

Textual Analysis of Stock Market Prediction Using Breaking Financial News: The AZFinText System

Robert P. Schumaker and Hsinchun Chen

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

Word Count: 7963

Abstract

Our research examines a predictive machine learning approach for financial news articles analysis using several different textual representations: Bag of Words, Noun Phrases, and Named Entities. Through this approach, we investigated 9,211 financial news articles and 10,259,042 stock quotes covering the S&P 500 stocks during a five week period. We applied our analysis to estimate a discrete stock price twenty minutes after a news article was released. Using a Support Vector Machine (SVM) derivative specially tailored for discrete numeric prediction and models containing different stock-specific variables, we show that the model containing both article terms and stock price at the time of article release had the best performance in closeness to the actual future stock price (MSE 0.04261), the same direction of price movement as the future price (57.1% directional accuracy) and the highest return using a simulated trading engine (2.06% return). We further investigated the different textual representations and found that a Proper Noun scheme performs better than the de facto standard of Bag of Words in all three metrics.

Categories and Subject Descriptors: G.1.10 [**Numerical Analysis**] Applications; H.4.2 [**Information Systems Applications**] Decision Support; I.2.7 [**Natural Language Processing**] Text analysis.

General Terms: Algorithms, Design, Experimentation, Measurement, Performance

Additional Key Words and Phrases: SVM, Stock Market, Prediction

Portions of this article were previously published as Schumaker, R., & Chen, H., (2006) Textual Analysis of Stock Market Prediction Using Financial News Articles. *12th Americas Conference on Information Systems (AMCIS)*

1 Introduction

Stock Market prediction has always had a certain appeal for researchers. While numerous scientific attempts have been made, no method has been discovered to accurately predict stock price movement. The difficulty of prediction lies in the complexities of modeling market dynamics. Even with a lack of consistent prediction methods, there have been some mild successes.

Stock Market research encapsulates two elemental trading philosophies; Fundamental and Technical approaches [25]. In Fundamental analysis, Stock Market price movements are believed to derive from a security's relative data. Fundamentalists use numeric information such as earnings, ratios, and management effectiveness to determine future forecasts. In Technical analysis, it is believed that market timing is key. Technicians utilize charts and modeling techniques to identify trends in price and volume. These later individuals rely on historical data in order to predict future outcomes.

One area of limited success in Stock Market prediction comes from textual data. Information from quarterly reports or breaking news stories can dramatically affect the share price of a security. Most existing literature on financial text mining relies on identifying a predefined set of keywords and machine learning techniques. These methods typically assign weights to keywords in proportion to the movement of a share price. These types of analyses have shown a definite, but weak ability to forecast the direction of share prices.

In this paper we experiment using several linguistic textual representations, including Bag of Words, Noun Phrases, and Named Entities approaches. We believe that combining more precise textual representations with past stock pricing information will yield improved predictability results.

This paper is arranged as follows. Section 2 provides an overview of literature concerning Stock Market prediction, textual representations, and machine learning techniques. Section 3 describes our research questions. Section 4 outlines our system design. Section 5 provides an overview of our experimental design. Section 6 expresses our experimental findings and discusses their implications. Section 7 delivers our experimental conclusions with a brief oratory on future directions for this stream of research.

2 Literature Review

When predicting the future prices of Stock Market securities, there are several theories available. The first is Efficient Market Hypothesis (EMH) [6]. In EMH, it is assumed that the price of a security reflects all of the information available and that everyone has some degree of access to the information. Fama's theory further breaks EMH into three forms: Weak, Semi-Strong, and Strong. In Weak EMH, only historical information is embedded in the current price. The Semi-Strong form goes a step further by incorporating all historical and currently public information in the price. The Strong form includes historical, public, and private information, such as insider information, in the share price. From the tenets of EMH, it is believed that the market reacts instantaneously to any given news and that it is impossible to consistently outperform the market.

A different perspective on prediction comes from Random Walk Theory [16]. In this theory, Stock Market prediction is believed to be impossible where prices are determined randomly and outperforming the market is infeasible. Random Walk Theory has similar theoretical underpinnings to Semi-Strong EMH where all public information is assumed to be available to everyone. However, Random Walk Theory declares that even with such information, future prediction is ineffective.

It is from these theories that two distinct trading philosophies emerged; the fundamentalists and the technicians. In a fundamentalist trading philosophy, the price of a security can be determined through the nuts and bolts of financial numbers. These numbers are derived from the overall economy, the particular industry's sector, or most typically, from the company itself. Figures such as inflation, joblessness, return on equity (ROE), debt levels, and individual Price to Earnings (PE) ratios can all play a part in determining the price of a stock.

In contrast, technical analysis depends on historical and time-series data. These strategists believe that market timing is critical and opportunities can be found through the careful averaging of historical price and volume movements and comparing them against current prices. Technicians also believe that there are certain high/low psychological price barriers such as support and resistance levels where opportunities may exist. They further reason that price movements are not totally random, however, technical analysis is considered to be more of an art form rather than a science and is subject to interpretation.

Both fundamentalists and technicians have developed certain techniques to predict prices from financial news articles. In one model that tested trading philosophies; LeBaron et. al. posited that much can be learned from a simulated stock market with simulated traders [15]. In their work, simulated traders mimicked human trading activity. Because of their artificial nature, the decisions made by these simulated traders can be dissected to identify key nuggets of information that would otherwise be difficult to obtain. The simulated traders were programmed to follow a rule hierarchy when responding to changes in the market; in this case it was the introduction of relevant news articles and/or numeric data updates. Each simulated trader was then varied on the timing between the point of receiving the information and reacting to it. The results were startling and found that the length of reaction time dictated a preference of trading

philosophy. Simulated traders that acted quickly formed technical strategies, while traders that possessed a longer waiting period formed fundamental strategies [15]. It is believed that the technicians capitalized on the time lag by acting on information before the rest of the traders, which lent this research to support a weak ability to forecast the market for a brief period of time.

