NORTHERN ARIZONA UNIVERSITY
The W. A. Franke College of Business

Evaluating Sentiment in Financial News Articles

Working Paper Series—11-10 | October 2011

Robert P. Schumaker[1], Yulei Zhang[2], Chun-Neng Huang[3] and Hsinchun Chen[3]

[1] Computer and Information Systems Dept
Cleveland State University, Cleveland, Ohio 44115, USA
rob.schumaker@gmail.com

[2] Computer Information Systems, The W. A. Franke College of Business,
Northern Arizona University, Flagstaff, Arizona 86011, USA
yulei.zhang@nau.edu

[3] Artificial Intelligence Lab, Department of Management Information Systems
The University of Arizona, Tucson, Arizona 85721, USA
{cnhuang, hchen}@eller.arizona.edu

Evaluating Sentiment in Financial News Articles

1. Introduction

The ability to predict stock market behavior has always had a certain appeal to researchers. Even though numerous attempts have been made, the difficulty has been the inability to capitalize on the behaviors of human traders. Behavioral patterns have not been fully defined and are constantly changing; thus making accurate predictions quite difficult. To further add to the uncertainty, there are two entirely opposed philosophies of stock market research; fundamental and technical analysis techniques (Technical Analysis, 2005). Fundamentalists seek to leverage a security's relative data, ratios and earnings, while technicians analyze charts and modeling techniques based on historical trading volume and pricing. The basic problem coalesceses into *does price history matter?*

With the advent of cheap computing and the ease of gathering information, the role of computers in stock prediction has increased dramatically. These systems have simply followed the trend of automating existing fundamental and/or technical strategies. Their goal is to achieve better returns than human traders by removing the elements of emotion and bias from trading (Jelveh, 2006). The negative aspect of these systems is that they lack intuition and will continue to buy even after unfavorable news events, such as losing a costly court battle. In order to work effectively, these systems require news events to be translated into numeric data before appropriate decisions can be made. This problem introduces serious lag-time into decisions and in some cases human analysts must override trades.

The motivation of this paper is to build and test a financial news article system that incorporates sentiment analysis techniques in its predictive arsenal. By adjusting the variables of article sentiment and tone of the article, we seek to find an optimal trading system.

This paper is arranged as follows. Section 2 provides an overview of literature concerning Stock Market prediction, textual representations and sentiment analysis techniques. Sections 3 and 4 describe our proposed approaches and the AZFinText system respectively. Section 5 provides an overview of our experimental design. Section 6 details our experimental findings and discusses their impact on stock market prediction. Section 7 delivers our conclusions and a brief discourse on future research directions.

2. Literature Review

As mentioned earlier, there have been two theories with a significant impact on predicting security prices, Efficient Market Hypothesis (EMH) and Random Walk Theory. In EMH, the price of a security is a reflection of complete market information and when a change occurs, the market instantly adjusts the price of the security to reflect this new information (Fama, 1964). EMH can vary the amount of information contained to encompass three distinct levels; the weak form, the semi-strong and the strong form. In weak EMH, only historical data is embedded within the current price. Semi-Strong EMH incorporates both historical and current public information into its prices. Strong EMH includes all pertinent information such as history, current public information and private information, such as Insider Trading. From EMH theory, it is the belief that markets behave efficiently and instantaneous price corrections make any prediction model useless.

Malkiel's Random Walk Theory is similar to Semi-Strong EMH where all information is contained within the current price and is worthless for future prediction. This theory slightly differs from EMH by maintaining that short-term price movements are indistinguishable from random activities (Malkiel, 1973). This short-term random activity thus produces unpredictable prices and makes it impossible to consistently outperform the market.

The ability to scrutinize the decisions of traders and uncover the micro-effects of trading behavior on the scale of a market exchange is extremely difficult. To lessen this difficulty and simultaneously test the impact of fundamental and technical trading strategies, LeBaron created an artificial stock market with simulated traders whose decisions can be dissected (LeBaron et al., 1999). LeBaron accomplished this by introducing new pieces of information into the market and then adjusting the amount of time between when an individual trader would receive information and act upon it. He discovered that traders with longer waiting times would form fundamental strategies (relying more heavily on company-specific performance data) while those with shorter waiting times developed technical strategies (such as timing trades). This study led to a more important contribution by uncovering a lag between the time that information is introduced to when the market would return to equilibrium. This delay in market behavior helped to dismiss the instantaneous correction tenets of EMH and lent support to the idea that markets could be forecast for short durations following the introduction of new information. Follow-up research into the limits of this predictive window led to the discovery of a twenty minute window of opportunity before and after a financial news article is released (Gidofalvi, 2001). Within this window, weak prediction of the direction of a stock price is possible.

