&A practitioners. To verify such a hypothesis, we conducted a follow-up study by inviting the remaining three M&A experts who participated in developing the gold standard of the test cases to carefully examine the system's recommendations archived on these two occasions. Only if all five experts agreed that a recommendation made by the system was an appropriate supplement to their original gold standard would that recommendation be included in the revised *TruePositive* set. The adjusted P@10 scores computed based on all experts' judgment are shown in the fourth and seventh columns of Table 6, respectively. These higher adjusted P@10 scores when compared to the P@10 scores (e.g., 43 percent higher on average) suggest that the ABIMA system can produce quality M&A recommendations not originally identified by our M&A experts. This provides another confirmation of the system's ability to provide complementary decision support to M&A practitioners.

According to the qualitative feedback of the M&A experts, the ABIMA system is quite useful for efficiently scanning a large volume of qualitative information and alerting users to important events that may affect their M&A decision making. Specifically, they found that the due diligence scorecard of the ABIMA system was very helpful in providing a quick preliminary assessment of potential M&A targets. The experts felt that the scorecard and the qualitative details of each recommended target provided complementary information to enhance their M&A decision-making processes. One expert also noted that the network diagram of a company or the entire industrial sector was a very useful decision support tool for M&A practitioners. In particular, the result of business

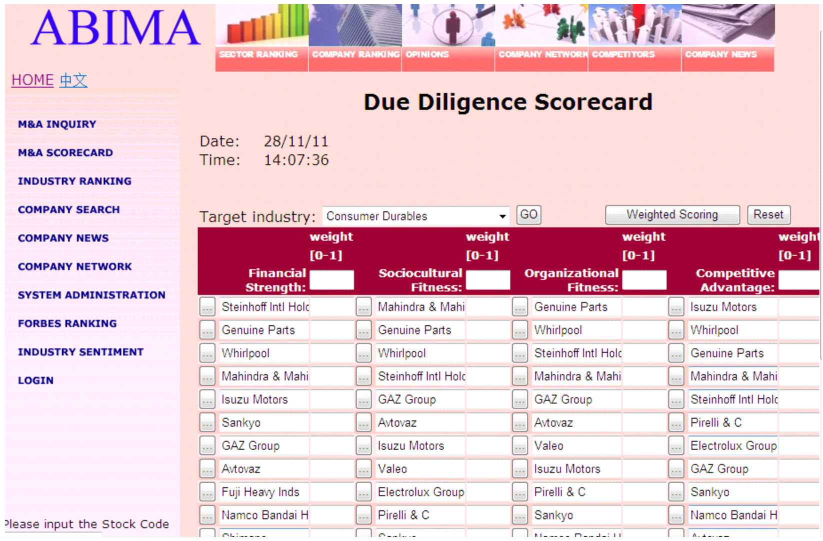

Figure 8. Due Diligence Scorecard for the SAIC Motor Case (November 2011)

relation mining helped in identifying a potential M&A target's relationships with other comparable companies when the "comparable net worth to market value" method (Reed et al. 2007) was applied during a full due-diligence inquiry. In contrast, the experts commented that they would like to see an improvement in the ABIMA prototype system that included information about more companies that did not appear in the Forbes 2000 list. This is a positive sign that points to the potential for upgrading the current prototype system to a fully operational commercial software package in the future.

Summary and Conclusions

Existing research shows that top management's tendency to neglect sociocultural and other nonfinancial factors has led to a poor track record of M&As. In this era of Web 2.0, a vast volume of user-generated qualitative data regarding the sociocultural and political–economic issues of industrial sectors or companies are readily available on the Internet. Top executives and M&A consultants have unprecedented opportunities to tap into valuable business intelligence (e.g., the sociocultural knowledge about a targeted market) by continuously scanning the Web 2.0 environment. Grounded in Porter's five forces model, one major contribution of our research is the design of a novel due diligence scorecard model that leverages collective web intelligence to enhance M&A decision making. Another important contribution of our work is the design and development of an adaptive BI 2.0 system, which is underpinned by a hierarchical evolutionary learning approach and unsupervised methods for sentiment analysis and business relation mining to operationalize the due diligence scorecard model. Specifically, NLP and unsupervised statistical learning techniques have been exploited to design the computational algorithms for domain-specific sentiment analysis and business relation mining for M&A intelligence discovery. By using the proposed BI 2.0 system to continuously scan the Web 2.0 environment, the due diligence scorecard is able to provide adaptive M&A decision support to top management.

The results of our controlled laboratory experiments show that the proposed domain-specific sentiment analysis method and the unsupervised business relation mining method are effective, outperforming other well-known baseline methods. Looking at the business context of Chinese companies' cross-border M&As, the results of our experiment and field tests confirm that the ABIMA prototype system can provide significant aid to decision makers engaging in M&A activities. Moreover, the recommendations generated by the system are adaptive with respect to changing business contexts. Essentially, all the research questions raised in this paper have been answered through our study. To the best of our knowledge, this is the very first adaptive BI 2.0 system successfully designed to support M&A activities. The managerial implication of our research is that top executives can apply the proposed BI 2.0 technology to enhance their corporate investment decision making. In particular, it is a very useful complementary decision support tool for cross-border corporate investment for which private information of the targeted markets may not be readily available. Furthermore, because the proposed BI 2.0 technology is underpinned by unsupervised statistical learning techniques, it has good potential to be applied in other business domains, such as financial risk identification, bankruptcy prediction, and investment portfolio management. The societal implication of our research is that fair corporate financial investment and trading may be promoted because organizations that lack private market information can tap into collective web intelligence to enhance their investment decision-making processes.

Our future work will involve evaluating both the effectiveness and the efficiency of the ABIMA system, based on more real-world M&A scenarios. We will conduct more field tests to evaluate the system-wide adaptation and the personalization mechanisms of the ABIMA system. We will invite dozens of M&A practitioners to try the ABIMA system for a few months. Based on the questionnaires returned by these M&A practitioners, we will further analyze the usability of the ABIMA system. The bilingual information processing capability of the system will also be evaluated in the context of inbound M&As in the Chinese financial market. Moreover, a more sophisticated business relation mining method will be developed to capture the specific directions of competitive and collaborative relationships. Both local and global relationships captured in a business network will be exploited to estimate a company's competitiveness.

Acknowledgments

The authors would like to sincerely thank the three anonymous reviewers, the associate editor, and the guest editors of the BI research special issue of *MIS Quarterly* for their constructive comments that were so helpful in improving the quality of the paper. The authors would particularly like to thank the following cross-border M&A experts who helped develop the M&A test cases for our experiments and offered constructive comments regarding the design of the ABIMA prototype system: Mr. Michael Wu (Senior Manager) of PricewaterhouseCoopers (China), Mr. Dylan Tey (Senior Manager) of PricewaterhouseCoopers (China), Mr. Li Shang Heng (Partner & CFA) of Shang Heng Fund L.P., and Ms. Ammi Hong (Senior Finance Manager & CPA) of CMHJ Partners. The work described in this paper was partially supported by a grant from the Research Grants Council of the Hong Kong Special Administrative Region, China (Project No. CityU 115910).