In similar research on real stock data and financial news articles, Gidofalvi gathered over 5,000 financial news articles concerning 12 stocks, and identified this brief duration of time to be a period of twenty minutes before and twenty minutes after a financial news article was released [9]. Within this period of time, Gidofalvi demonstrated that there exists a weak ability to predict the direction of a security before the market corrects itself to equilibrium. One reason for the weak ability to forecast is because financial news articles are typically reprinted throughout the various news wire services. Gidofalvi posits that a stronger predictive ability may exist in isolating the first release of an article. Using this twenty minute window of opportunity and an automated textual news parsing system, the possibility exists to capitalize on stock price movements before human traders can act.

2.1 Textual Representation

There are a variety of methods available to analyze financial news articles. One of the more common methods is to apply a vector representation where article terms are indexed and then weighted. Selecting article terms can be as simple as tokenizing and using each word in the document. This technique assigns importance to determiners and prepositions which have little contribution to the overall meaning of the article. One method of circumventing these problems is to use a Bag of Words approach. In this approach, a list of semantically empty stop-words are removed from the article (e.g.; the, a, and for). The remaining terms are then used as the textual representation. The Bag of Words approach has been used as the de facto standard of financial

article research primarily because of its simple nature and its ability to produce a suitable representation of the text.

Building upon the Bag of Words approach, another tactic is to use a subset of terms as features [19], which can address issues related to article scaling while still encompassing the important concepts of an article [27]. One such method using this approach is Noun Phrasing. Noun Phrasing is accomplished through the use of a syntax where parts of speech (i.e., nouns) are identified through the aid of a lexicon and aggregated using syntactic rules on the surrounding parts of speech, forming noun phrases.

A third method of article representation is Named Entities. This technique builds upon Noun Phrases by using lexical semantic/syntactic tagging where nouns and noun phrases can be classified under predetermined categories [22]. This contrasts with using a differential approach, where concepts can be determined using a distributional analysis [14]. An example of the predetermined category approach is the MUC-7 framework of entity classification, where categories include date, location, money, organization, percentage, person and time. The entity tagging procedure can happen in a number of ways. Typically, successful taggers have large lexicons of sample entities and/or word patterns, which may include both syntax and lexical information. Lexicons and pattern information can be used to designate features or machine learning approaches or be incorporated in rule in more of a rules-based approach. The large quantities of patterns are considered relatively cheap or shallow knowledge to obtain. Thus reusability of the extraction rules is not a priority. When input text is matched to the stored extraction patterns the corresponding input text gets assigned an entity tag. Because of the constrained categories, Named Entities in effect provide the smallest coverage of the document,

but identify very specific types of phrases, which may or may not be helpful for stock price prediction.

Both Noun Phrases and Named Entities have shown limited success through previous comparison trials of tagging accuracy between differing algorithms. However, their usage as wide-scale textual representations for machine learning purposes remains somewhat unknown.

2.2 Machine Learning Algorithms

Like textual representation, there are also a variety of machine learning algorithms available. Almost all techniques start off with a technical analysis of historical security data by selecting a recent period of time and performing linear regression analysis to determine the price trend of the security. From there, a Bag of Words analysis is used to determine the textual keywords. Some keywords such as ‘earnings’ or ‘loss’ can lead to predictable outcomes which are then classified into stock movement prediction classes such as up, down, and unchanged. Much research has been done to investigate the various techniques that can lead to stock price classification. Table 1 illustrates a Stock Market prediction landscape of the various machine learning techniques.

Algorithm	Classification	Source Material	Examples
Genetic Algorithm	2 categories	Undisclosed number of chatroom postings	Thomas & Sycara, 2002
Naïve Bayesian	3 categories	Over 5,000 articles borrowed from Lavrenko	Gidofalvi et al. 2001
	5 categories	38,469 articles	Lavrenko et al. 2000
	5 categories	6,239 articles	Seo et al. 2002
SVM	3 categories	About 350,000 articles	Fung et al. 2002
	3 categories	6,602 articles	Mittermayer, 2004

Table 1. Prior algorithmic research

From Table 1, several items become readily noticeable. The first of which is that a variety of algorithms have been used. The second is that almost all instances commonly classify

predicted stock movements into a set of classification categories, not a discrete price prediction. Lastly, not all of the studies were conducted on financial news articles, although a majority was.

The first technique of interest is the Genetic Algorithm. In this study, discussion boards were used as a source of independently generated financial news [26]. In their approach, Thomas and Sycara attempted to classify stock prices using the number of postings and number of words posted about an article on a daily basis. It was found that positive share price movement was correlated to stocks with more than 10,000 posts. However, discussion board postings are quite susceptible to bias and noise.

Another machine learning technique, Naïve Bayesian, represents each article as a weighted vector of keywords [23]. Phrase co-occurrence and price directionality is learned from the articles which lead to a trained classification system. One such problem with this style of machine learning is from a company mentioned in passing. An article may focus its attention on some other event and superficially reference a particular security. These types of problems can cloud the results of training by unintentionally attaching weight to a casually-mentioned security.

One of the more interesting machine learners is Support Vector Machines (SVM). In the work of Fung et. al., regression analysis of technical data is used to identify price trends while SVM analysis of textual news articles is used to perform a binary classification in two predefined categories; stock price rise and drop [7]. In cases where conflicting SVM classification ensues, such as both rise and drop classifiers are determined to be positive, the system returns a 'no recommendation' decision. From their research using 350,000 financial news articles and a simulated Buy-Hold strategy based upon their SVM classifications, they showed that their technique of SVM classification was mildly profitable.

Mittermayer also used SVM in his research to find an optimal profit trading engine [18]. While relying on a three tier classification system, this research focused on empirically establishing trading limits. It was found that profits can be maximized by buying or shorting stocks and taking profit on them at 1% up movement or 3% down movement. This method slightly beat random trading by yielding a 0.11% average return.