2.1 Financial News Articles

New information is introduced into the market all the time. While a variety of information sources can all move a stock price, e.g., rumors, eavesdropping and scandals; financial news articles are considered more stable and a more trustworthy source. This stability has prompted some to declare news to be another form of commodity (Mowshowitz, 1992) that can have differing values (Raban & Rafaeli, 2006). However, the exact relationship between financial news articles and stock price movement is complex. Even when the information contained in financial news articles can have a visible impact on a security's price (Gidofalvi, 2001; Lavrenko et al., 2000a; Mittermayer, 2004; Wuthrich et al., 1998), sudden price movements can still occur from other sources, such as large unexpected trades (Camerer & Weigelt, 1991).

The first challenge of a textual financial prediction system is to manage the large amounts of textual information that exist for securities. This material can include required reports such as periodic SEC filings, press releases and financial news articles reporting both unexpected events and routine news alike. These textual documents can then be parsed using Natural Language Processing (NLP) techniques to identify specific article terms or phrases most likely to cause dramatic share price changes, such as "factory exploded" would probably indicate a price plunge in the near future. By automating this process, machines can take advantage of arbitrage opportunities faster than human counterparts by repeatedly forecasting price fluctuations and executing immediate trades.

Obtaining timely financial documents from reputable Web sources is a critical step and there are a variety of financial news aggregation sites that provide this service. One of these sites is Comtex which offers real-time financial news in a subscription format. Another source is PRNewsWire, which offers free real-time and subscription-based services. Yahoo! Finance is a third such source and is a compilation of 45 different news sources including the Associated Press, Financial Times and PRNewsWire among others. This source provides a variety of perspectives and timely news stories regarding financial markets.

2.2 Textual Representation

Once the financial news articles have been gathered, we need to represent their important features in machine-friendly form. One technique is a Bag of Words approach which has been extensively used in textual financial research (Gidofalvi, 2001; Lavrenko et al., 2000a). This process involves removing the meaningless stopwords such as conjunctions and declaratives from the text and using what remains as the textual representation. While the Bag of Words method has been popular, it suffers from noise issues associated with seldom-used terms and problems of scalability, where immense computational power is required for large datasets. An improved representational system is Noun Phrases. This representation retains only the nouns and noun phrases from a document and can adequately represent the important article concepts (Tolle & Chen, 2000). As a result, this technique uses fewer terms and can handle article scaling better than Bag of Words. A third representational technique is Named Entities, which is an extension of Noun Phrases. It functions by selecting the proper nouns of an article that fall within well-defined categories. This process uses a semantic lexical hierarchy (Sekine & Nobata, 2004) as well as a syntactic/semantic tagging process (McDonald et al., 2005) to assign candidate terms to categories. The exact categorical definitions are described in the Message Understanding Conference (MUC-7)

Information Retrieval task and encompass the entities of date, location, money, organization, percentage, person and time. Named Entities allows for better generalization of previously unseen terms and does not possess the scalability problems associated with a semantics-only approach. A fourth representational technique is Proper Nouns. This method functions as an intermediary between Noun Phrases and Named Entities where it exists as a subset of Noun Phrases by selecting specific nouns and also as a superset of Named Entities without the constraint of pre-defined categories. This representation removes the ambiguity associated with proper nouns that could be represented by more than one named entity category or fall outside one of the seven defined categories. In a comparison study using these four representational techniques, it was found that the Proper Noun representation was much more effective in representing textual financial news articles (Schumaker & Chen, 2006).

Assigning a representational mechanism is not sufficient to address scalability issues associated with large datasets. A common solution is to introduce a term frequency threshold that uses a term frequency cut-off to represent article terms that appear more frequently (Joachims, 1998). This technique not only eliminates noise from lesser used terms, but also reduces the number of features to represent. Once scalability issues have been addressed, the data needs to be prepared in a more machine-friendly manner. Machine learning algorithms are unable to process raw article terms and require an additional layer of representation. One popular method is to represent article terms in binary where the term is either present or not in a given article (Joachims, 1998). This solution leads to large but sparse matrices where the number of represented terms throughout the dataset will greatly outnumber the terms used in an individual article.