References

Abbasi, A. and Chen, H. 2008. "Cybergate: A Design Framework and System for Text Analysis of Computer-Mediated Communication," *MIS Quarterly* (32:4), pp. 811-837.

Abbasi, A., Chen, H., and Salem, A. 2008. "Sentiment Analysis in Multiple Languages: Feature Selection for Opinion Classification in Web Forums," *ACM Transactions on Information Systems* (26:3), Article 12.

Abbasi, A., Chen, H., Thoms, S., and Fu, T. 2008. "Affect Analysis of Web Forums and Blogs Using Correlation Ensembles," *IEEE Transactions on Knowledge and Data Engineering* (20:9), pp. 1168-1180.

Albright, K. S. 2004. "Environmental Scanning: Radar for Success," *Information Management Journal* (38:3), pp. 38-44,

Andrade, G., Mitchell, M., and Stafford, E. 2001. "New Evidence and Perspectives on Mergers," *Journal of Economic Perspectives* (15:2), pp. 103-120.

Antweiler, W., and Frank, M. 2004. "Is All That Talk Just Noise? The Information Content of Internet Stock Message Boards," *Journal of Finance* (59:3), pp. 1259-1294.

Arazy, O., and Woo, C. 2007. "Enhancing Information Retrieval Through Statistical Natural Language Processing: A Study of Collocation Indexing," *MIS Quarterly* (31:3), pp. 525-546.

Archak, N., Ghose, A., and Ipeirotis, P. 2007. "Show Me the Money! Deriving the Pricing Power of Product Features by Mining Consumer Reviews," in *Proceedings of the 13th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, San Jose, CA, pp. 56-65.

Archak, N., Ghose, A., and Ipeirotis, P. 2011. "Deriving the Pricing Power of Product Features by Mining Consumer Reviews," *Management Science* (57:8), pp. 1485-1509.

Armano, G., Marchesi, M., and Murru, A. 2005. "A Hybrid Genetic–Neural Architecture for Stock Indexes Forecasting," *Information Science* (170), pp. 3-33.

Bao, S., Li, R., Yu, Y., and Cao, Y. 2008. "Competitor Mining with the Web," *IEEE Transactions on Knowledge and Data Engineering* (20:10), pp. 1297-1310.

Barboza, D., and Wines, M. 2009. "Mining Giant Scraps China Deal," *New York Times*, June 4 (http://www.nytimes.com/2009/06/05/business/global/05mine.html?_r=2).

Barari, M., and Mitra, S. 2008. "Power Law Versus Exponential Law in Characterizing Stock Market Returns," *Atlantic Economic Journal* (36:3), pp. 377-379.

Barney, J. 1991. "Firm Resources and Sustained Competitive Advantage," *Journal of management* (17:1), pp. 99-120.

Batagelj, V., and Mrvar, A. 1998. "Pajek: Program for Large Network Analysis," *Connections* (21:2), pp. 47-57.

Bernstein, A., Clearwater, S., Provost, F. 2003. "The Relational Vector–Space Model and Industry Classification," in *Proceedings of Workshop on Learning Statistical Models from Relational Data*, pp. 8-18.

Bhagwan, V., Grandison, T., and Gruhl, D. 2009. "Sound Index: Charts for the People, by the People," *Communications of the ACM* (52:9), pp. 64-70.

Birkinshaw, J., Bresman, H., and Hakanson, L. 2000. "Managing the Post-Acquisition Integration Process: How the Human Integration and Task Integration Processes Interact to Foster Value Creation," *Journal of Management Studies* (37), pp. 395-425.

Bollen, J., Maoa, H., and Zengb, X. 2011. "Twitter Mood Predicts the Stock Market," *Journal of Computational Science* (2:1), pp. 1-8.

Brandenburger, A. M., and Nalebuff, B. J. 1996. *Co-opetition*, New York: Doubleday.

Brody, S., and Elhadad, N. 2010. "An Unsupervised Aspect–Sentiment Model for Online Reviews," in *Proceedings of the 2010 Annual Conference of the North American Chapter of the Association for Computational Linguistics*, pp. 804-812.

Calipha, R., Tarba, S., and Brock, D. 2010. "Mergers and Acquisitions: A Review of Phases, Motives, and Success Factors," in *Advances in Mergers and Acquisitions* (Volume 9), C. Cooper and S. Finkelstein (eds.), Bingley, UK: Emerald Group Publishing, pp. 1-24.

Calix, R. A., Mallepudi, S. A., Chen, B., and Knapp, G. M. 2010. "Emotion Recognition in Text for 3-D Facial Expression Rendering," *IEEE Transactions on Multimedia* (12:6), pp. 544-551.

Calvo, R. A., and D'Mello, S. 2010. "Affect Detection: An Interdisciplinary Review of Models, Methods, and Their Applications," *IEEE Transactions on Affective Computing* (1:1), pp. 18-37.

Cartwright, S., and Cooper, C. L. 1993. "The Role of Culture Compatibility in Successful Organizational Marriage," *The Academy of Management Executive* (7), pp. 57-70.

Chen, H. 2010. "Business and Market Intelligence 2.0," *IEEE Intelligent Systems* (25:1), pp. 68-71.

Chen, S. H. 2002. *Evolutionary Computation in Economics and Finance*, Heidelberg: Physica-Verlag.

Choo, C. W. 1999. "The Art of Scanning the Environment," *Bulletin of the American Society for Information Science and Technology* (25), pp. 21-24.

Chung, W., Chen, H., and Nunamaker, Jr., J. F. 2005. "A Visual Framework for Knowledge Discovery on the Web: An Empirical Study of Business Intelligence Exploration," *Journal of Management Information Systems* (21:4), pp. 57-84.

Cunningham, H. 2002. "GATE: A General Architecture for Text Engineering," *Computers and the Humanities* (36), pp. 223-254.

Darwin, C. 1936 (original 1859). *The Origin of Species*, New York: Modern Library.

Das, S. 2010. "The Finance Web: Internet Information and Markets," *IEEE Intelligent Systems* (25:2), pp. 74-78.

Das, S., and Chen, M. 2007. "Yahoo! for Amazon: Sentiment Extraction from Small Talk on the Web," *Management Science* (53:9), pp. 1375-1388.

Delgado, M., Zuben, F., and Gomide, F. 2004. "Coevolutionary Genetic Fuzzy Systems: A Hierarchical Collaborative Approach," *Fuzzy Sets and Systems* (141:1), pp. 89-106.

