The Stock Sonar — Sentiment Analysis of
Stocks Based on a Hybrid Approach

[1,2]**Ronen Feldman**, [2]**Benjamin Rosenfeld**, [2]**Roy Bar-Haim** and [2]**Moshe Fresko**

[1]School of Business Administration, The Hebrew University of Jerusalem, Jerusalem, ISRAEL
[2]Digital Trowel, Airport City, ISRAEL
ronen.feldman@huji.ac.il,{grur,roy,moshe}@digitaltrowel.com

Abstract

The Stock Sonar (TSS) is a stock sentiment analysis application based on a novel hybrid approach. While previous work focused on document level sentiment classification, or extracted only generic sentiment at the phrase level, TSS integrates sentiment dictionaries, phrase-level compositional patterns, and predicate-level semantic events. TSS generates precise in-text sentiment tagging as well as sentiment-oriented event summaries for a given stock, which are also aggregated into sentiment scores. Hence, TSS allows investors to get the essence of thousands of articles every day and may help them to make timely, informed trading decisions. The extracted sentiment is also shown to improve the accuracy of an existing document-level sentiment classifier.

1 Introduction

The Efficient Market Hypothesis states that stock prices already reflect all known information, and are instantly adjusted in response to new information. However, various studies, e.g. (Chan 2003; Tetlock, Saar-Tsechansky, and Macskassy 2008) have found that stock prices under-react to news stories and company events, and therefore investors may be able to achieve abnormal returns by following signals conveyed in news articles. This coincides with the common practice of many investors, who regard financial news and blogs as a major source of information for their trading.

These sources include factual information about companies, such as positive and negative business events (deals, lawsuits, new products etc.), as well as more subjective information such as analyses, opinions, speculations and rumors. These various types of information collectively determine whether an article is positive or negative with respect to a given stock. Notably, this view of stock sentiment extends the common view in sentiment analysis literature, by which sentiment is associated only with subjective utterances (Pang and Lee 2008).

The volume of articles published daily makes automatic analysis of financial content a much needed tool for investment decision making. Most previous work on stock sentiment analysis focused on document-level sentiment classification (as we discuss in more detail in Section 5). By con-

trast, TSS provides precise *sentiment extraction*: highlighting of positive and negative expressions within the article text, as well as extraction of positive and negative business events. Extracted sentiment provides the user an *explanation* for the article score as well as an effective *summary* of multiple news articles. As we show in this paper, the sentiment extracted by TSS can also be used to improve document-level sentiment classifiers.

Another limitation of previous methods is concerned with the level of linguistic analysis required to correctly predict sentiment. Current systems usually employ sentiment lexicons and machine-learning algorithms that operate at the word or phrase level. Such methods typically fail to model compositional expressions, e.g. correctly classifying *"reducing losses"* as positive, but *"reducing forecasts"* as negative. Furthermore, it is often necessary to go beyond the phrase level to obtain the correct sentiment. Consider the following sentence:

Toyota announces voluntary recall of their highly successful top selling 2010 model-year cars.

Phrase level sentiment would be misled by the expressions *"highly successful"* and *"top selling"*, and classify this sentence as positive, or, at best, as neutral, if these expressions are balanced by the negative expression *"voluntary recall"*. However, correct sentiment analysis should recognize the whole sentence as a negative business event (product recall). In many cases it is also necessary to correctly recognize the semantic role of the analyzed company in the event. For example, in the following sentence:

ArvinMeritor (ARM), a maker of integrated systems, rose after it won an antitrust suit against electrical power gear maker Eaton (ETN).

it is necessary to correctly identify the winner and the loser in the lawsuit for correct sentiment assignment.

In order to cope with the inherent linguistic complexity and allow high precision sentiment extraction, we have developed a hybrid sentiment analysis approach, which combines three linguistic components: (a) a wide-coverage sentiment lexicon (b) Patterns for modeling phrase-level compositional expressions, and (c) Semantic event extractor for business events. These three components are implemented in a novel framework which combines knowledge engineering with state-of-art machine learning.