Once these financial news articles have been represented, learning algorithms can then begin to identify patterns of predictable behavior. One accepted method, Support Vector Regression (SVR), is a regression equivalent of Support Vector Machines (SVM) but without the aspect of classification (Vapnik, 1995). Like SVM, SVR attempts to minimize its fitting error while maximizing its goal function by fitting a regression estimate through a multi-dimensional hyperplane. This method is also well-suited to handling textual input as binary representations and has been used in similar financial news studies (Schumaker & Chen, 2006; Tay & Cao, 2001).

2.3 Sentiment Analysis

In general, sentiment analysis is concerned with the analysis of direction-based text, i.e., text containing opinions and emotions (Abbasi et al., 2008). Sentiment classification studies attempt to determine whether text is objective or subjective, as well as whether the subjective parts contain either positive or negative sentiments. This classification into positive and negative sentiments is a common two-class problem (Pang et al., 2002; Turney, 2002). Additional variations include classifying sentiments as opinionated/subjective or factual/objective (Wiebe et al., 2001; Wiebe et al., 2004). Some studies have attempted to classify emotions, including happiness, sadness, anger, horror, etc., instead of sentiments (Grefenstette et al., 2004; Mishne, 2005; Subasic & Huettner, 2001).

Recently, the relationship between sentiment analysis and stock market movement, or online product sales, has been examined in several studies. Some of them focus on a formal channel, such as financial press releases and financial news (Davis et al., 2006; Devitt & Ahmad, 2007; Tetlock, 2007). While others focus on user-generated content (Antweiler & Frank, 2004; Das & Chen, 2007; Ghose et al., 2007). Tetlock measured the interactions between the media and the stock market using daily content from the Wall Street Journal. The results show that negative sentiment from these boards may be predictive of future downward moves in firm value (Tetlock, 2007). Devitt and Ahmad explore positive and negative polarity in financial news and found it to be consistent with human judgment (Devitt & Ahmad, 2007). Davis et al. investigated the effects of optimistic and pessimistic language used in press releases on future firm performance (Davis et al., 2006). Antweiler and Frank examined the bullishness of messages and found that chat forums can predict market volatility (Antweiler & Frank, 2004). Das and Chen extracted investor sentiment from stock message boards and found evidence of a relationship with stock prices (Das & Chen, 2007). Ghose et al. found that user feedback had a positive relationship with sales (Ghose et al., 2007).

One such tool to measure article tone and polarity is OpinionFinder. OpinionFinder aims to identify subjective sentences and to mark the various aspects of the subjectivity in these sentences, including positive or negative sentiments (Wilson et al., 2005a). This tool is well known and tested. It was developed by Wiebe's group based on a series of publications, such as the subjective sentence classifier (Riloff & Wiebe, 2003; Wiebe & Riloff, 2005), and the polarity classifier (Wilson et al., 2005b). It also has relatively good performance. When compared against the MPQA Opinion Corpus, OpinionFinder has an accuracy of 74%, subjective precision of 78.4%, subjective recall of 73.2% and a subjective F-measure of 75.7% as compared to baseline accuracy of 55.3%.

3. Research Questions

From these gaps, we have formulated several research questions. The first of which is:

- Does Objectivity/Subjectivity impact news article prediction?

While we know from the literature that the tone of an article can influence future price, can the addition of measures of financial article objectivity/subjectivity provide AZFinText with improved prediction capabilities.

As a follow-up to this question, we also ask:

- Does Positive/Negative Subjectivity impact news article prediction?

4. System Design

In order to evaluate our research questions, we designed the AZFinText system. Figure 1 illustrates the AZFinText system design.