Deresky, H. 2011. *International Management: Managing Across Borders and Cultures, Text and Cases* (7th ed.), New York: Prentice Hall.

El Sawy, O. A. 1985. "Personal Information Systems for Strategic Scanning in Turbulent Environments: Can the CEO Go On-Line?," *MIS Quarterly* (9:1), pp. 53-60.

Emott D. T. 2011. *Practitioner's Complete Guide to M&As*, Hoboken, NJ: John Wiley & Sons.

Esuli, A., and Sebastiani, F. 2005. "Determining the Semantic Orientation of Terms through Gloss Classification," in *Proceedings of the 14th ACM International Conference on Information and Knowledge Management,* Bremen, Germany, pp. 617-624.

Farhoomand, A. F., and Drury, D. H. 2002. "Managerial Information Overload," *Communications of the ACM* (45:10), pp. 127-131.

Finkel, J. R., Kleeman, A., and Manning, C. D. 2008. "Efficient, Feature-Based, Conditional Random Field Parsing," in *Proceedings of the Annual Meeting of the Association for Computational Linguistics*, pp. 959-967.

Foo, S., and Li, H. 2004. "Chinese Word Segmentation and its Effect on Information Retrieval," *Information Processing and Management* (40:1), pp. 161-190.

Galpin T., and Herndon M. 2007. *The Complete Guide to Mergers and Acquisitions: Process Tools to Support M&A Integration at Every Level*, San Francisco: John Wiley & Sons.

Garbuio, M., Lovallo, D., and Horn, J. 2010. "Overcoming Biases in M&A: A Process Perspective," in *Advances in Mergers and Acquisitions* (Volume 9), C. Cooper and S. Finkelstein (eds.), Bingley, UK: Emerald Group Publishing, pp. 83-104.

Gnyawali, D. R., and Madhavan, R. 2001. "Cooperative Networks and Competitive Dynamics: A Structural Embeddedness Perspective," *Academy of Management Review* (26:3), pp. 431-446.

Goldberg D. 1989. *Genetic Algorithms in Search, Optimization and Machine Learning*, Reading, MA: Addison-Wesley.

Granka, L., Joachims, T., and Gay, G. 2004. "Eye-Tracking Analysis of User Behavior in WWW Search," in *Proceedings of the 27th Annual International ACM SIGIR Conference on Research and Development in Information Retrieval*, pp. 478-479.

Gupta, V., Hanges, P. J., Dorfman, P. 2002. "Cultural Clusters: Methodology and Findings," *Journal of World Business* (37), pp. 11-15.

Harding, D., and Rouse, T. 2007. "Human Due Diligence," *Harvard Business Review* (85:4), pp. 124-131.

He, W., and Lyles, M. A. 2008. "China's Outward Foreign Direct Investment," *Business Horizons* (51:6), pp. 485-491.

Hevner, A., March, S., Park, J., and Ram, S. 2004. "Design Science in Information Systems Research," *MIS Quarterly* (28:1), pp. 75-105.

Holland, J. H. 1992. "Genetic Algorithms," *Scientific American* (267:1), pp. 66-72.

Howson, P. 2006. *Commercial Due Diligence: The Key to Understanding Value in an Acquisition*, Burlington, VT: Gower.

Huang, H., Pasquier, M., Quek, C. 2009. "Financial Market Trading System with a Hierarchical Coevolutionary Fuzzy Predictive Model," *IEEE Transactions on Evolutionary Computation* (13:1), pp. 56-70.

Islam, A., and Inkpen, D. 2008. "Semantic Text Similarity Using Corpus-Based Word Similarity and String Similarity," *ACM Transactions on Knowledge Discovery from Data* (2:2), Article 10.

Jo, H., Han, I, and Lee, H. 1997. "Bankruptcy Prediction Using Case-Based Reasoning, Neural Networks, and Discriminant Analysis," *Expert Systems With Applications* (13:2), pp. 97-108.

Jo, Y., and Oh, A. 2011. "Aspect and Sentiment Unification Model for Online Review Analysis," in *Proceedings of the 4th International Conference on Web Search and Data Mining*, pp. 815-824.

Kanayama, H., and Nasukawa, T. 2006. "Fully Automatic Lexicon Expansion for Domain-Oriented Sentiment Analysis," in *Proceedings of the 2006 Conference on Empirical Methods in Natural Language Processing*, pp. 355-363.

Kaplan, R. S., and Norton, D. P. 1996. "Using the Balanced Scorecard as a Strategic Management System," *Harvard Business Review*, January-February, pp.75-85.

Kindo, T., Yoshida, H., Morimoto, T., and Watanabe, T. 1997. "Adaptive Personal Information Filtering System that Organizes Personal Profiles Automatically," in *Proceedings of the 15th International Joint Conference on Artificial Intelligence*, pp. 716-721.

Kuo, R. J., Chen, C. H., and Hwang, Y. C. 2001 "An Intelligent Stock Trading Decision Support System through Integration of Genetic Algorithm Based Fuzzy Neural Network and Artificial Neural Network," *Fuzzy Sets Systems* (118), pp. 21-45.

Lau, R. Y. K. 2003. "Context-Sensitive Text Mining and Belief Revision for Intelligent Information Retrieval on the Web," *Journal of Web Intelligence and Agent Systems* (1:3-4), pp. 151-172.

Lau, R. Y. K., Bruza, P. D., and Song, D. 2008. "Towards a Belief Revision Based Adaptive and Context Sensitive Information Retrieval System," *ACM Transactions on Information System* (26:2), pp. 8.1-8.38.

Lau, R. Y. K., Lai, C. L., and Li, Y. 2009. "Leveraging the Web Context for Context-Sensitive Opinion Mining," in *Proceedings of the 2009 IEEE International Conference on Computer Science and Information Technology,* pp. 467-471.

Lau, R. Y. K., Lai, C. L., Ma, J., and Li, Y. 2009. "Automatic Domain Ontology Extraction for Context-Sensitive Opinion Mining," in *Proceedings of the 30th International Conference on Information Systems*, Phoenix, AZ, Completed Research Paper 171.

Lau, R. Y. K., Song, D., Li, Y., Cheung, C. H., and Hao, J. X. 2009c. "Towards A Fuzzy Domain Ontology Extraction Method for Adaptive e-Learning," *IEEE Transactions on Knowledge and Data Engineering* (21:6), pp. 800-813.

Lau, R. Y. K., Tang, M., Wong, O., Milliner, S., and Chen, Y. 2006. "An Evolutionary Learning Approach for Adaptive Negotiation Agents," *International Journal of Intelligent Systems* (21:1), pp. 41-72.

Ma, Z., Pant, G., Sheng, O. R. L. 2009. "A Network-Based Approach to Mining Competitor Relationships from Online News," in *Proceedings of the 30th International Conference on Information Systems*, Phoenix, AZ, Completed Research Paper 59.