2 Architecture

TSS collects thousands of articles from thousands of sources every day. The articles are collected via stock-specific RSS feeds, so that each article is associated with one or more stocks. The same news is often repeated by multiple sources. Currently we do not attempt to identify these duplicates, except for the obvious case where articles have the same title and date. However, we assume that the number of times a story is repeated is indicative for its significance, and therefore keeping these duplicates is beneficial.

The collected articles are first cleaned so that the main body of the article is maintained and the extraneous content (such as ads, links to other stories, etc.) is deleted. The module in charge of the extraction of the main textual content from the HTML pages is based on a supervised machine learning approach and a visual training module as described in (Rosenfeld, Feldman, and Ungar 2008). The output of this module is plain text. Each article is analyzed separately for each of its associated stocks. We shall refer to the stock for which the article is currently analyzed as the *main company*. Based on the extracted sentiment, an article score is computed, and article scores are then aggregated into daily scores for each ticker. The rest of this section details sentiment extraction and scoring.

2.1 The CARE Extraction Platform

TSS is based on a hybrid approach to sentiment analysis: it combines *dictionary*-based sentiment, compositional phrase-level *patterns*, and semantic *events*. These three components are implemented using our information extraction platform, termed *CARE (CRF Assisted Relation Extraction)*. Due to space and scope limitations, we give here only a brief overview of the platform. CARE is a hybrid Machine Learning/Knowledge Engineering-based system designed to extract complex relations from natural language text. It is a direct descendant of TEG (Rosenfeld et al. 2004), an extraction engine based on partial, relation specific, sentence parsing. CARE may be viewed as the discriminative version of (a generative) TEG in the same sense in which discriminative sequence classifiers are improved versions of HMM-based sequence classifiers.

CARE is based on Weighted Context Free Grammars (WCFG), interpretable and trainable as CRFs. It also features a flexible interface between the parsing component and the token classification components such as part-of-speech (POS) tagger and NER (Named Entity Recognizer). This interface allows the grammar to selectively modify the classification results and to adapt generically trained sequence classifiers to specific domains of relation extraction.

Each of the three TSS components is implemented as a rulebook (grammar written in the CARE language), and these rulebooks are fed together to CARE's parsing component. This has several advantageous consequences. The first is that the grammars representing each of the components are processed as a single weighted grammar, allowing the analyses generated by each component to compete for the best parse. Thus, the lexical items, compositional patterns and semantic events are all simultaneously considered when

Sentiments		
	Positive	**Negative**
Adjectives	attractive, superior	inefficient, risky
Verbs	invents, advancing	failed, lost
Nouns	opportunity, success	weakness, crisis
Multi-word expressions	exceeding expectations, falling into place	chapter 11,pull back
Neutral expressions	in the worst case, best practice	
Sentiments Modifiers		
Emphasis	Huge, incredible, highly	
De-emphasis	mostly, quite	
Reversal	far from, cut, no	

Table 1: Lexicon components

evaluating the sentiment of each token. Furthermore, because the three components operate within the CARE framework, they can utilize its CRF classifier, which flexibly connects them to the state of the art NER and POS taggers.

The TSS rulebook was developed by a team of three linguistic engineers, assisted by two financially-trained domain experts, over a period of five months. We now turn to a short description of its individual components.

2.2 Dictionary-Based Sentiment

Our dictionary-based component aims at wide coverage of generic sentiment. It relies on lists of positive and negative words, or expressions. We started out with the domain-specific McDonald's Lexicons[1], and the much larger general-purpose "Harvard Inquirer"[2] list. These lexicons were revised, filtered and extended by our domain experts. Our lexicon currently contains around 2,000 words and expressions. The main components of our lexicon are illustrated in Table 1. Sentiment modifiers are operators that change either the magnitude (emphasis or de-emphasis) or the polarity of the sentiments (reversal). For example, *"highly successful"* is more positive than *"successful"*, whereas *"mostly successful"* is less positive, and *"far from successful"* is negative. Sentiment modifiers are processed right to left, at post-processing time, to obtain the overall expression sentiment.