Many of the prior studies were classification oriented with questions asked such as; *will this article cause the stock price to increase/decrease?* These studies were all tests of directional movement and not the predictors of stock prices. Discrete prediction from numeric trends is hardly new. However, the application of this regression technique to SVM mechanics is rather recent [8]. One such method is Sequential Minimal Optimization (SMO) [21], where many of the scalability problems from using large training sets has been obviated through a more simplistic SVM solving technique. This combination of techniques has led to the completely numeric prediction studies for futures contracts [24], but discrete prediction has not been coupled with a systematic study of various textual analysis methods before.

From prior studies on the textual representation of documents, Joachims posits that limiting the inclusion of features to three or more instances per document will avoid the problem of unmanageably large feature spaces [10]. Extending this to textual representation, each feature is further represented in binary as either a zero or one; the term is either present or not present in the article [28]. This simple representational scheme is easy to implement and will lead to a sparse dataset with many zero features.

Applying these regression based methods and textual representation techniques to a supervised machine learning algorithm such as SVM can lead to a trained system with discrete numeric output.

Evaluation of output has been generally focused on only one of the following three metrics; measures of Closeness, Directional Accuracy, or Simulated Trading. In measures of Closeness, the estimated value from machine learning is compared against the actual value in a Mean Squared Error (MSE) measure [20]. Directional Accuracy was the more common measure of previous financial studies, where the direction of the predicted value is compared with the movement direction of the actual value [4]. Whereas Simulated Trading initiates a simple trading engine to capitalize on large predicted value differences [13].

2.3 Financial News Article Sources

In real-world trading applications, the amount of textual data available to stock market traders is staggering. This data can come in the form of required shareholder reports, government-mandated forms, or news articles concerning a company’s outlook. Articles and reports are also routinely cross-posted in many different locations leading to problems of uniqueness and database selection [5]. Reports of an unexpected nature can lead to wildly significant changes in the price of a security. Table 2 illustrates some examples of textual financial data.

Textual Source	Types	Examples	Description
Company Generated Sources	SEC	8K	Reports on significant changes
	Reports	10K	Annual reports
Independently Generated Sources	Analyst Created	Recommendations	Buy/Hold/Sell assessments
		Stock Alerts	Alerts for share prices
	News Outlets	Financial Times	Financial News stories
		Wall Street Journal	Financial News stories
	News Wire	PRNewsWire	Breaking financial news articles
		Yahoo Finance	45 financial news wire sources
Discussion Boards	The Motley Fool	Forum to share stock-related information	

Table 2. Examples of textual financial data

Textual data can arise from two sources; company generated and independently generated sources. Company generated sources such as quarterly and annual reports can provide a rich linguistic structure that if properly read can indicate how the company will perform in the future [11]. This textual wealth of information may not be explicitly shown by financial ratios but rather encapsulated in forward-looking statements or other textual locations. Independent sources such as analyst recommendations, news outlets, and wire services can provide a more balanced view of the company and have a lesser potential to bias news reports. Discussion boards can also provide independently generated financial news; however, they can be suspect sources.

News outlets can be differentiated from wire services in several different ways. One of the main differences is that news outlets are centers that publish available financial information at specific time intervals. Examples include Bloomberg, Business Wire, CNN Financial News, Dow Jones, Financial Times, Forbes, Reuters, and the Wall Street Journal [3, 23]. In contrast, news wire services publish available financial information as soon as it is publicly released or discovered. News wire examples include PRNewsWire, which has free and subscription levels for real-time financial news access, and Yahoo Finance, which is a compilation of 45 news wire services including the Associated Press and PRNewsWire. Besides their relevant and timely release of financial news articles, news wire articles are also easy to automatically gather and are an excellent source for computer-based algorithms.

Stock Quotations are also an important source of financial information. Quotes can be divided into various increments of time from minutes to days, however, one minute increments provide sufficient granularity for machine learning.

While previous studies have mainly focused on the classification of stock price trends, none has been discovered to harness machine learning to determine a discrete stock price prediction based on breaking news articles. Prior techniques have relied solely on a Bag of Words approach and not other textual representations. Finally, there is no consensus on what information to include in a model that will lead to better performance. From these gaps in the research we form the crux of our study with the following questions.

3 Research Questions

Given that prior research in textual financial prediction has focused solely on the classification of stock price direction, we ask whether the prediction of discrete values is possible. This leads to our first research question.

- How effective is the prediction of discrete stock price values using textual financial news articles?

We expect to find that discrete prediction from textual financial news article is possible. Since prior research has indicated that certain keywords can have a direct impact on the movement of stock prices, we believe that predicting the magnitude of these movements is likely.

Prior research into stock price classification has almost exclusively relied on a Bag of Words approach. While this de facto standard has led to promising results, we feel that other textual representation schemes may provide better predictive ability, leading us to our second research question.

- Which combination of textual analysis techniques is most valuable in stock price prediction?

Since prior research has not examined this question before, we are cautious in answering such an exploratory issue. However, we feel that other textual representation schemes may serve to better distill the article into its essential components.

4 System Design

From these questions we developed the AZFinText system illustrated in Figure 1.

