

Directional Prediction of Stock Prices using Breaking News on Twitter

Hana Alostad, Hasan Davulcu

*School of Computing, Informatics and Decision Systems Engineering, Arizona State University
Tempe, Arizona 85287–8809*

E-mail: {Hana.Alostad,Hasan.Davulcu}@asu.edu

Abstract. Stock market news and investing tips are popular topics in Twitter. In this paper, first we utilize a 5-year financial news corpus comprising over 50,000 articles collected from NASDAQ website matching 30 stock components in Dow Jones Index (DJI) to train a directional stock price prediction system based on news content. Next, we proceed to show that information in articles indicated by breaking Tweet volumes leads to a statistically significant boost in hourly directional prediction accuracies for the DJI stock prices mentioned in these articles. Secondly, we show that using document-level sentiment extraction does not yield a statistically significant boost in the directional predictive accuracies in the presence of other 1-gram keyword features. Thirdly we test the performance of the system on several time-frames and identify 4 hour time-frame for both the price charts and for Tweet breakout detection as the best time-frame. Finally, we develop a set of price momentum based trade exit rules to cut losing trades early and to allow winning trades run longer. We show that the Tweet volume breakout based trading system with price momentum based exit rules not only improve the winning accuracies and the return on investment, but they also lower the maximum drawdown and achieve highest overall return over maximum drawdown.

Keywords: stock prediction, text mining, breaking news, Twitter analysis, Twitter volume spikes

1. Introduction

Online social networks, like Twitter, are enabling people who are passionate about trading and investing to break critical financial news faster and they also go deeper into relevant areas of research and sources leading to real-time insights. Recently Twitter has been used to detect and forecast civil unrest [12], criminal incidents [30], box-office revenues of movies [9], and seasonal influenza [8].

Stock market news and investing tips are popular topics in Twitter. In this paper, first we utilize a 5-year financial news corpus comprising over 50,000 articles collected from the NASDAQ website for the 30 stock symbols in Dow Jones Index (DJI) to train a directional stock price prediction system based on news content. Next we collect over 750,000 Tweets during a 6 month period in 2014 that mention at least one of the 30 DJI stock symbols. We utilize the 68-95-99.7 rule, also known as the three-sigma rule or empirical rule [25], to

define a simple method for detecting hourly stock symbol related Tweet volume breakouts. Then we proceed to test our hypothesis to determine if “information in articles indicated by breaking Tweet volumes will lead to a statistically significant boost in the hourly directional prediction accuracies for the prices of DJI stocks mentioned in these articles”.

The contributions of the paper can be summarized as follows:

- Firstly, we show that sparse logistic regression [19] for this text based classification task with 1-gram keyword features filtered by a Chi2 [18] feature selection algorithm lead to the best overall directional prediction accuracy among a set of other classifiers and feature sets.
- Secondly, we show that using document-level sentiment extraction does not yield to a statistically significant boost in the predictive accuracies in the presence of other 1-gram keyword features.

Table 1
Summary of Previous Research Results

Reference	Data set			Time-frame	Period	Prediction	Algorithm	Accuracy
	Stock Price	Online News	Tweets					
[24]	✓	✗	✗	Daily	9 Yrs	Direction	Naive Bayes	90%
[13]	✓	✗	✗	Monthly	2 Yrs	Direction	Logistic Regression	83%
[26]	✓	✗	✗	Daily	13 Yrs	Price	Linear Regression	2.54 (RMSE)
[23]	✓	✓	✗	2 Hrs	4 Yrs	Direction	SVM	83%
[14]	✓	✓	✗	Daily	14 Yrs	Direction	SVM	79%
[15]	✓	✓	✗	Daily	1 Year	Direction	SVM	61%
[27]	✓	✓	✗	20 Min	1 Mo	Price	SVR	51%,
[16]	✓	✓	✗	Daily	1 Year	Direction	Neural Network	3.70 (RMSE)
[10]	✓	✗	✓	Daily	10 Mos	Direction	Neural Network	88%
[29]	✓	✗	✓	Daily	2 Mos	Direction	Decision Tree	77%
[22]	✓	✗	✓	Daily	3 Mos	Direction	Liner regression	68%
[21]	✓	✗	✓	Daily	1 Year	Price	Bayesian	0.3% (daily)

- Thirdly, we show that information in articles indicated by Tweet volume breakouts leads to a statistically significant boost in the hourly directional prediction accuracies for the DJI stocks mentioned in articles linked by Tweets.
- Fourthly, we compare the performance of breaking Tweet volumes based trading system on different time-frames. We identify 4 hour time-frame for both price charts and for Tweet volume breakouts detection as the best time-frame.
- Finally, we develop a set of price momentum based trade exit rules to cut losing trades early and to allow the winning trades run longer. We show that the Tweet volume breakout based trading system with the momentum based trade exit rules not only improve the average winning accuracy and the return on investment, but they also lower the maximum drawdown and yield the highest overall return over maximum drawdown (RoMaD).

The rest of the paper is organized as follows. In Section 2 we present related work. In Section 3 we present the problem definition for the directional prediction of stock prices. The design of experiments to evaluate the performance of various trading systems and strategies are presented in Section 4. In Section 5 we describe the experimental data sets that we used and the simu-

lated financial backtesting results for all experiments. In Section 6 we conclude the paper and discuss future work.

2. Related Work

Table 1 contains a summary of previous research findings related to stock price and direction prediction; in particular, we discuss data sets, time-frames, markets and algorithms used used for prediction and resulting overall accuracies.

These systems have different prediction time-frames and goals. Some of them predict stock price for the intended time-frame like [26], [27] and [21]. Time frames vary between next 20 minutes to up to one month. Works such as [10], [14], [15], [16], [22], [24], and [29] predict stock price direction for the next day. [23] aims to predict the price direction every 2-hours, and [13] aims to predict monthly direction.