2.3 Patterns

Patterns capture phrase-level sentiments. Comparing to dictionary-based sentiment they are more structured, compositional, and much more specific. Our patterns model expressions that convey positive and negative financial information. For example, the following sentence

> **The shares fell 13%,** *despite* **a record quarter sales growth** *and* **better-than-expected earnings**.

contains a negative pattern, followed by two positive patterns. Let us consider the first expression. It can be generalized as *[financial parameter] [direction of change] [mag-*

---

[1]http://www.nd.edu/~mcdonald/Word_Lists.html
[2]http://www.wjh.harvard.edu/~inquirer/homecat.htm

Category	Example (polarity)
Legal	ArvinMeritor (ARM), a maker of integrated systems, rose after it won an antitrust suit against electrical power gear maker **Eaton (ETN)**. ⊖
Analyst Recommendation	On June 23, Caris & Co. reiterated its "buy" rating on **CRM** and increased its price target to $115 from $100. ⊕
Financial	**Western Digital** earnings climb 35% ⊕
Stock Price Change	**SandRidge Energy** fell 3.0 percent to $6.24. ⊖
Deals	The U.S. Army's Mission Installation Contracting Command has awarded **Northrop Grumman Corporation (NYSE:NOC)** a contract to provide logistics support at Fort Eustis, Va. ⊕
Mergers and Acquisitions	On Monday, Ramius LLC has offered to acquire all the outstanding shares of **CYPB** for $4.00 per share in cash. ⊕
Partnerships	Stennis has partnered with **Orbital Sciences Corporation** to test the AJ 26 engines that will power the first stage of the company's Taurus II space launch vehicle. ⊕
Product	**The fast-food giant** recalled 12 mil drinking glasses that contain cadmium ⊖
Employment	**Heidrick & Struggles** Appoints New Hedge Fund Leadership Team. ⊕

Table 2: Types of extracted events. Main company is marked in bold. ⊕/⊖ represent positive/negative polarity.

nitude], where "financial parameter" may also be sales, revenue, losses, and so on, and *"direction of change"* may be *"rose"*, *"decreased"* etc. The sentiment is obtained by the composition of these pieces. For instance, if the financial parameter is *"expenses"* or *"debts"*, a downward change is positive, while for parameters such as *"sales"* and the *"stock price"* it would be negative.

2.4 Events

Business events, such as launching a new product, (or, by contrast, product recall), signing a new deal, filing a lawsuit against the company, workforce change and so on, have strong positive or negative impact on the company and how it is perceived by the stock market, and are therefore crucial for our TSS system. Events are the most complex objects we aim to capture, and consumed most of the grammar development time. The event types currently addressed by the system are listed and illustrated in Table 2. All events must include the main company (either a direct mention or an anaphor). We aim to avoid catching historic events, e.g. *"Google acquired YouTube **in 2006**"*, or ***"In the same month**, Repower signed a deal with Everpower"*, mainly by considering the event's immediate context. Events get higher weights than other sentiment types, to ensure that part of the event is not recognized as a pattern or a lexical sentiment.

Figure 1 shows a much simplified fragment of the Product event rules. Rule weights are omitted, and there are many other syntactic simplifications for the sake of readability. These rules catch product events of the type *"Teva an-*

nounced today its new treatment for multiple sclerosis was found effective in clinical trials". It defines in (1) a positive product event. The primary "anchors" these rules utilize are *"treatment"* nouns (5), followed by a positive adjective (6).

2.5 Sentiment Relevance

When analyzing sentiment for a given company, it is crucial to assess that the sentiment indeed refers to that company. O'Hare et al. (2009) suggested to consider only a window of N words around each mention of the main company in the article, and showed that it improves polarity prediction as compared to considering the whole article for the company. Our experiments on a training corpus confirmed that identification of relevant sections is crucial to obtain reasonable precision, and that the distance from a mention of the main company is a good predictor of relevance. However, we found two additional cues for relevant that were not considered by O'Hare at el.:

1. *Directionality* - sentiments that appear after the main company are more likely to be relevant than sentiments preceding it.