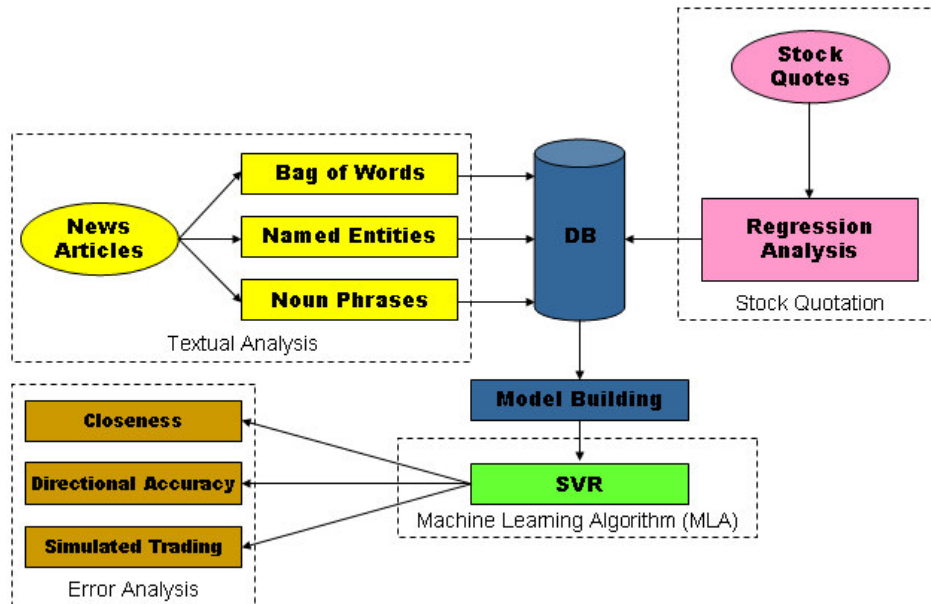


Figure 1. AZFinText system design

In this design, each financial news article is represented using three textual analysis techniques; Bag of Words, Noun Phrases, and Named Entities. These representations identify the important article terms and store them in the database. To limit the size of the feature space, we selected terms that occurred three or more times in a document [10].

To perform our textual analysis we chose a modified version of the Arizona Text Extractor (AzTeK) system which performs semantic/syntactic word level tagging as well as phrasal aggregation. AzTeK’s Noun Phrasing component works by using a syntactic tagger to identify and aggregate the document’s noun phrases and was found to have an 85% F-measure

for both precision and recall, which is comparable to other tools [27]. The Entity extractor portion goes one step further by assigning hybrid semantic/syntactic tags to document terms and phrases in one of the seven predefined categories of date, location, money, organization, percentage, person and time [17]. These entities are then identified through the usage of a lexicon. Although the AzTeK system was selected due to availability, it performs adequately for noun phrase and named entity extraction. However, there are many other such systems as reported in the Message Understanding Conference [17], that can be adopted for financial news text analysis.

Stock Quotes are gathered on a per minute basis for each stock. When a news article is released, we estimate what the stock price would be 20 minutes after the article was released. To do this we perform linear regression on the quotation data using an arbitrary 60 minutes prior to article release and extrapolate what the stock price should be 20 minutes in the future.

To test the types of information that need to be included, we developed four different models and varied the data given to them. The first model, Regress, was a simple linear regression estimate of the +20 minute stock price. Assuming that breaking financial news articles have no impact on the movement of stock prices, we would expect a reasonable performance from this model. While we acknowledge the obvious violation of Random Walk Theory, within such a compressed amount of time weak predictive ability remains [9]. Next the three models use the supervised learning of SVM regression to compute their +20 minute predictions. Model M1, uses only extracted article terms for its prediction. While no baseline stock price exists within this model, we chose it because of its frequent usage in prior studies on directional classification of stock prices. Model M2, uses extracted article terms and the stock price at the time the article was released. We feel that given a baseline of stock price that this

model will fare better. Model M3, uses extracted terms and a regressed estimate of the +20 minute stock price. This model may lead to better predictive results should the article terms have no impact on the movement of the stock price. All three models rely on using article terms in their prediction. SVM learns what terms lead to share price changes and adjust their weights depending on the severity of price changes.

To illustrate how the AZFinText system works, we offer a sample news article [2] and step through the logic of our system.

Schwab shares fell as much as 5.3 percent in morning trading on the New York Stock Exchange but later recouped some of the loss. San Francisco-based Schwab expects fourth-quarter profit of about 14 cents per share two cents below what it reported for the third quarter citing the impact of fee waivers a new national advertising campaign and severance charges. Analysts polled by Reuters Estimates on average had forecast profit of 16 cents per share for the fourth quarter. In September Schwab said it would drop account service fees and order handling charges its seventh price cut since May 2004. Chris Dodds the company's chief financial officer in a statement said the fee waivers and ad campaign will reduce fourth-quarter pre-tax profit by \$40 million while severance charges at Schwab's U.S. Trust unit for wealthy clients will cut profit by \$10 million. The NYSE fined Schwab for not adequately protecting clients from investment advisers who misappropriated assets using such methods as the forging of checks and authorization letters. The improper activity took place from 1998 through the first quarter of 2003 the NYSE said. This case is a stern reminder that firms must have adequate procedures to supervise and control transfers of assets from customer accounts said Susan Merrill the Big Board's enforcement chief. It goes to the heart of customers expectations that their money is safe. Schwab also agreed to hire an outside consultant to review policies and procedures for the disbursement of customer assets and detection of possible misappropriations the NYSE said. Company spokeswoman Alison Wertheim said neither Schwab nor its employees were involved in the wrongdoing which she said was largely the fault of one party. She said Schwab has implemented a state-of-the-art surveillance system and improved its controls to monitor independent investment advisers. According to the NYSE Schwab serves about 5 000 independent advisers who handle about 1.3 million accounts. Separately Schwab said October client daily average trades a closely watched indicator of customer activity rose 10 percent from September to 258 900 though total client assets fell 1 percent to \$1.152 trillion. Schwab shares fell 36 cents to \$15.64 in morning trading on the Big Board after earlier falling to \$15.16. (Additional reporting by Dan Burns and Karey Wutkowski)