Figure 1. The AZFinText System

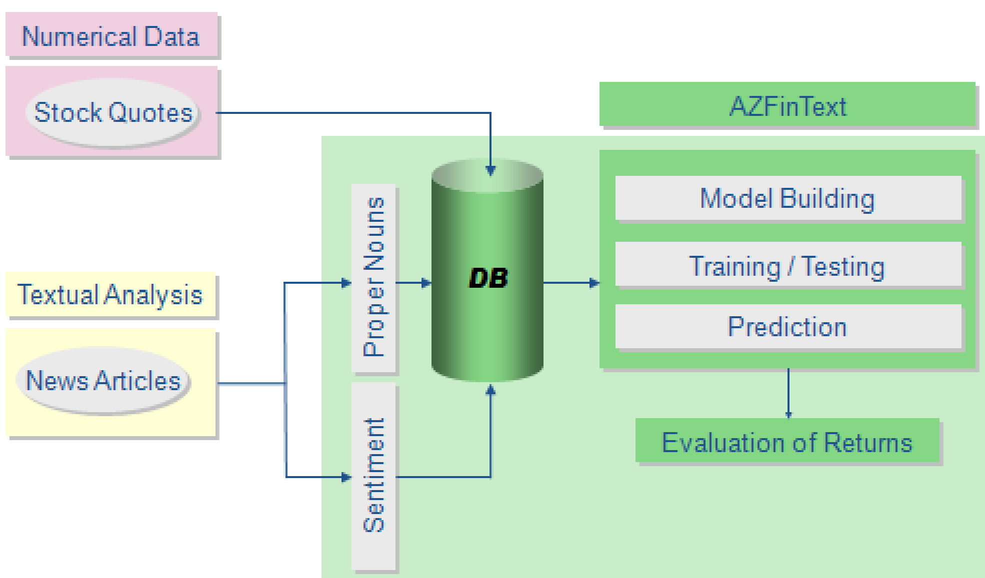

From the AZFinText system design in Figure 1, there are several major components. The first component is Numerical Data that gathers stock price data in one minute increments from a commercially available stock price database. The second component is Textual Analysis. This component gathers financial news articles from Yahoo! Finance and represents them by their proper nouns as well as by the sentiment of the article. This module further limits extracted features to three or more occurrences in any document, which reduces the noise from rarely used terms (Joachims, 1998).

Once the data is gathered, AZFinText makes +20min price predictions for each financial news article. From prior empirical testing, it was found that including the proper noun representation and the stock price at the time the news article was released, provided AZFinText with superior predictive performance compared to other textual representations and different pieces of price information (Schumaker & Chen, 2006).

At the Model Building stage of AZFinText, we partitioned the data to best answer the research questions. In the case of the Objectivity/Subjectivity, we used the OpinionFinder tool to make a determination of the overall tone of the article, i.e., is the article more objective or more subjective. In the case of Positive/Negative Subjectivity, we again used OpinionFinder to determine the overall polarity of the article.

For the machine learning algorithm we chose to implement the SVR Sequential Minimal Optimization (Platt, 1999) function through Weka (Witten & Eibe, 2005). This function allows discrete numeric prediction instead of classification. We selected a linear kernel and ten-fold cross-validation. A similar prediction method was employed in the forecasting of futures contracts (Tay & Cao, 2001).

AZFinText is then trained on the data and issues price predictions for each financial news article encountered. Evaluations are then made regarding the effect of stock returns in terms of the models generated.

5. Experimental Design

For the experiment, we selected a consecutive five week period of time to serve as our experimental baseline. This period of research was from Oct. 26, 2005 to Nov. 28, 2005 and incorporates twenty-three trading days. The five-week period of study was selected because it gathered a comparable number of articles in comparison to prior studies: 6,602 for Mittermayer (Mittermayer, 2004) and 5,500 for Gidofalvi (Gidofalvi, 2001). We also observe that the five-week period chosen did not have unusual market conditions and was a good testbed for our evaluation. In order to identify the companies with the most likelihood of having quality financial news, we limited our scope of activity to only those companies listed in the S&P 500 as of Oct. 3, 2005. Articles gathered during this period were restricted to occur between the hours of 10:30am and 3:40pm. Even though trading starts at 9:30am, we felt it important to reduce the impact of overnight news on stock prices and selected a period of one-hour to allow these prices to adjust. The 3:40pm cut-off was selected to disallow any +20 minute stock predictions to occur after market hours. A further constraint to reduce the effects of confounding variables was introduced where two articles on the same company cannot exist within twenty minutes of each other or both will be discarded. The above processes filtered the 9,211 candidate news articles gathered during this period to 2,802, where the majority of discarded articles occurred outside of market hours. Similarly, 10,259,042 per-minute stock quotations were gathered during this period. This large testbed of time-tagged articles and fine-grain stock quotations allow us to perform a systematic evaluation.