Ma, Z., Sheng, O. R. L., and Pant, G. 2009. "Discovering Company Revenue Relations from News: A Network Approach," *Decision Support Systems* (47:4), pp. 408-414.

Madura, J., Vasconcellos, G., and Kish, R. 1991. "A Valuation Model for International Acquisitions," *Management Decision* (29:4), pp. 31-38.

March, S. T., and Storey, V. C. 2008. "Design Science in the Information Systems Discipline: An Introduction to the Special Issue on Design Science Research," *MIS Quarterly* (32:4), pp. 725-730.

Marks, M. L., and Mirvis, P. H. 2001. "Making Mergers and Acquisitions Work: Strategic and Psychological Preparation," *Academy of Management Executive* (15:2), pp. 80-92.

Maynard, D., Tablan, V., Ursu, C., Cunningham, H., and Wilks, Y. 2001. "Named Entity Recognition from Diverse Text Types," in *Proceedings of the 2001 Conference on Recent Advances in Natural Language Processing*, Tzigov Chark, Bulgaria.

McEwen, T. 2008. "Environmental Scanning and Organizational Learning in Entrepreneurial Ventures," *The Entrepreneurial Executive* (13:1), pp. 1-16.

Miller, G., Beckwith, R., Fellbaum, C., Gross, D., and Miller, K. 1990. "Introduction to WordNet: An On-Line Lexical Database," *Journal of Lexicography* (3:4), pp. 234-244.

Murphy, T. 2005. "CNOOC Abandons $18.5B Unocal Offer," *Forbes.com Business News Bulletin*, August 8 (http://www.forbes.com/2005/08/02/cnooc-unocal-oil-cx_tm_0802video.html).

Olsson, B. 2001. "Co-evolutionary Search in Asymmetric Spaces," *Information Sciences* (133:3-4), pp. 103-125.

Ong, T. H., Chen, H., Sung, W. K., and Zhu, B. 2005. "Newsmap: A Knowledge Map for Online News," *Decision Support Systems* (39:4), pp. 583-597.

Oreilly, T. 2007. "What Is Web 2.0: Design Patterns and Business Models for the Next Generation of Software," *Communications & Strategies* (1), pp. 17.

Ounis, I., Macdonald, C., and Soboroff, I. 2008. "Overview of the TREC-2008 Blog Track", in *Proceedings of the 17th Text Retrieval Conference*, Gaithersburg, MD.

Pang, B., Lee, L., and Vaithyanathan, S. 2002. "Thumbs Up? Sentiment Classification Using Machine Learning Techniques," in *Proceedings of the 2002 Conference on Empirical Methods in Natural Language Processing*, Philadelphia, pp. 79-86.

Pant, G., and Sheng, O. R. L. 2009. "Avoiding the Blind Spots: Competitor Identification Using Web Text and Linkage Structure," in *Proceedings of the 30th International Conference on Information Systems*, Phoenix, AZ, Completed Research Paper 57.

Peffers, K., Tuunanen, T, Rothenberger, M., and Chatterjee, S. 2008. "A Design Science Research Methodology for Information Systems Research," *Journal of Management Information Systems* (24:3), pp. 45-77.

Porter, M. E. 1980. *Competitive Strategy: Techniques for Analyzing Industries and Competitors*, New York: Free Press.

Porter, M. E. 1985. *Competitive Advantage: Creating and Sustaining Superior Performance*, New York: Free Press.

Ptaszynski, M., Maciejewski, J., Dybala, P., Rzepka, R., and Araki, K. 2010. "CAO: A Fully Automatic Emoticon Analysis System Based on Theory of Kinesics," *IEEE Transactions on Affective Computing* (1:1), pp.46-59.

Raman, T. V. 2009. "Toward 2^W, Beyond Web 2.0," *Communications of the ACM* (52:2), pp. 52-59.

Reed, S. F., Lajoux, A., and Nesvold, H. P. 2007. *The Art of M&A: A Merger/Acquisition/Buyout Guide* (4th ed.), New York: McGraw-Hill.

Riloff, E., Wilson, T., Hoffmann, P., Somasundaran, S., Kessler, J., Wiebe, J., Choi, Y., Cardie, C., and Patwardhan, S. 2005. "Opinionfinder: A System for Subjectivity Analysis," in *Proceedings of the 2005 Human Language Technology Conference and Conference on Empirical Methods in Natural Language Processing*, pp. 34-35.

Ruiz-Torrubiano, R., and Suarez, A. 2010. "Hybrid Approaches and Dimensionality Reduction for Portfolio Selection with Cardinality Constraints," *IEEE Computational Intelligence Magazine* (5:2), pp. 92-107.

Salton, G., and McGill, H. J. 1983. *Introduction to Modern Information Retrieval*, New York: McGraw-Hill.

Sarawagi, S. 2006. "Efficient Inference on Sequence Segmentation Models," in *Proceedings of the ACM International Conference on Machine Learning*, pp. 793-800.

Shimizu, K., Hitt, M. A., Vaidyanath, D., and Pisano, V. 2004. "Theoretical Foundations of Cross-Border Mergers and Acquisitions: A Review of Current Research and Recommendations for the Future," *Journal of International Management* (10), pp. 307-353

Srivastava, J., and Cooley, R. 2003. "Web Business Intelligence: Mining the Web for Actionable Knowledge," *INFORMS Journal on Computing* (15:2), pp. 191-207.

Stahl, G., and Voigt, A. 2008. "Do Cultural Differences Matter in Mergers and Acquisitions? A Tentative Model and Examination," *Organization Science* (19:1), pp. 160-176.

Subrahmanian, V., and Reforgiato, D. 2008. "AVA: Adjective-Verb-Adverb Combinations for Sentiment Analysis," *IEEE Intelligent Systems* (23:4), pp. 43-50.

Tan, H., and Ai, Q. 2010. "China's Outward Mergers and Acquisitions in the 21st Century: Motivations, Progress and the Role of the Chinese Government," in *Advances in Mergers and Acquisitions* (Volume 9), C. Cooper and S. Finkelstein (eds.), Bingley, UK: Emerald Group Publishing, pp. 25-50.

Tetenbaum, T. J. 1999. "Beating the Odds of Merger and Acquisition Failure: Seven Key Practices that Improve the Chance for Expected Integration and Synergies," *Organizational Dynamics* (28:2), pp. 22-36.

Tetlock, P., Saar-Tsechansky, P. M., and Macskassay, S. 2008. "More than Words: Quantifying Language to Measure Firms' Fundamentals," *Journal of Finance* (63:3), pp. 1437-1467.

Thet, T., Na, J., and Khoo, C. 2010. "Aspect-Based Sentiment Analysis of Movie Reviews on Discussion Boards," *Journal of Information Science* (36:6), pp. 823-848.