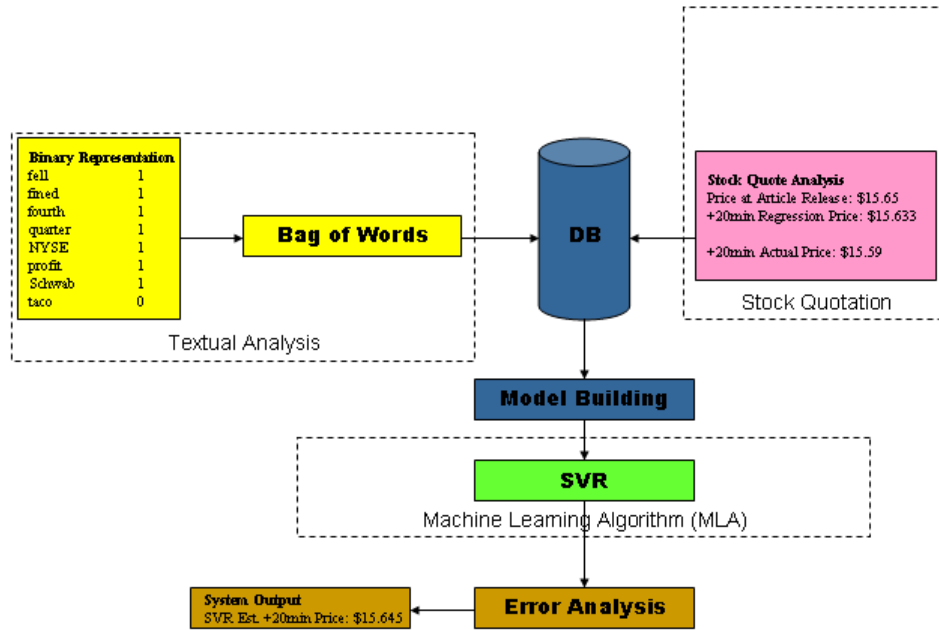


Figure 2. Example AZFinText representation

The first step in our system is to extract the text from each article using our three textual representations, independently. This builds up 3 separate corpora for each representation. Then each article is passed through the system, one at a time, as is shown in Figure 2 with the prior Schwab article using the Bag of Words representation. In the Textual Analysis box, extracted terms are represented in binary as either present or not in this article. Supposing our corpus also contained the term *Reuters* that appeared in many different articles but not this one, the term is given a zero for not being present in the current article. For stock quotation data, we lookup what the stock price was at the time of article release (\$15.65), calculate a regression estimate of the +20 minute stock price over the past hour (\$15.633) and lookup the actual +20 minute stock price for training and later evaluation (\$15.59). This data is then taken to the Model Building stage where the various models are given their appropriate data. Following that, machine learning takes place and an estimate of the +20 minute stock price is produced (\$15.645). We

can see from the stock prices given in this example that Schwab's share price dropped six cents while the model estimate figures a more conservative half penny drop.

5 Experimental Design

For our experiment we picked a research period of Oct. 26 to Nov. 28, 2005 to gather news articles and stock quotes. We further focused our attention only on companies listed in the S&P 500 as of Oct. 3, 2005. We acknowledge that several mergers and acquisitions did take place during this period of time; however, this only had an effect on less than 2% of the stocks tracked. In order to eliminate the 'company in passing' problem, we gathered the news articles from Yahoo Finance using a company's stock ticker symbol. This resulted in articles on 484 of the 500 companies listed in the S&P 500. Articles were further constrained to a time frame of one hour after the stock market opened to twenty minutes before the market closed. This period of time allows for sufficient data to be gathered for prior regression trend analysis and future estimation purposes. We further limited the influence of articles such that we did not use any two or more articles that occurred within twenty minutes of each other. This measure eliminated several possible avenues of confounding results.

By performing these actions we gathered 9,211 candidate financial news articles and 10,259,042 stock quotes over the five-week period. From this pool of news articles we analyzed them using the three textual representations and retained only those terms that appeared three or more times in an article, which results in a differing number of articles. The filtering process resulted in the following breakdown:

- Bag of Words used 4,296 terms from 2,839 articles
- Noun Phrases used 5,283 terms from 2,849 articles
- Named Entities used 2,856 terms from 2,620 articles

Article and stock quote data was then processed by a Support Vector Machine derivative, using Sequential Minimal Optimization [21] in a form of regression, which can handle discrete number analysis [24].

Following training, we chose three evaluation metrics; Closeness, Directional Accuracy, and a Simulated Trading Engine. The Closeness metric evaluated the difference between the predicted value and the actual stock price, measured using Mean Squared Error (MSE). Directional Accuracy measured the up/down direction of the predicted stock price compared with the actual direction of the stock price. While the inclusion of Directional Accuracy may not seem intuitive given the measure of Closeness, it is possible to be close in prediction yet predict the wrong direction of movement. This leads us to a third evaluation measure using a Simulated Trading Engine that invests \$1,000 per trade and follows simple trading rules. The rules implemented by our trading engine are a modified version of those proposed by Mittermayer to maximize short-term trading profit [18]. Our Simulated Trading Engine evaluates each news article and will buy/short the stock if the predicted +20 minute stock price is greater than or equal to 1% movement from the stock price at the time the article was released. Any bought/shorted stocks are then sold after 20 minutes. This assumes a zero transaction cost which is consistent with the research of Lavrenko [12, 13] and Mittermayer [18] who argue that trading in volume will offset the costs of trading.

6 Experimental Findings and Discussion

In order to answer our research questions on the effectiveness of discrete stock prediction and the best textual representation; we tested our three models against a regression-based predictor using the three dimensions of analysis; measures of Closeness, Directional Accuracy and a Simulated Trading Engine. Table 3 shows the results of the Closeness measures where

smaller numbers indicate less error in price prediction. Table 4 illustrates Directional Accuracy, where 50.0% could be achieved by chance alone. Finally, Table 5 displays the returns obtained from the Simulated Trading Engine.