AZFinText's predictions were then analyzed against a three metric evaluation of Closeness, Directional Accuracy and a simple Trading Engine. Closeness, or how close AZFinText's predicted +20min value was to the actual +20min price, is measured in terms of Mean Squared Error (MSE) where $MSE = (1/n)\Sigma(\text{Predicted} - \text{Actual})^2$ (Cho et al., 1999). Directional Accuracy is simply how often AZFinText was correct in predicting the price direction of the +20min stock (Gidofalvi, 2001). For a Trading Engine, AZFinText utilized a modified version of Lavrenko's Trading Engine (Lavrenko et al., 2000a) that examines the percentage return of the stock. When a stock demonstrates an expected movement exceeding 1%, then $1,000 worth of that stock is either bought or shorted and then disposed of after twenty minutes. This modified version differs from Lavrenko's original design in regards to the dollar amount of stock bought. We further assume zero transaction costs, consistent with Lavrenko.

An example of AZFinText in operation is shown in Figure 2.

Figure 2. AZFinText Textual Example

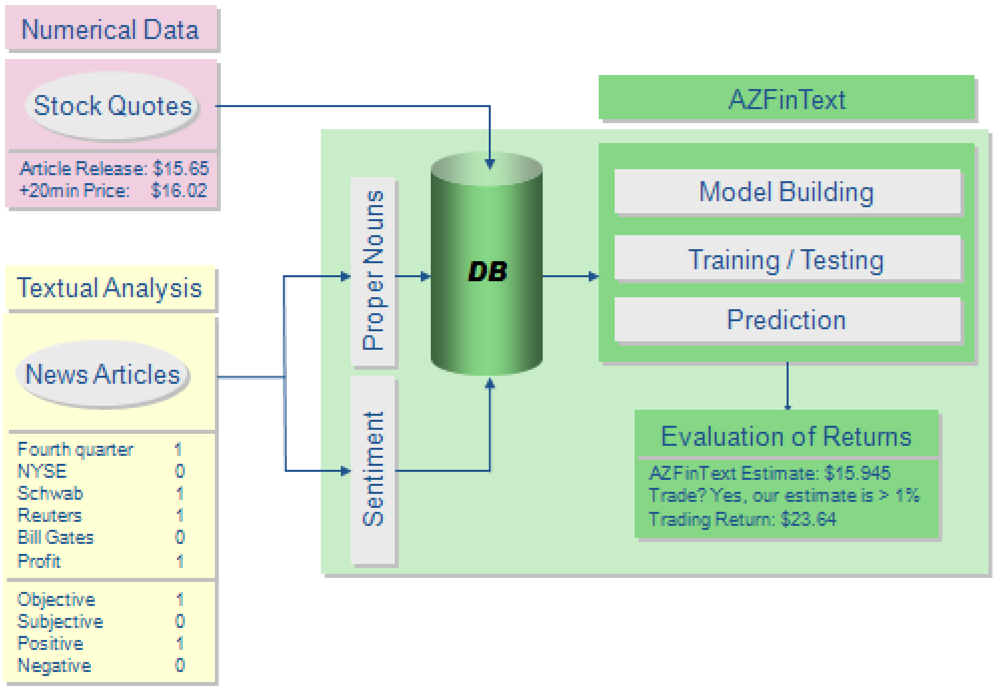

The first task is to extract financial news articles. The entire corpus of financial news articles is represented by their Proper Nouns in binary. If a particular Proper Noun feature is present in the article, that feature is given a 1, else a 0 and then stored in the database. Similarly, each financial news article is run through OpinionFinder to identify its overall tone, is it more Objective or Subjective, and its polarity, is the article more Positive or Negative. In tandem, stock quotations gathered on a per minute basis are stored. To build a model, we first pair together the representational Proper Nouns and stock quotation at the time the article was released, for each financial news article. Then, depending upon the particular model that is tested, data is aggregated and passed to our machine learning component for training and testing. Stock price predictions are then made for each financial news article and passed along to the evaluation instruments.

From the example above, AZFinText derived a prediction price of $15.945 which is greater than 1% of the stock price at the time the article was released, $15.65. Our trading engine makes a trade and disposes of it in twenty minutes time, for a trade return of $23.64 or 2.36%.

6. Experimental Results

To answer our first research question of *does Objectivity/Subjectivity impact news article prediction*, we tested five models of varying levels of financial news prediction and article tone from none (Baseline), all (Tone), majority objective (Objective), majority subjective (Subjective), majority neutral (Neutral). The results of the models are presented in Table 1.