2. *Other entities* - entities that appear between the main company mention and the sentiment are good indicators for irrelevance.

Eventually, we implemented a relevance strategy that considered only sentiments that appear after the main company (considering its closest mention) in the same sentence, with no other entities in between (this was approximated by considering any capitalized word).

2.6 Scoring

Each extracted sentiment was given either a positive or a negative score, according to its polarity and type. Event weights were set according to the significance of each event type, as estimated by our domain experts. For instance, analyst recommendations have the heaviest weight of 10. Different subtypes of the same event category may also get different weights. For example, a product upgrade gets a positive score of 2, whereas product recall gets a negative score of 10. In addition, sentiments extracted from the headline are given a higher weight than sentiments in the article body. Let P and N be the sum of positive and negative scores in an article, respectively. The *article score* is then defined as

$$S = \frac{P - N}{P + N + 3} \quad (1)$$

The resulting score is between -1 and 1, where positive (negative) score indicates positive (negative) sentiment polarity. Note that the denominator diminishes the sentiment magnitude in the case of mixed sentiment, where there are many positive and negative sentiments in the same article. The following scores are computed for each company every day:

1. *Positive/Negative impact* is simply the sum of positive/negative scores (P,N) for a company in the given day.

2. *Daily score* is the same as the article score, with P and N summing over all main company articles in a given date.

(1)	*PositiveProductEvent* → *CompanyAnnounced* Tokens(0,3) *TreatmentIsPositive*
(2)	*CompanyAnnounced*→ [*Time*] *MainCompany* Tokens(0,2) (announced|said) [*Time*]
(3)	*Time* → yesterday|today
(4)	*TreatmentIsPositive* → *TreatNoun* Tokens(0,5) [is|are|was|were] [found] *PositiveTreatmentAdjective*
(5)	*TreatNoun* → treatment|therapy
(6)	*PositiveTreatmentAdjective* → safe|robust|effective|efficient|positive|(commercially viable)|beneficial

Figure 1: An excerpt from the Product rulebook. [. . .] is optional.'|' specifies alternatives. Token(i,j) is any i to j tokens. Nonterminals are capitalized and *emphasized*.

3. *Composite score* is computed by summing over all the articles collected for the company, while weighting older articles with a decay factor.

3   Application

The Stock Sonar application allows registered users to view today's most positive and negative articles, to analyze sentiment for selected stocks, and to create and follow a portfolio. The basic TSS edition is deployed as a public website[3], where users can currently register at no fee. In addition, we offer premium content for institutional users through partnerships with leading commercial content providers such as Dow Jones. We also provide an API to the sentiment engine, which produces XML annotation for a given article and company name/ticker. The API is useful for algorithmic trading companies, who want to incorporate our sentiment signals into their trading algorithms.

The TSS application has three main sections:

Impact Feed: Shows today's most positive and negative articles. Allows filtering by index (S&P 500, Dow Jones Industrial Average etc.), by sector (Financial, Technology, Healthcare etc.), or by considering only stocks in the user's portfolio.

Analyzer: Provides several tools for analyzing the sentiment of a given stock. (a) *Sentiment Graphs:* show positive and negative impact, and daily and composite score over time, and compare them to the stock price chart. (b) *Articles:* lists articles containing positive and negative sentiment for the stock, along with the article's score. Clicking on the article's headline displays the article body with positive and negative sections highlighted. (c) *Events:* lists a summary of positive and negative events extracted for the stock. Can be filtered by the event type.

Portfolio: Allows users to easily follow selected stocks by saving them in a portfolio.