MSE	Regress	M1	M2	M3
Bag of Words	0.07279	930.87	0.04422	0.12605
Noun Phrases	0.07279	863.50	0.04887	0.17944
Named Entities	0.07065	741.83	0.03407	0.07711
Average	0.07212	848.15	0.04261	0.12893

Table 3. Closeness results

Directional Accuracy	Regress	M1	M2	M3
Bag of Words	54.8%	52.4%	57.0%	57.0%
Noun Phrases	54.8%	56.4%	58.0%	56.9%
Named Entities	54.2%	55.0%	56.4%	56.7%
Totals	54.6%	54.6%	57.1%	56.9%

Table 4. Directional Accuracy results

Trading Engine	Regress	M1	M2	M3
Bag of Words	-1.81%	-0.34%	1.59%	0.98%
Noun Phrases	-1.81%	0.62%	2.57%	1.17%
Named Entities	-2.26%	-0.47%	2.02%	2.97%
Totals	-1.95%	-0.05%	2.06%	1.67%

Table 5. Simulated Trading Engine results

A. *Model M2 with articles terms and baseline stock price performed the best.*

From looking at the average results in Table 3, Model M2 which used both article terms and the stock price at the time of article release, had the lowest MSE score (0.04261) of any of the models (p-values < 0.05). This result signifies that Model M2's predictions were closer to the actual +20 minute stock price than any of the other models including linear regression (Regress). Looking deeper into the results, we find that Model M2 performed better than Regress in each of the three textual representations, which supports Gidofalvi's claim of weak short-term predictability.

Model M1, which used only article terms, had a difficult time in its estimation of future stock prices with an average Closeness score of 848.15. While this model may have been appropriate for prior classification-only studies, this unpleasant value was somewhat expected given the lack of baseline stock prices.

The other item of interest is that Model M2's Named Entities representation had the best performance at 0.03407 (p-values < 0.05). We will further investigate the effects of textual representation in a later section.

Turning our attention back to Model M2, we examined the weighting scheme that SVM assigned to the training variables. The stock price at the time of article release was given a weight of 0.9997 by SVM, while the article terms had a combined weight of 0.0003. While the weighting of article terms may appear superficially light, these terms are important because they provide the final touches to the estimated +20 minute stock price. If we were to rely on stock price alone without article terms, we would have the values of Regress, and Model M2 which used both the stock price and article terms performed better than Regress. This signifies that the 0.0003 combined weighting of article terms is an important element in providing more accurate price results. If instead we used a regressed price estimate plus article terms, we would have Model M3, and M2 performed better than M3. However, article terms alone were not sufficient in estimating the future stock price as demonstrated by Model M1.

In order to gain insight into the performance of our results, we can compare them to Pai and Lin whom conducted a similar study on forecasting stock prices [20]. In their study, they attempted stock price prediction one day in advance, using a small set of stocks and only close of day prices. They managed an MSE score of 0.3001 and comparing that to our average MSE scores in Table 3, our findings were an order of magnitude better.

In evaluating the Directional Accuracy results of Table 4, we again note that Model M2 performed better on average (57.1%) than the other models (p-values < 0.05). Regress did not perform so well (54.6%), which would seem to indicate that unexpected stock swings were captured by article terms. Comparing our results to previous studies shows that our values are somewhat reasonable. Cho et. al. which used 100 days of training articles and 392 keywords, had an average directional accuracy of 46.8% [4].

In the Simulated Trading results of Table 5, Model M2 using article terms and the stock price at the time of article release again had the best performance at 2.06% return (p-values < 0.05). This result would imply that Model M2 was better able to capitalize on trading opportunities given article terms and baseline stock price. Comparing this model against Model M1 with a trading return of -0.05%, we see that using article terms alone were insufficient. The results from Regress (-1.95% return) were unexpected. We believe that this finding was from the news articles themselves affecting major changes in the share price of stocks. Correlating our results with prior studies, Lavrenko et. al. claimed a 2% return from tracking four stocks over a forty day period [12]. In a similar study, Lavrenko et. al. expanded the number of stocks to 127 over the same 40 day period and had a much lower return of 0.23% [13]. Both of these studies used essentially the same trading mechanism as we did which leads to an interesting observation; that perhaps more stocks lead to lower returns, although our study tracked 500 companies over 23 trading days. In a third simulated trading study, Mittermayer obtained a 0.11% average return using all of the stocks from Nasdaq, NYSE, and AMEX over a one year period of time. From our results, it would appear that our system is achieving fairly reasonable results at a 2.09% return.

Compiling all of the model performances together, Model M2 using article terms and the stock price at the time of article release performed best in all three metrics; measures of Closeness (0.04261), Directional Accuracy (57.1%), and Simulated Trading (2.06% return). This model was better able to capture stock price movements and further bolsters the idea of weak short-term predictability. Our results were also inline with those from prior studies and mostly performed better. With some tweaking to how we classify directional movement, we feel that our system could produce better Directional Accuracy results as well.

B. A superset of Named Entities was the best textual representation.

To answer our second research question, *which combination of textual analysis techniques is most valuable in stock price prediction*, we compare the averages of each textual representation using our three metrics. Table 6 presents the results of Closeness measures, Table 7 displays Directional Accuracy and Table 8 illustrates the Simulated Trading Engine.

MSE	Regress	M1	M2	M3	Average
Bag of Words	0.07279	930.87	0.04422	0.12605	232.77789
Noun Phrases	0.07279	863.50	0.04887	0.17944	215.95020
Named Entities	0.07065	741.83	0.03407	0.07711	185.50404

Table 6. Closeness results

Directional Accuracy	Regress	M1	M2	M3	Totals
Bag of Words	54.8%	52.4%	57.0%	57.0%	55.3%
Noun Phrases	54.8%	56.4%	58.0%	56.9%	56.5%
Named Entities	54.2%	55.0%	56.4%	56.7%	55.6%

Table 7. Directional Accuracy results

Trading Engine	Regress	M1	M2	M3	Totals
Bag of Words	-1.81%	-0.34%	1.59%	0.98%	0.10%
Noun Phrases	-1.81%	0.62%	2.57%	1.17%	0.64%
Named Entities	-2.26%	-0.47%	2.02%	2.97%	0.57%

Table 8. Simulated Trading Engine results

From these tables, Named Entities had the lowest score in measures of Closeness (185.50404) and Noun Phrases had the better score in both Directional Accuracy (56.5%) and Simulated Trading (0.64%), all p-values < 0.05. These seemingly confusing results were not as clear-cut as our Model selection in the previous section as no one textual representation dominated the results.