Table 1. Article Tone results

	Baseline	Tone	Objective	Subjective	Neutral
# Articles	2,802	2,802	2,662	61	79
Closeness	**0.0516**	0.0565	0.0544	0.103	0.0930
Direction	50.4%	49.8%	49.5%	**59.0%**	53.2%
Trading	2.41%	2.00%	2.03%	**3.30%**	0.42%

Although Tone was essentially Baseline plus sentiment analysis, it performed worse against Baseline in all three metrics. Although it may seem that the addition of sentiment variables harmed AZFinText's predictive capability, when we broke apart Tone into its three constituent parts of Objective, Subjective and Neutral, it became apparent that several of Tone's components were depressing Tone's score. Objective articles were performing worse than chance in Directional Accuracy (49.5%) versus Baseline and Neutral articles had poor Trading Returns (0.42%) versus Baseline. By contrast, Subjective articles performed better. From these results, Baseline, which did not include any sentiment analysis in its model, had the best Closeness score of 0.0516. Subjective articles had the best Directional Accuracy (59.0%) and Trading Return (3.30%). We further believe that the author's use of subjectivity in the financial news article may have influenced market trading immediately following article release. These results were all significant as all values versus Baseline's values had p-values < 0.05.

To answer our second research question of *does Positive/Negative subjectivity impact news article prediction*, we again tested five models, varying the polarity of the article from not included (Baseline), all (Polarity), majority positive (Positive), majority negative (Negative), majority neutral (Neutral). The results of the models are presented in Table 2.

Table 2. Article Polarity results

	Baseline	Polarity	Positive	Negative	Neutral
# Articles	2,802	2,802	619	1,077	1,106
Closeness	**0.0516**	0.0556	0.0521	0.0576	0.0557
Direction	50.4%	49.4%	50.1%	**50.9%**	47.6%
Trading	2.41%	2.29%	1.73%	**3.04%**	1.98%

Again the model of Polarity which included whether the article was positive, negative or neutral, performed worse against Baseline in all three metrics. We suspect that the addition of the polarity variables was detrimental to AZFinText's predictions. However, by breaking apart Polarity into its component pieces, interesting results appeared. As shown in Table 2, Baseline again performed best in measures of Closeness (0.0516). Negative subjective articles performed best in Directional Accuracy (50.9%) and Trading Returns (3.04%). We believe that this may be a psychological reflection of market dynamics because negative emotions can have a larger and more lasting impact than positive or neutral ones. These results were all significant as all values versus Baseline's values had p-values < 0.05 except for Positive Closeness which was statistically equivalent to Baseline.

To pursue why negative articles were easier to predict, we analyzed the three component pieces of Polarity (Positive, Negative and Neutral) versus AZFinText's Directional Accuracy, as shown in Table 3.

Table 3. Article Polarity versus Directional Accuracy

Correct Predictions	Positive	Negative	Neutral
Upswings	46.0%	**52.4%**	**49.5%**
Downswings	**53.5%**	49.1%	46.0%

From this table, AZFinText worked best at predicting downswings of price in Positive Polarity articles (53.5%) and price upswings on both Negative and Neutral Polarity articles (52.4% and 49.5% respectively). All p-values < 0.05. AZFinText exhibited seemingly counter-intuitive results where positive articles were easier to predict decreasing and negative and neutral articles were easier to predict increasing in value. We believe that perhaps this result can be attributable to market psychology behaving in a contrarian manner, e.g., see good news, sell… see bad news, buy.

An example of this contrarian behavior is illustrated below with a PRNewsWire article on "Goodyear, Tire Industry Association Join Forces on Certification."

> The industry-first collaboration between a tire company and the international industry trade association includes two one-day seminars in 48 cities next year. Tire technicians from the Goodyear Dunlop and Kelly dealer network as well as Wingfoot Commercial Tire Centers are eligible for the Goodyear/TIA joint certification. Steve McClellan Goodyear's vice president of commercial tire systems said the customized training combines TIA s commercial tire service program and Goodyear's hands-on experience. TIA's certification process offers technicians valuable information that will help them perform their jobs better and with more care particularly as we embark on an aggressive campaign to expand our service business through a strong network of tire servicing outlets. McClellan said Goodyear seeks to improve its service business revenue which is more stable than cyclical commercial tire sales. Knowledgeable tire technicians across more than 2,000 commercial tire centers will deliver unrivaled service to our customers in their quest to lower costs. The service equation becomes a win-win for Goodyear and its dealer network. Al Cohn manager of strategic initiatives for Goodyear s commercial tire systems said dealer training is a major initiative. Our vision is to do more than just manufacture quality tires. We want to create business solutions that are measurable, repeatable and sustainable. That means helping dealers to work with fleets to manage their tire costs from original equipment to replacement and retreads and delivering service and value along the way. In addition the course exceeds OSHA training requirements for improved safety awareness. And that's why Goodyear/TIA technician certification makes sense. Through this collaboration the training provides a competitive advantage for independent dealers to grow their tire and service business, he said. Goodyear/TIA certification also may reduce dealer workers compensation costs, Cohn added. Goodyear commercial tire systems offer complete products and services to the trucking industry including a full range of original equipment and replacement tires. In addition the company's cradle-to-grave tire and service network includes retreading tire management tools and business solutions for tomorrow's trucking fleets. For more information on Goodyear's line of commercial tires, go to http://www.goodyear.com/truck