We illustrate a typical use of the TSS application through an example. Consider the impact graph produced by the TSS system for *Clinical Data Inc.* (NASDAQ:CLDA) between January 15th and January 27th, 2011 (Figure 2). We see that a prominent peak of positive impact on January 18 precedes a surge in the stock price, from $15.03 at the close of January 21 to $25.17 at the close of January 24.

By clicking on each peak the user can see the articles and events that contributed to that impact score. The main contributor to Jan. 18's impact score was an analyst recommendation event. The extracted event and its source article are

---

Figure 3: An extracted event, and its source article.

shown in Figure 3. In that article the analyst firm says about Clinical Data's antidepressant, Vilazodone: *"We remain optimistic that Vilazodone will receive FDA approval by Monday, January 24th"*. Spotting this event and article by TSS reveals an excellent opportunity for investors. Indeed, three days later, on January 21, FDA approved this drug, and a dramatic jump in the stock price followed.

4   Evaluation

In this section we report two experiments which evaluate the sentiment extracted by TSS. Section 4.1 shows that TSS extracts sentiment expressions with high precision. We then show in Section 4.2 that the sentiment extracted by TSS can also improve the prediction of document-level sentiment over current bag-of-words classifiers.

4.1   Precision of Sentiment Extraction

Our TSS system extracts substantial amount of sentiment: on a sample of 10,000 articles, it extracted sentiment from 3,808 of them. In order to assess the quality of extracted sentiment, we checked the sentiment polarity of phrases extracted from a random sample of articles. 200 extracted phrases from all sources (dictionaries, patterns, and events) were judged in total. In addition, the precision of a 100 randomly sampled extracted events was checked, for which we required the correctness of both sentiment and the event type. Judgments were carried out by one of our domain experts. The results are presented in Table 3. The results show that overall, sentiments are extracted with good precision
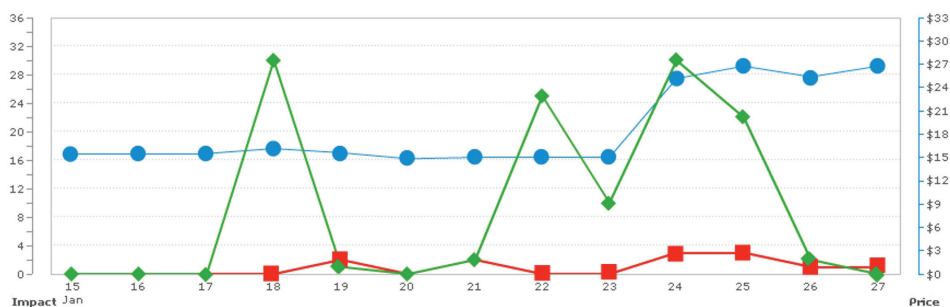
Figure 2: Sentiment graph for Clinical Data Inc., Jan. 15-27, 2011. Positive impact (◆), negative impact (■), and stock price (●) are shown.

	Total	Correct	Precision
Sentiments	200	165	82.5%
Events	100	95	95.0%

Table 3: Precision of extracted sentiment and events

and in particular, the precision of events is very high.

4.2 Document-Level Polarity Classification

Next, we evaluate sentiment polarity classification of whole articles by TSS, and compare it to a previous approach, which relies solely on machine learning. TSS determines article polarity according to the sign of the article score, as defined in Section 2.6. An article is classified as neutral if either its score is zero, or no sentiment is extracted for it. As a baseline, we implemented the classifier described in (O'Hare et al. 2009), as the problem they address seems closest to ours. O'Hare et al. classify article polarity with respect to a given company, using a bag-of-words classifier. As we described in Section 2.5, they extract features only from a windows of N words around each company mention (in their experiments they found that a window of N words was better than N sentences or N paragraphs).