However, it must be noted that these averaged results contain noise from previously failed models. If we were to focus only on the textual results for Model M2 and discard the other models, Noun Phrases performed the best in 2 of the 3 metrics and Named Entities in the remaining one.

These results ran contrary to our expectations. We had assumed that a Named Entity representation would generate better performance because of its ability to abstract the article terms and discard the noise of terms picked up by both Bag of Words and Noun Phrases. This MUC-7 textual representation was not sufficient to adequately model our article terms and lead us to ask the question, *what were the differences between Noun Phrases and Named Entities?* The answer was that Named Entities are essentially specialized Proper Nouns. The AzTeK system we used for part of speech tagging, identifies select terms in one of seven categories; date, location, money, organization, percentage, person and time [17]. Words in these categories are basically a subset of Noun Phrases. We believe that expanding the number of categories for Named Entities will lead to a better representational scheme.

In order to investigate this we took a subset of terms from Noun Phrases that were tagged as Proper Nouns and introduced a fourth, hybrid, textual representation of Proper Nouns. This selection of terms is a comparable superset of Named Entities but without the Entity categories.

Proper Nouns captured 3,710 article terms from 2,809 articles compared to the 5,283 terms in 2,849 articles for Noun Phrases and 2,856 terms in 2,620 articles for Named Entities.

To give the reader an understanding of what types of terms would be captured as Proper Nouns and not Named Entities, we refer back to the sample news article immediately preceding Figure 2. It is important to remember that Named Entities are derived using a semantic lexicon of previous input. Therefore, terms such as NYSE, which do not appear as a Named Entity, will be depicted in a Proper Noun representation.

Restating the metrics in terms of Model M2 to clear up some of the noise from the other failed models, we introduce the following data. Table 9 shows measures of Closeness, Table 10, Directional Accuracy and Table 11, Simulated Trading.

MSE	M2
Bag of Words	0.04422
Noun Phrases	0.04887
Proper Nouns	0.04433
Named Entities	0.03407

Table 9. Closeness results

Directional Accuracy	M2
Bag of Words	57.0%
Noun Phrases	58.0%
Proper Nouns	58.2%
Named Entities	56.4%

Table 10. Directional Accuracy results

Trading Engine	M2
Bag of Words	1.59%
Noun Phrases	2.57%
Proper Nouns	2.84%
Named Entities	2.02%

Table 11. Simulated Trading Engine results

The first item of interest is that the Proper Nouns subset performed better than Noun Phrases in all three metrics; 0.04433 to 0.04887 in measures of Closeness, 58.2% to 58.0% in Directional Accuracy and 2.84% to 2.57% in Simulated Trading (all p-values < 0.05). This would seem to back up our initial expectation that a more abstract textual representation would perform better. In comparison to Named Entities, Proper Nouns performed better in 2 of the 3 metrics, Directional Accuracy and Simulated Trading whereas Named Entities had better success at measures of Closeness. This would indicate that the direction we have undertaken is perhaps correct, but is still in need of refinement. We would suggest that future research should evaluate expanding the number of entity categories and evaluating the optimal mix for business-related news articles.

Overall, Bag of Words performed poorly by comparison. While this textual representation may be the de facto standard used in other studies, its weak performance is believed to arise from its reliance on too many noisy article terms. Noun Phrases performed much better with good performance in both Directional Accuracy and Simulated Trading. However, it suffered from poor Closeness measures. We believe that this is the result of using a better tuned representational scheme of news articles, as compared to the Bag of Words approach. Yet Noun Phrases still possessed some elements of noise which led to less than desirable Closeness scores. Named Entities had some problems and did not perform as expected. While this representation had the best Closeness score in prediction accuracy, it was unable to translate those gains into both Directional Accuracy and Simulated Trading returns. This is probably the result of using a limited set of Entity categories which was unable to fully represent the content of financial news articles. Finally, Proper Nouns had the better performance results. While this textual representation can be thought of as the hybrid go-between for Noun Phrases

and Named Entities, it had a solid performance on both Directional Accuracy and Simulated Trading. This result is likely attributable to Proper Nouns adequately using the article terms in a manner that was freer of the noise plaguing Noun Phrases and free of the constraining categories used by Named Entities.

7 Conclusions and Future Directions

Our first conclusion was that Model M2, using both article terms and the stock price at the time of article release, had a dominating performance in all three metrics; measures of Closeness at 0.04261, Directional Accuracy at 57.1% and Simulated Trading at a 2.06% return. These results were the direct consequence of this model's ability to capitalize on the article terms and stock price for machine learning.

Our second conclusion was that Proper Nouns had the better textual representation performance. While it performed best in 2 of the 3 metrics, Directional Accuracy at 58.2% and Simulated Trading at 2.84%, it pulled up short on measures of Closeness, 0.04433, as compared to Named Entities with 0.03407, all p-values < 0.05. However, this subset representation performed better than its parent, Noun Phrases, in all three metrics. We believe that Proper Nouns can attribute its success to being freer of the term noise used by Noun Phrases and free of the constraining categories used by Named Entities. Although more research into what constitutes an optimum mix of entity categories is encouraged.

Future research includes using other machine learning techniques such as Relevance Vector Regression, which promises to have better accuracy and fewer vectors in classification [1]. It would also be worthwhile to pursue expanding the selection of stocks outside of the S&P 500. While the S&P 500 is a fairly stable set of companies, perhaps more volatile and less tracked companies may provide interesting results. Another worthwhile approach would be to

test a model based on article terms and percentage of stock price change. While our models relied on fixed stock prices that traded within a consistent range, penny stocks with wild fluctuations may prove worthy of further research. Lastly, while we trained our system on the entire S&P 500, it would be a good idea to try more selective article training such as industry groups or company peer group training and examine those results in terms of prediction accuracy.

Finally, there are some caveats to impart to readers. While the findings presented here are certainly interesting, we acknowledge that they rely on a small dataset. Using a larger dataset would help offset any market biases that are associated with using a compressed period of time, such as the effects of cyclic stocks, earnings reports, mergers and other unexpected surprises.