Turney, P., and Littman, M. 2003. "Measuring Praise and Criticism: Inference of Semantic Orientation from Association," *ACM Transactions on Information Systems* (21:4), pp. 315-346.

Valitutti, A., Strapparava, C., and Stock, O. 2004. "Developing Affective Lexical Resources," *Psychology* (2:1), pp. 61-83.

Weber, Y., Shenkar, O., and Raveh, A. 1996. "National and Corporate Cultural Fit in Mergers/Acquisitions: An Exploratory Study," *Management Science* (42:8), pp.1215-1227.

Wu, C., and Xie, N. 2010. "Determinants of Cross-Border Merger and Acquisition Performance of Chinese Enterprises," *Social and Behavioral Sciences* (2:5), pp. 6896-6905.

Wilson, T., Wiebe, J., and Hwa, R. 2006. "Recognizing Strong and Weak Opinion Clauses," *Computational Intelligence* (22:2), pp. 73-99.

Yan, X. Lau, R. Y. K., Song, D., Li, X., and Ma, J. 2011. "Towards a Semantic Granularity Model for Domain-Specific Information Retrieval," *ACM Transactions on Information Systems* (29:3), Article 15.

Yasai-Ardekani, M., and Nystrom, P. C. 1996. "Designs for Environmental Scanning Systems: Tests of a Contingency Theory," *Management Science* (42:2), pp. 187-204.

About the Authors

Raymond Y. K. Lau is an assistant professor in the Department of Information Systems at City University of Hong Kong. He received his Ph.D. in Information Technology from Queensland University of Technology. His research interests include business intelligence, information retrieval, text mining, social media, and agent-mediated e-Commerce. He is the author of over 100 articles in refereed international journals and conference proceedings, including *ACM Transactions on Information Systems*, *Journal of Management Information Systems*, *IEEE Transactions on Knowledge and Data Engineering*, *IEEE Internet Computing, Computational Intelligence*, ACM SIGIR, ACM CIKM, and the International Conference on Information Systems. He is a member of the AIS, a senior member of both the IEEE and the ACM.

Stephen S. Y. Liao is a professor and the director of the Advanced Transportation Information Systems (ATIS) Research Center at the City University of Hong Kong. He received a bachelor's degree from Beijing University and a Ph.D. from the University of Aix-Marseille III and Institute of France Telecom. His research has been published in journals including *MIS Quarterly*, *Decision Support Systems*, *Communications of the ACM*, *Information Science*, *Computer Software and Applications* and various IEEE transactions. He actively works in the areas of intelligent transportation systems and is currently the project coordinator of several Innovation Technology Funds projects funded by the Hong Kong SAR government. His current research interests include the use of information technology in mobile commerce applications and intelligent business systems, especially intelligent transportation systems.

Kam-Fai Wong is a professor in the Department of Systems Engineering and Engineering Management, and the Associate Dean (External Affairs) of the Faculty of Engineering, Chinese University of Hong Kong. He is also the director of the Centre for Innovation and Technology. He received his Ph.D. from Edinburgh University, Scotland, in 1987. His research interests focus on Internet programming and applications, Chinese computing, parallel databases and information retrieval. He has published over 200 technical papers in these areas. He is the founding editor-in-chief of *ACM*

Transactions on Asian Language Processing and a member of the editorial boards of *International Journal on Computer Processing of Oriental Languages*, and *International Journal on Computational Linguistics and Chinese Language Processing*. He is the author of *Introduction to Chinese Natural Language Processing*, the first book covering this topic in English. Due to his contributions to the information technology sector of Hong Kong, he was awarded the Medal of Honour (MH) by the Hong Kong SAR Government in 2011. He is a fellow of HKIE, IET, and BCS, and a member of the ACM, CLCS, and IEEE-CS.

Dickson K. W. Chiu is a research fellow in the Department of Information Systems at City University of Hong Kong. He received his Ph.D., and M.Sc. in Computer Science from the Hong Kong University of Science and Technology, and a B.Sc. (Hons.) degree in Computer Studies from the University of Hong Kong. His research interests include business intelligence, service science, and design science with a cross-disciplinary approach. He has published over 150 refereed papers in international journals and conference proceedings. He received the best paper award at the 37th Hawaii International Conference on System Sciences in 2004. He is the founding editor-in-chief of *International Journal on Systems and Service-Oriented Engineering* and serves in the editorial boards of several international journals. He is a senior member of both the ACM and the IEEE, and a life member of the Hong Kong Computer Society.

MIS Quarterly

SPECIAL ISSUE: BUSINESS INTELLIGENCE RESEARCH

WEB 2.0 ENVIRONMENTAL SCANNING AND ADAPTIVE DECISION SUPPORT FOR BUSINESS MERGERS AND ACQUISITIONS

Raymond Y. K. Lau, and Stephen S. Y. Liao
Department of Information Systems, City University of Hong Kong, Tat Chee Avenue, Kowloon,
Hong Kong SAR, CHINA {raylau@cityu.edu.hk} {issliao@cityu.edu.hk}

K. F. Wong
Department of Systems Engineering and Engineering Management, The Chinese University of Hong Kong,
Shatin, N.T., Hong Kong SAR, CHINA {kfwong@se.cuhk.edu.hk}

Dickson K. W. Chiu
Department of Information Systems, City University of Hong Kong, Tat Chee Avenue, Kowloon,
Hong Kong SAR, CHINA {kwchiu@cityu.edu.hk}

Appendix A

A Hierarchical Coevolution Genetic Algorithm for Adaptive M&A Decision Support

Figure A1 depicts an overview of the coevolution process controlled by the proposed HCGA algorithm. The Level II (high-level) population consists of individuals representing the feasible values of the application parameters. The Level I (low-level) populations represent two sets of system parameters (i.e., two species). One of the low-level populations represents the system parameters that drive the sentiment analysis process, whereas the second low-level population represents the system parameters that control the business relation mining process. The HCGA algorithm controls the evolutionary processes among all of the populations. At the end of the coevolution process, a set of near-optimal application parameter values and low-level NLP features (e.g., the use of specific sentiment lexicons) with respect to a particular M&A situation are obtained to refine the M&A target scoring function (i.e., the decision support mechanism).

The fitness of an individual from each population is assessed in terms of a high-level measure, that is, the system's performance on M&A target recommendations. In particular, the precision regarding making top 10 recommendations (i.e., P@10) is assessed to determine an individual's fitness. The rationale for using P@10 (Ounis et al. 2008) as the fitness function instead of other quantitative measures such as ROI is that it may take years for an M&A deal to generate the anticipated ROI, and thus this kind of information may not be available to assess the system's performance. The P@10 score is computed with respect to a set of training M&A cases retrieved from the real-world or recommended by M&A experts. The fitness of an individual (chromosome) c is defined by the following:

$$fitness(c) = \frac{TruePositive_{top-10}}{FalsePositive_{top-10} + FalsePositive_{top-10}}$$

where $TruePositive_{top\text{-}10}$ and $FalsePositive_{top\text{-}10}$ represent the true positive and false positive recommendations at the top 10 positions, respectively. In other words, the system's M&A scoring module (i.e., the decision support model) is invoked to generate the top 10 recommendations whenever the fitness of a chromosome needs to be evaluated.