For this experiment we used a corpus of *(company, article)* pairs. That is, for each article it was determined which company is the main company for sentiment analysis. We considered only articles where a mention of the main company was found by our system. Each article was labeled with one of the following categories: {*highly negative, negative, neutral, positive, highly positive*}. For our evaluation, we ignored the distinction between *highly-positive* and *positive* (and *highly negative/negative*), ending up with a three-class labeling. The corpus was tagged jointly by our domain experts and a team of last-year undergraduate students in financial and business-related disciplines. The corpus contains 2,132 articles, split into a training set of 1,358 articles, and a test set of 774 articles. The optimal window size for the classifier was found to be 25 words, based on cross-validation results on the training set. Using this windows size, an SVM classifier was trained on the training set.

The two leftmost columns in Table 4 compare our TSS results with the classifier results on the test corpus. Looking at the various metrics (precision, recall, F-measure) of

		TSS	Classifier	Combined
Positive	R	46.71	**90.63**	88.82
	P	**75.36**	57.36	61.38
	F1	57.67	70.26	**72.59**
Negative	R	45.79	37.38	**51.40**
	P	60.49	59.70	**61.80**
	F1	52.13	45.98	**56.12**
Neutral	R	**75.00**	24.69	31.28
	P	46.10	65.93	**67.26**
	F1	**57.10**	35.93	42.70
Accuracy		56.64	58.74	**62.41**

Table 4: Three-way article polarity classification. Precision (P), recall (R), and F-measure (F1) are reported. Best results are shown in bold.

the different classes, neither of the methods is consistently better than the other. However, the classifier outperforms TSS on overall accuracy (percentage of correct classifications), which was the metric O'Hare et al. used for evaluation. Most misclassifications confuse positive or negative with neutral, while only a small fraction of them confuses positive with negative. Indeed, TSS and the classifier seem complementary: TSS extract precise sentiment and can handle more complex expressions, but does not take advantage of the training data to adjust its scoring function, and lacks the broad coverage that a bag-of-words classifier has. Thus, it makes sense to combine these two approaches.

We combined TSS and the classifier by adding the positive and negative counts for events, patterns and dictionaries (six numbers in total) as additional features for the classifier, after applying discretization of the scores (e.g. a positive score of 3 was translated into three binary features: $p1$, $p2$, $p3$).

The rightmost column in Table 4 shows the combined results. We can see that it improves the classifier results in almost every metric, and the overall accuracy is better than both the classifier and TSS alone. The most prominent improvement was for negative articles. This illustrates the added value of our knowledge-based hybrid approach for improving sentiment classification.

5   Related Work

Previous academic research focused on document-level sentiment classification of financial news, blogs and stock message boards, as well as on stock prediction from these sources, either via sentiment classification or directly (Lavrenko et al. 2000; Das and Chen 2001; Koppel and Shtrimberg 2004; Devitt and Ahmad 2007; O'Hare et al. 2009; Schumaker and Chen 2010). These works rely on pre-defined sentiment lexicons, learning from manually classified training texts (or using stock prices as labels), or some combination of the two.

Moving on to commercial sentiment analysis tools, *RavenPack* also classifies company sentiment at the document level, using sentiment lexicons and learning from labeled documents (Mitra, Mitra, and diBartolomeo 2008). *Lexalytics* (2010) and SAS (2011) go beyond document classification and highlight sentiment within the document. They aim to match the extracted phrases to the entities or concepts being analyzed. Lexalytics' engine automatically learns positive and negative phrases by considering expressions corresponding to certain POS patterns (e.g. adjective-noun), and learning their polarity from co-occurrence with known positive and negative expressions, using web statistics. SAS Sentiment Analysis allows writing of linguistic rules for one or several matches of a term, regular expressions and POS tags, along with boolean operators expressing constraints such as the distance and occurrence of concepts in relation to other words.

TSS differs from these approaches in the following respects. Most previous work did not aim at in-text sentiment tagging. Tools that do provide this capability (Lexalyitcs, SAS) are limited to generic sentiment. By contrast, the CARE platform allows us to model and extract much more complex semantic objects such as compositional patterns and business events, using the expressive power of WCFGs. Events and patterns complement generic sentiment with factual information, which is crucial for inferring the article's polarity, and provide the user with an effective news summary.