This news article had a stock price of $15.39 when the article came out. AZFinText predicted it's +20min stock price to be $15.136, a decline in price. Goodyear's actual +20min stock price was $15.22, a clear decline. This article was tagged by OpinionFinder as a Positive Polarity article. Since the price was stable prior to the article release, it would appear contrarian trading was taking place.

7. Conclusions

From our investigation we found several interesting results. The first of which was that AZFinText was best able to predict Subjective articles in Directional Accuracy (59.0% to 50.4%) and Trading Returns (3.30% to 2.41%), but not Closeness (0.103 versus 0.0516). We felt that the subjectivity of the articles may have influenced trading behavior. The second notable result was that AZFinText was best able to predict Negative Subjective articles in Directional Accuracy (50.9% to 50.4%) and Trading

Returns (3.04% to 2.41%), but not Closeness (0.0576 versus 0.0516). We believe that these results are attributable to investors reacting more strongly to negative articles. The third notable result was that AZFinText found evidence of Contrarian activity. AZFinText was better able to predict downswings in Positive articles (53.5%) and upswings in Negative and Neutral articles (52.4% and 49.5% respectively).

We would also suggest several future directions for this course of research. The first of which is to investigate the role of verbs and adverbs as a textual representation method. Perhaps this representational scheme will lead to better predictivity than Proper Nouns alone. It would be beneficial in future studies to draw comparisons between various textual representations. The second future direction would be train article tone and polarity separately. In prior work evaluating AZFinText on both momentum and contrarian portfolios, it was found that by training each portfolio separately, AZFinText was better able to pick up on nuances within the set. Applying that to this research, we feel a value added follow up study that investigates article tone and polarity in separate training sets may yield increased predictivity.

References

Abbasi, A., H. Chen, et al. 2008. Sentiment analysis in multiple languages: Feature selection for opinion classification in Web forums. ACM Transactions on Information Systems 26(3).

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

Camerer, C. & K. Weigelt 1991. Information Mirages in Experimental Asset Markets. Journal of Business 64(4): 463-493.

Cho, V., B. Wuthrich, et al. 1999. Text Processing for Classification. Journal of Computational Intelligence in Finance 7(2).

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

Davis, A., J. Piger, et al. 2006. Beyond the numbers: an analysis of optimistic and pessimistic language in earnings press releases. Technical Report. Federal Reserve Bank of St. Louis.

Devitt, A. & K. Ahmad 2007. Sentiment Polarity Identification in Financial News: A Cohesion-Based Approach. Association of Computational Linguistics, Prague, Czech Republic.

Fama, E. 1964. The Behavior of Stock Market Prices. The Journal of Business 38(1): 34-106.

Ghose, A., P. Ipeirotis, et al. 2007. Opinion Mining Using Econometrics: A Case Study on Reputation Systems. Association of Computational Linguistics, Prague, Czech Republic.

Gidofalvi, G. 2001. Using News Articles to Predict Stock Price Movements. Department of Computer Science and Engineering. University of California, San Diego.

Grefenstette, G., Y. Qu, et al. 2004. Coupling Niche Browsers and Afect Analysis for an Opinion Mining Application. 7th International Conference on "Recherche d'Information Assistee par Ordinateur", Avignon, France.

Jelveh, Z. 2006. How a Computer Knows What Many Managers Don't. The New York Times.

Joachims, T. 1998. Text Categorization with Support Vector Machines: Learning with Many Relevant Features. European Conference on Machine Learning, Chemnitz, Germany.

Lavrenko, V., M. Schmill, et al. 2000a. Language Models for Financial News Recommendation. International Conference on Information and Knowledge Management, Washington, DC.