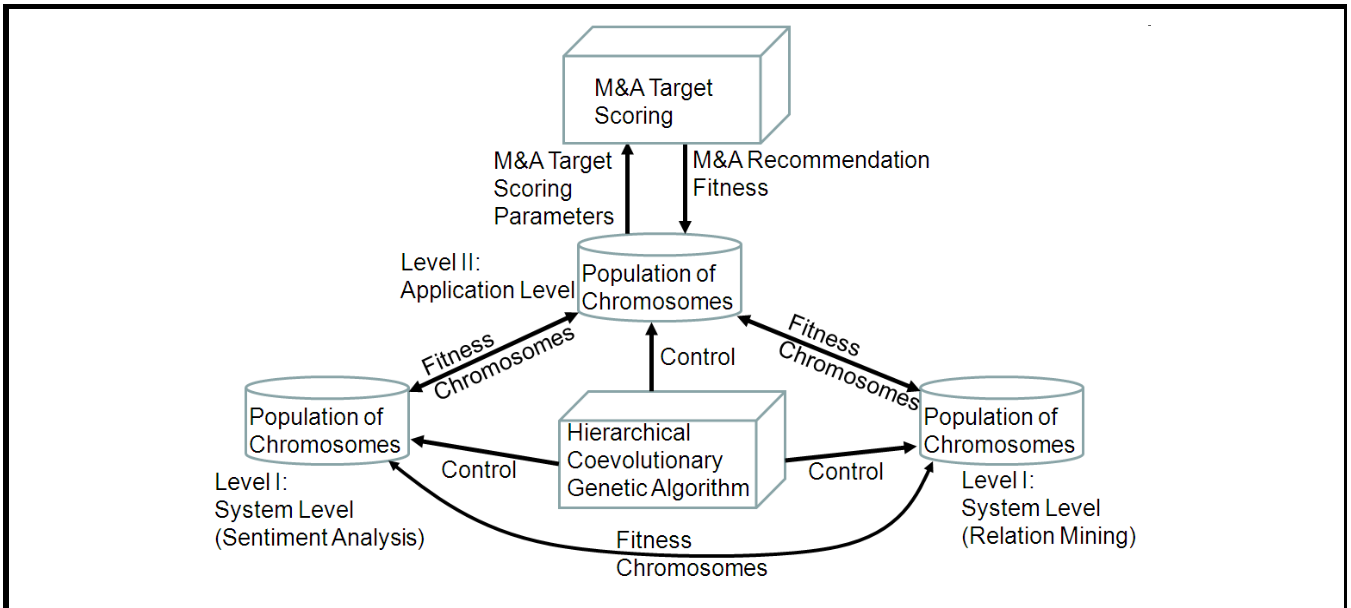

Figure A1. Overview of the Hierarchical Coevolutionary Process

When an HCGA-coordinated evolutionary process takes place, it always begins with the high-level population. To assess the fitness of each individual in the Level II population, the fittest individual from each Level I population is passed to the Level II population. In other words, the sets of system parameters and the set of application parameters are combined to drive the eventual M&A target scoring process. The P@10 score of the resulting M&A recommendations is used to assess the fitness of each individual. Similarly, the fittest individual of the Level II population is passed to a Level I population when the fitness of an individual in a Level I population is assessed. In addition, the fittest individual is exchanged between the two Level I populations. These interactions among the species (populations) drive the coevolution process. The advantage of the coevolutionary approach is that a large solution space can be divided into subspaces for a parallel and diversified search, which improves both the efficiency and the effectiveness of the heuristic search process (Delgado et al. 2004; Olsson 2001).

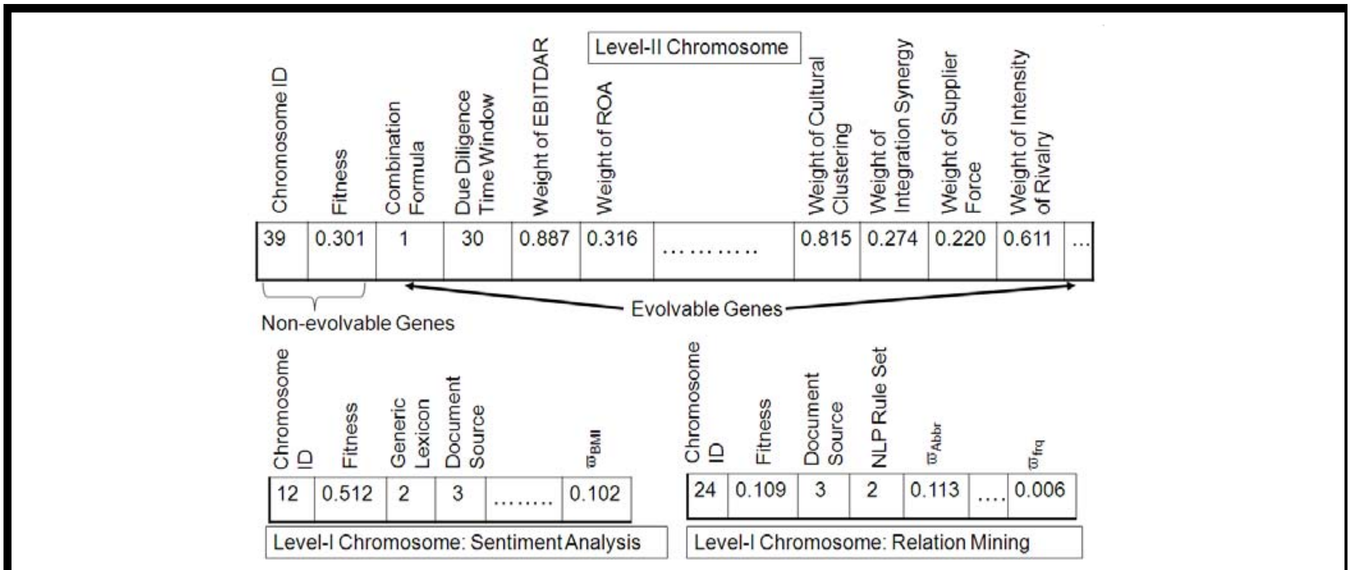

Figure A2. Gene Encoding of the HCGA Algorithm

Algorithm HCGA(Avg_{fit}, Max_{gen}, NP, $P_{size}[x]$, $p_m[x]$, $p_c[x]$, $w_e[x]$)