The SONAR surveillance application (Goldberg et al. 2003) also extracts events from news articles, which partially overlap with TSS event types, for the purpose of detecting insider trading and fraud in stock markets. Extraction rules were written in the DIAL language (Feldman and Sanger 2006), an early ancestor of CARE. SONAR differs considerably from TSS in terms of intended use, extracted information and its presentation to the user.

6   Conclusion

This paper introduced *The Stock Sonar*, a novel stock sentiment analysis application. TSS combines the common practice of dictionary-based sentiment analysis with more complex linguistic analysis that captures compositional patterns and semantic events. Events and patterns enrich stock sentiment analysis with crucial factual information that was missing in previous work, and provide a useful news summary. The CARE knowledge engineering platform allows modeling of language subtleties and complex linguistic struc-

tures, which are currently beyond the reach of pure machine-learning methods. TSS goes beyond article-level sentiment classification and provides precise in-text highlighting of positive and negative expressions. In particular, it was shown to extract business events with very high precision.

In future work we plan to investigate the predictive power of TSS for stock prices, and look for better ways to incorporate TSS tagging with classification-based methods.

References

Chan, W. S. 2003. Stock price reaction to news and no-news: drift and reversal after headlines. *Journal of Financial Economics* 70(2):223 – 260.

Das, S., and Chen, M. 2001. Yahoo! for Amazon: Extracting market sentiment from stock message boards. In *Proceedings of APFA*.

Devitt, A., and Ahmad, K. 2007. Sentiment polarity identification in financial news: A cohesion-based approach. In *Proceedings of ACL*.

Feldman, R., and Sanger, J. 2006. *The Text Mining Handbook: Advanced Approaches in Analyzing Unstructured Data*. Cambridge, MA, USA: Cambridge University Press.

Goldberg, H. G.; Kirkland, J. D.; Lee, D.; Shyr, P.; and Thakker, D. 2003. The NASD Securities Observation, New Analysis and Regulation system (SONAR). In *IAAI*, 11–18.

Koppel, M., and Shtrimberg, I. 2004. Good news or bad news? Let the market decide. In *Proceedings of the AAAI Spring Symposium on Exploring Attitude and Affect in Text: Theories and Applications*.

Lavrenko, V.; Schmill, M.; Lawrie, D.; Ogilvie, P.; Jensen, D.; and Allan, J. 2000. Mining of concurrent text and time series. In *Proceedings of KDD-2000 Workshop on Text Mining*.

Lexalytics. 2010. How we do sentiment analysis. http://www.lexalytics.com/resources/how-we-do-sentiment-analysis-page1.

Mitra, L.; Mitra, G.; and diBartolomeo, D. 2008. Equity portfolio risk (volatility) estimation using market information and sentiment. http://www.ravenpack.com/research/Equity_portfolio_risk_est.pdf.

O'Hare, N.; Davy, M.; Bermingham, A.; Ferguson, P.; Sheridan, P.; Gurrin, C.; and Smeaton, A. F. 2009. Topic-dependent sentiment analysis of financial blogs. In *Proceedings of TSA'09*.

Pang, B., and Lee, L. 2008. Opinion mining and sentiment analysis. *Found. Trends Inf. Retr.* 2.

Rosenfeld, B.; Feldman, R.; Fresko, M.; Schler, J.; and Aumann, Y. 2004. TEG: a hybrid approach to information extraction. In *Proceedings of CIKM*.

Rosenfeld, B.; Feldman, R.; and Ungar, L. 2008. Using sequence classification for filtering web pages. In *Proceeding of CIKM*.

SAS. 2011. SAS sentiment analysis. http://www.sas.com/resources/factsheet/104357_0610.pdf.

Schumaker, R. P., and Chen, H. 2010. A discrete stock price prediction engine based on financial news. *Computer* 43:51–56.

Tetlock, P. C.; Saar-Tsechansky, M.; and Macskassy, S. 2008. More than words: Quantifying language to measure firms' fundamentals. *The Journal of Finance* 63(3):1437–1467.