LeBaron, B., W. B. Arthur, et al. 1999. Time Series Properties of an Artificial Stock Market. Journal of Economic Dynamics and Control 23(9-10): 1487-1516.

Malkiel, B. G. 1973. A Random Walk Down Wall Street. W.W. Norton & Company Ltd., New York.

McDonald, D. M., H. Chen, et al. 2005. Transforming Open-Source Documents to Terror Networks: The Arizona TerrorNet. American Association for Artificial Intelligence Conference Spring Symposia, Stanford, CA.

Mishne, G. 2005. Experiments with Mood Classification in Blog Posts. SIGIR, Salvador, Brazil.

Mittermayer, M. 2004. Forecasting Intraday Stock Price Trends with Text Mining Techniques. Hawaii International Conference on System Sciences, Kailua-Kona, HI.

Mowshowitz, A. 1992. On the Market Value of Information Commodities. The Nature of Information and Information Commodities. Journal of the American Society for Information Science 43(3): 225-232.

Pang, B., L. Lee, et al. 2002. Thumbs up?: Sentiment classification using machine learning techniques. Association for Computational Linguistics, Philadelphia, PA.

Platt, J. C. 1999. Fast Training of Support Vector Machines using Sequential Minimal Optimization. In Advances in Kernel Methods: Support Vector Learning, B. Scholkopf, C. Burges & A. Smola. MIT Press, Cambridge, MA, 185-208.

Raban, D. & S. Rafaeli 2006. The Effect of Source Nature and Status on the Subjective Value of Information. Journal of the American Society for Information Science and Technology 57(3): 321-329.

Riloff, E. & J. Wiebe 2003. Learning Extraction Patterns for Subjective Expressions. Conference on Empirical Methods in Natural Language Processing, Sapporo, Japan.

Schumaker, R. P. & H. Chen 2006. Textual Analysis of Stock Market Prediction Using Financial News Articles. Americas Conference on Information Systems, Acapulco, Mexico.

Sekine, S. & C. Nobata 2004. Definition, Dictionaries and Tagger for Extended Named Entity Hierarchy. Language Resources and Evaluation Conference, Lisbon, Portugal.

Subasic, P. & A. Huettner 2001. Affect Analysis o Text Using Fuzzy Semantic Typing. IEEE Transactions on Fuzzy Systems 9(4): 483-496.

Tay, F. & L. Cao 2001. Application of Support Vector Machines in Financial Time Series Forecasting. Omega 29: 309-317.

Technical Analysis 2005. The Trader's Glossary of Technical Terms and Topics. Retrieved Mar. 15, 2005.

Tetlock, P. 2007. Giving Content to Investor Sentiment: The Role of Media in the Stock Market. The Journal of Finance 62(3): 1139-1168.

Tolle, K. M. & H. Chen 2000. Comparing Noun Phrasing Techniques for Use with Medical Digital Library Tools. JASIS 51(4): 352-370.

Turney, P. D. 2002. Thumbs Up or Thumbs Down? Semantic Orientation Applied to Unsupervised Classification of Reviews. Association for Computational Linguistics (ACL'02), Philadelphia, PA.

Vapnik, V. 1995. The Nature of Statistical Learning Theory. Springer, New York.

Wiebe, J. & E. Riloff 2005. Creating Subjective and Objective Sentence Classifiers from Unannotated Texts. Sixth International Conference on Intelligent Text Processing and Computational Linguistics, Mexico City, Mexico.

Wiebe, J., T. Wilson, et al. 2001. Identifying Collocations for Recognizing Opinions. Association for Computational Linguistics, Toulouse, France.

Wiebe, J., T. Wilson, et al. 2004. Learning Subjective Language. Computational Linguistics 30(3): 277-308.

Wilson, T., P. Hoffmann, et al. 2005a. OpinionFinder: a system for subjectivity analysis. Human LanguageTechnology Conference, Vancouver, Canada.

Wilson, T., J. Wiebe, et al. 2005b. Recognizing Contextual Polarity in Phrase-level Sentiment Analysis. Conference on Human Language Technology and Empirical Methods in Natural Language Processing, Vancouver, Canada.

Witten, I. H. & F. Eibe 2005. Data Mining: Practical Machine Learning Tools and Techniques. Morgan Kaufmann, San Francisco.

Wuthrich, B., V. Cho, et al. 1998. Daily Stock Market Forecast from Textual Web Data. IEEE International Conference on Systems, Man, and Cybernetics, San Diego, CA.