Inputs: Avg_{fit} /* the average fitness threshold to terminate the algorithm

Max_{gen} /* the maximum no. of generations re-produced

NP /* no. of populations

$P_{size}[x]$ /* the population size of a specific population x

$P_m[x]$ /* the mutation probability of a specific population x

$P_c[x]$ /* the crossover probability of a specific population x

$w_e[x]$ /* the elitism rate of a specific population x

Output: $c_{out}[x]$ /* the fittest chromosomes returned

Main Procedure:

1. $gen := 1$; /* initialize the generation index
2. for $x = 1$ to NP
3. randomly assign feasible value to each gene of each chromosome according to population size $P_{size}[x]$;
4. end
5. exchange chromosomes among high-level and low-level populations;
6. compute the fitness of each chromosome in each population $P[x]^{gen}$ for generation gen according to the fitness function;
7. while $gen <= Max_{gen}$ and $Fitness(P[x]^{gen})$ of each population $< Avg_{fit}$
8. for $x = 1$ to NP
9. select $w_e[x]$ percentage of fittest chromosomes of a population $P[x]^{gen}$ to build $P[x]^{gen+1}$ starting from the high-level population;
10. while $size(P[x]^{gen+1}) <= P_{size}[x]$
11. apply roulette-wheel selection to select two chromosomes to produce $P[x]^{gen+1}$;
12. generate a random number r for each pair of selected chromosome;
13. if $r < P_c[x]$, apply two point crossover to the selected chromosomes;
14. generate a random number r for each evolvable gene of each chromosome;
15. if $r < P_m[x]$, apply mutation to the evolvable gene of each selected chromosome;
16. end
17. end
18. exchange the fittest chromosomes among high-level and low-level populations;
19. compute the fitness of each chromosome in each population $P[x]^{gen+1}$ for generation $gen+1$ according to the fitness function;
20. $gen := gen + 1$;
21. end
22. for $x = 1$ to NP
23. return the fittest chromosome $c_{out}[x]$ from population x;
24. end
25. exit

Figure A3. The HCGA Algorithm

Figure A2 shows the gene encoding of individuals in both Level I and Level II populations. Basically, a Level II chromosome carries the genes representing various application parameters, such as the scoring formula for each due diligence dimension, the time window of due diligence, the weights of various financial metrics, the weights of various sociocultural factors, the weights of various factors related to business fitness, and the weights of various factors related to competitive advantage of an M&A target. All together, there are 36 evolvable genes of each Level II chromosome. A Level I chromosome carries the genes related to various system parameters and low-level NLP features (e.g., a specific generic sentiment lexicon to be used) controlling the sentiment analysis process or the business relation mining process. There are 9 evolvable genes of the Level I chromosome representing the system parameters controlling sentiment analysis, and 12 that are encoding the parameters for business relation mining. Decimal gene encoding is used for both Level I and Level II chromosomes in our system (Goldberg 1989; Lau et al. 2006).

Figure A3 shows the computational details of the HCGA algorithm. At the beginning of a coevolutionary process (i.e., the first generation), the chromosomes at each level are initialized by randomly assigning feasible values to each evolvable gene. Then, the fitness of each individual (chromosome) in a population is evaluated starting from level II. During fitness evaluation, the fittest individuals among the populations are exchanged according to the interaction pattern depicted in Figure A1. For the first generation, a randomly chosen individual from each Level I population is passed to the Level II population for fitness computation because the fitness of each individual of a Level I population has not yet been determined. For the subsequent generations, only the fittest individuals are exchanged among the populations. For each population, standard genetic operators such as selection, crossover, and mutation are applied to reproduce individuals of the next generation (Goldberg 1989; Huang et al. 2009; Lau et al. 2006). Moreover, the elitism rate w_e (i.e., the elitism factor) is applied to directly transfer a certain percentage of the fittest chromosomes from the current generation to the next generation P^{i+1} in order to retain the fittest chromosomes that represent good solutions for a problem domain (Goldberg 1989; Lau et al. 2006).

Roulette wheel selection (Goldberg 1989; Huang et al. 2009; Maiti and Maiti 2008) is applied to choose relatively fitter chromosomes from the current generation to produce individuals of the next generation. This type of selection is analogous to a roulette wheel, where each individual occupies an area on the wheel. The larger area the individual occupies, the more likely the ball will land there (i.e., the individual will be chosen). To tie the fitness of a chromosome to its probability of being chosen for reproduction of the next generation, a probabilistic selection function is defined according to the following:

$$\Pr(c) = \frac{fitness(c)}{\sum_{i=1}^{P_{size}} fitness(c_i)}$$

where c is a chromosome under consideration and is the fitness function defined earlier. P_{size} is the predefined size of a population P. To implement roulette wheel selection, a random number r in the unit interval is generated for each chromosome under consideration. If the selection probability $\Pr(c)$ of a chromosome c is greater than the random number (i.e., $\Pr(c) > r$), the corresponding chromosome is selected for reproduction. In other words, a fitter chromosome has a better chance to be selected for reproduction. However, chromosomes with low fitness values may also have a chance to be selected to maintain a good balance between exploitation- and exploration-oriented search.

After two chromosomes are selected for re-production, the genetic operation of two-point crossover (Goldberg 1989; Lau et al. 2006; Ruiz-Torrubiano 2010) is applied according to a predefined crossover probability p_c. Specifically, a random number r in the unit interval is generated for the pair of chromosomes under consideration. If $r < p_c$ is true, a two-point crossover is applied to the selected pair of chromosomes. Basically, two points along the evolvable genes of the pair are randomly selected. Then, the genes within the two-point boundary are exchanged between the two parent chromosomes to produce two child chromosomes. If $r < p_c$ is not established, the crossover operation will not be applied to the pair.

Each chromosome of the selected pair is then considered for the mutation operation after the crossover operation. First, a random number r in the unit interval is generated for each evolvable gene of each chromosome of the selected pair. If $r < p_m$ is true, where p_m is a predefined mutation rate, a mutation operation will be applied to the particular gene. With decimal gene encoding, the current value of the selected gene is replaced by another feasible gene value randomly. The evolutionary process (i.e., selection, crossover, and mutation) is repeated until the number of individuals of the next generation reaches the predefined number P_{size}. The aforementioned evolutionary process is applied to each population from high-level to low-level. If the average fitness of each population reaches a predefined threshold AVG_{fit}, or the number of generations reproduced exceeds the maximum number of generations MAX_{gen}, the HCGA algorithm will be terminated. At this stage, the fittest chromosome from each population is chosen to drive the operation of the M&A target scoring module of the ABIMA system. Table A1 lists the genetic parameters of the HCGA algorithm; these parameters were applied to the experiments reported in this paper.