References

- [1] Bishop, C.M. and M.E. Tipping, Bayesian Regression and Classification. *Advances in Theory: Methods, Models and Applications*, ed. G.H. J. Suykens, S. Basu, C. Micchelli, and J. Vandewalle. Vol. 190 of NATO Science Series III: Computer & Systems Sciences. 2003, Amsterdam: IOS Press.
- [2] Burns, D. and K. Wutkowski, Schwab to miss forecast, fined by NYSE, in *Excerpted from: http://biz.yahoo.com/rb/051115/financial_schwab.html?.v=3*. Nov. 15, 2005: Retrieved from Yahoo! News.
- [3] Cho, V., Knowledge Discovery from Distributed and Textual Data, in *Computer Science*. 1999, The Hong Kong University of Science and Technology: Hong Kong.
- [4] Cho, V., B. Wuthrich, and J. Zhang, Text Processing for Classification. *Journal of Computational Intelligence in Finance*, 1998. 26.
- [5] Conrad, J.G. and J.R.S. Claussen, Early user - system interaction for database selection in massive domain-specific online environments. *ACM Transactions on Information Systems*, 2003. 21(1): 94-131.
- [6] Fama, E., The Behavior of Stock Market Prices, in *Graduate School of Business*. 1964, University of Chicago.

- [7] Fung, G.P.C., J.X. Yu, X. Yu, and W. Lam. News Sensitive Stock Trend Prediction. *Pacific-Asia Conference on Knowledge Discovery and Data Mining (PAKDD)*. 2002. Taipei, Taiwan.
- [8] Gao, J.B., S.R. Gunn, C.J. Harris, and M. Brown, A Probabilistic Framework for SVM Regression and Error Bar Estimation. *Machine Learning*, 2002. 46(1 - 3): 71-89.
- [9] Gidofalvi, G., Using News Articles to Predict Stock Price Movements. 2001, University of California, San Diego: Department of Computer Science and Engineering.
- [10] Joachims, T., Text Categorization with Support Vector Machines: Learning with Many Relevant Features, in *Proceedings of the 10th European Conference on Machine Learning*. 1998, Springer-Verlag. 137-142.
- [11] Kloptchenko, A., T. Eklund, J. Karlsson, B. Back, H. Vanharanta, and A. Visa, Combining Data and Text Mining Techniques for Analysing Financial Reports. *Intelligent Systems in Accounting, Finance & Management*, 2004. 12(1): 29-41.
- [12] Lavrenko, V., M. Schmill, D. Lawrie, and P. Ogilvie. Mining of Concurrent Text and Time Series. *The Sixth ACM SIGKDD International Knowledge Discovery and Data Mining (KDD)*. 2000b. Boston, MA.
- [13] Lavrenko, V., M. Schmill, D. Lawrie, P. Ogilvie, D. Jensen, and J. Allan. Language Models for Financial News Recommendation. *Proceedings of the 9th International Conference on Information and Knowledge Management*. 2000a.
- [14] Le Moigno, S., J. Charlet, D. Bourigault, P. Degoulet, and M.-C. Jaulent, Terminology Extraction from Text to Build an Ontology in Surgical Intensive Care, in *AMIA Symposium*. 2002: San Antonio, TX.
- [15] LeBaron, B., W.B. Arthur, and R. Palmer, Time Series Properties of an Artificial Stock Market. *Journal of Economic Dynamics and Control*, 1999. 23(9-10): 1487-1516.
- [16] Malkiel, B.G., A Random Walk Down Wall Street. 1973, New York: W.W. Norton & Company Ltd.
- [17] McDonald, D.M., H. Chen, and R.P. Schumaker. Transforming Open-Source Documents to Terror Networks: The Arizona TerrorNet. *American Association for Artificial Intelligence Conference Spring Symposia*. 2005. Stanford, CA.
- [18] Mittermayer, M.-A. Forecasting Intraday Stock Price Trends with Text Mining Techniques. *Proceedings of the 37th Hawaii International Conference on Social Systems*. 2004. Hawaii.

- [19] Moldovan, D., M. Pasca, S. Harabagiu, and M. Surdeanu, Performance issues and error analysis in an open-domain question answering system. *ACM Transactions on Information Systems*, 2003. 21(2): 133-154.
- [20] Pai, P.-F. and C.-S. Lin, A hybrid ARIMA and support vector machines model in stock price forecasting. *Omega*, 2005. 33(6): 497-505.
- [21] Platt, J.C., Fast training of support vector machines using sequential minimal optimization, in *Advances in kernel methods: support vector learning*. 1999, MIT Press. 185-208.
- [22] Sekine, S. and C. Nobata. Definition, dictionaries and tagger for Extended Named Entity Hierarchy. *Proceedings of the LREC*. 2003.
- [23] Seo, Y.-W., J. Giampapa, and K. Sycara, Text Classification for Intelligent Portfolio Management. 2002, Carnegie Mellon University: Robotics Institute.
- [24] Tay, F. and L. Cao, Application of Support Vector Machines in Financial Time Series Forecasting. *Omega*, 2001. 29: 309-317.
- [25] Technical-Analysis, The Trader's Glossary of Technical Terms and Topics. 2005, <http://www.traders.com>.
- [26] Thomas, J.D. and K. Sycara. Integrating Genetic Algorithms and Text Learning for Financial Prediction. *Genetic and Evolutionary Computation Conference (GECCO)*. 2002. Las Vegas, NV.
- [27] Tolle, K.M. and H. Chen, Comparing Noun Phrasing Techniques for Use with Medical Digital Library Tools. *Journal of the American Society for Information Science*, 2000. 51(4): 352-370.
- [28] Vanschoenwinkel, B., A discrete Kernel Approach to Support Vector Machine Learning in Language Independent Named Entity Recognition. 2003, <http://citeseer.ist.psu.edu/682269.html>: Computational Modeling Lab, Vrije Universiteit, Brussel.