Table A1. Parameter Settings of the HCGA Algorithm

HCGA Parameters	Level II Population (Application Parameter)	Level I Population (Sentiment Analysis Parameter)	Level I Population (Relationship Mining Parameter)
Size of population	$P_{size}[1] = 90$	$P_{size}[2] = 40$	$P_{size}[3] = 40$
Gene encoding	decimal	decimal	decimal
Elitism rate	$w_e[1] = 10\%$	$w_e[2] = 10\%$	$w_e[3] = 10\%$
Crossover probability	$P_c[1] = 0.83$	$P_c[2] = 0.83$	$P_c[3] = 0.83$
Mutation probability	$P_m[1] = 0.05$	$P_m[2] = 0.05$	$P_m[3] = 0.05$
Type of crossover	two-point	two-point	two-point
Type of mutation	uniform	uniform	uniform
Max number of generations	$Max_{gen} = 500$		
Max average fitness	$Avg_{fit} = 0.95$		

References

Delgado, M., Zuben, F., and Gomide, F. 2004. "Coevolutionary Genetic Fuzzy Systems: A Hierarchical Collaborative Approach," *Fuzzy Sets and Systems* (141:1), pp. 89-106.

Goldberg D. 1989. *Genetic Algorithms in Search, Optimization and Machine Learning*, Reading, MA: Addison-Wesley.

Huang, H., Pasquier, M., and Quek, C. 2009. "Financial Market Trading System with a Hierarchical Coevolutionary Fuzzy Predictive Model," *IEEE Transactions on Evolutionary Computation* (13:1), pp. 56-70.

Lau, R. Y. K., Tang, M., Wong, O., Milliner, S., and Chen, Y. 2006. "An Evolutionary Learning Approach for Adaptive Negotiation Agents," *International Journal of Intelligent Systems* (21:1), pp. 41-72.

Maiti, A. K., and Maiti, M. 2008. "Discounted Multi-Item Inventory Model Via Genetic Algorithm with Roulette Wheel Selection, Arithmetic Crossover and Uniform mutation in Constraints Bounded Domains," *International Journal of Computer Mathematics* (85:9), pp. 1341-1353.

Olsson, B. 2001. "Co-evolutionary Search in Asymmetric Spaces," *Information Sciences* (133:3-4), pp. 103-125.

Ounis, I., Macdonald, C., and Soboroff, I. 2008. "Overview of the TREC-2008 Blog Track," in *Proceedings of the 17th Text Retrieval Conference*, Gaithersburg, MD (http://trec.nist.gov/pubs/trec17/).

Ruiz-Torrubiano, R., and Suarez, A. 2010. "Hybrid Approaches and Dimensionality Reduction for Portfolio Selection with Cardinality Constraints," *IEEE Computational Intelligence Magazine* (5:2), pp. 92-107.

Appendix B

Affect Analysis for Sociocultural Fitness Assessment of M&A Targets

Affect analysis is very useful for modeling financial phenomena from both theoretical and pragmatic perspectives (Bollen et al. 2011; Deresky 2011). Recently, affect analysis has been successfully applied to predict the movement of the Dow Jones Industrial Average (Bollen et al. 2011). The main function of our affect analysis module is to estimate the sociocultural fitness of the targeted M&A companies, the targeted industrial sectors, or the entire targeted nation. For instance, affect analysis can be applied to assess the public's feelings (e.g., happiness or fear) about a potential M&A deal after it is announced. Specifically, the WordNet-affect lexicon (Valitutti et al. 2004) is applied to build our affect analysis module. From among the big six classes of affect (i.e., anger, fear, happiness, sadness, surprise, and neutral) that are often applied to affect analysis (Calix et al. 2010), four of them—anger, fear, happiness, and sadness—are used to estimate the emotion score of a potential M&A deal. The affect classes of surprise and neutral are not used by our affect analysis module because our preliminary experiments show that these classes cannot improve the performance of M&A affect analysis.

Each token of a financial document (e.g., a financial news article or an investor comment about a potential M&A deal) is matched against the WordNet-affect lexicon to identify its emotion class. Then, the emotion score of the document is computed according to the following:

$$emotion(d) = \frac{|happy| - (|anger| + |fear| + |sad|)}{|happy| + |anger| + |fear| + |sad|}$$

where *happy*, *anger*, *fear*, and *sad* represent the sets of emotional indicators extracted from a financial document *d*, which covers the potential M&A deal. With respect to the predefined time window of due diligence (i.e., an input parameter of an M&A query), each emotion score is then weighted using an exponential decay function (Barari and Mitra 2008; Jo et al. 1997). In particular, we apply the following exponential decay function to weight the affect scores:

$$emotion(d,t) = emotion(d) \times e^{-\left(\frac{t_{current} - t}{\tau/2}\right)}$$

where emotion(d) the term is the original affect score of a document d (i.e., without weighting), and $emotion(d, t)$ is the weighted emotion score at time point t. The term τ is the due diligence time window specified in months, and the term $(t_{current} - t)$ is the elapsed time (in months) between the time t when a financial document containing affects is posted and the time $t_{current}$ when M&A target scoring is conducted. For an emotion score derived from a document posted in the same month when M&A target scoring is conducted, the elapsed time $(t_{current} - t)$ is zero.

The weighted emotion score of a potential M&A deal *emotin*(*deal*) is the mean of the weighted emotion scores of the set of relevant financial documents D containing affects about a deal over each time point of the due diligence window, and it is defined by

$$emotion(deal) = \frac{\sum_{i=0}^{\tau-1} \sum_{j=1}^{|D|} emotion(d_j, t_i)}{\sum_{i=0}^{\tau-1} \left| D_{t_i} \right|}$$

where $emotion(d_j, t_i)$ represents the emotion score of a document d_j at time point t_i. The term D_i refers to the set of relevant documents at each time point t_i. Finally, the sociocultural fitness of an M&A target is estimated by taking into account the weighted emotion score of the potential M&A deal and other sociocultural factors.

References

Barari, M., and Mitra, S. 2008. "Power Law Versus Exponential Law in Characterizing Stock Market Returns," *Atlantic Economic Journal* (36:3), pp. 377-379.

Bollen, J., Maoa, H., and Zengb, X. 2011. "Twitter Mood Predicts the Stock Market," *Journal of Computational Science* (2:1), pp. 1-8.

Calix, R. A., Mallepudi, S. A., Chen, B., and Knapp, G. M. 2010. "Emotion Recognition in Text for 3-D Facial Expression Rendering," *IEEE Transactions on Multimedia* (12:6), pp. 544-551.

Deresky, H. 2011. *International Management: Managing Across Borders and Cultures, Text and Cases* (7th ed.), New York: Prentice Hall.

Jo, H., Han, I, and Lee, H. 1997. "Bankruptcy Prediction Using Case-Based Reasoning, Neural Networks, and Discriminant Analysis," *Expert Systems with Applications* (13:2), pp. 97-108.

Valitutti, A., Strapparava, C., and Stock, O. 2004. "Developing Affective Lexical Resources," *Psychology* (2:1), pp. 61-83.