Related systems collected their input data from various sources and exchanges: [27], [22], and [21] collected stock news, Tweets and price charts related to S&P 500 companies. [29] collected Tweets and stock price data related to Nasdaq stocks, [10] collected Tweets and stock price charts related to Dow Jones Industrial Average (DJIA), [15] collected one year

of data related to Microsoft company. [13] collected stock price charts from Shenzhen Development Stock A (SDSA) exchange. [23] collected currency price and news data related to foreign exchange market (Forex). [24] collected stock price data from CNX Nifty, S&P BSE Sensex exchanges and finally [26] collected thirteen years of stock price charts data related to Goldman Sachs Group Inc.

[13], [24], and [26] used only stock prices as input to predict stock price or stock direction with varying accuracies between 83% and 90%. [14], [15], [16], and [23] are examples of papers which utilize news as well as stock prices to predict price direction with varying accuracies between 51% and 83%.

[22] made correlation analysis between the stock price and the Tweet volume, and used it to predict stock market direction with 68% accuracy. Following work by [21] analyzed Tweet spikes in combination with price action based technical indicators, such as price breakout direction, as an input to a Bayesian classifier for stock price prediction, yielding a daily average gain of approximately 0.3% during a period of 55 days generating a total gain of 15%. [10] used extracted sentiment information from Twitter data and a neural network classifier to predict Dow Jones Industrial average (DJIA) daily price direction with 88% accuracy. [29] also used sentiment information extracted from Twitter as input to a decision tree classifier to predict price direction for four companies in NASDAQ stock exchange with an average accuracy of 77% reported as APPL at 77%, GOOG at 77%, MSFT at 69% and AMZN at 85% during a two month evaluation period.

3. Problem Definition

The correction effect of online news articles covering company related events, announcements and technical analyst reports on the stock price may take some time to show. Depending on the severity and impact of the news announcement this period may vary between few minutes to an hour, and the effect may sometimes determine the trend direction of the financial instrument for upcoming weeks or months.

One way to measure the impact of news on stock prices is to analyze the trading volume following the news announcement. Another indicator of news impact is the diffusion rates and volumes of messages on social media containing the stock symbol and news links of interest.

Twitter provides a suitable platform to investigate properties of such information diffusion. Diffusion analysis can harness social media to investigate “viral Tweets” to create early-warning indicators that can signal if a breakout started to emerge in its early stages. In this paper, we utilize the 68-95-99.7 rule to define a simple method of Tweet volume breakouts. In statistics, the 68-95-99.7 rule, also known as the three-sigma rule or empirical rule [25], states that in a normal distribution nearly all values lie within three standard deviations (σ) of the mean (μ). We utilize a fixed sized sliding window (of length 20 hour intervals that was determined experimentally), to compute a running average and standard deviation for the hourly volumes of Tweets that mention a stock symbol. Then, we identify *breakout* signals within a time-series of hourly Tweet volumes for each stock symbol whenever its hourly Tweet volume exceeds $(\mu(20) + 2\sigma)$. We consider a breakout as an indication that traders or technical analysts are sharing some exciting or important new information regarding the company or a group of companies. Next, we collect the URL links mentioned within the Tweets for breaking-news hour and we design a pair of experiments to test the hypothesis whether “information in news indicated by breaking Tweet volumes will lead to statistically significant boost in the directional prediction accuracy for the prices of the related stock symbols mentioned in these articles”.

Our system has the following characteristics:

1. Input Data: Hourly stock price charts of 30 stocks comprising the Dow Jones Index (DJI), online stock news articles for a 5 year period spanning 2010 to 2014 from NASDAQ¹ news website, the Tweets related to the 30 stock symbols collected from Twitter Streaming API² spanning a 6 months period between March 2014 until September 2014, and online news articles mentioned in Tweets during the breaking news hours.
2. Prediction Time-Frame: The collected data is analyzed and predictions are made on hourly bases.
3. Prediction Goal: To predict the hourly price direction for the stocks mentioned in Tweets during breaking news hours.

The distinguishing features of our system compared to systems mentioned in the related work section are: (1) [21] used Twitter volume spikes alongside stock price-based technical indicators for stock price turn-

¹<http://www.nasdaq.com/symbol/ibm/news-headlines>

²<https://dev.twitter.com/streaming/overview>

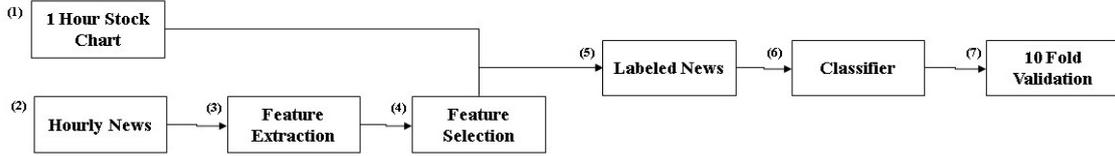


Fig. 1. Illustration of System Architecture of Experiment-1

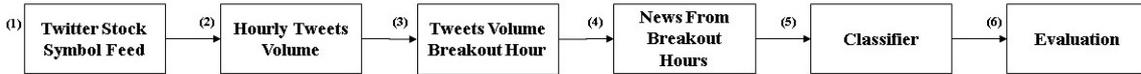


Fig. 2. Illustration of System Architecture of Experiment-2

ing point prediction where as our system utilizes textual content of the news mentioned in Tweets during breaking Twitter volume hours to predict the hourly direction of the stock price following a breakout period. (2) [10] and [29] used extracted sentiment information alongside stock price-based technical indicators to determine if sentiment information leads to a boost in directional prediction accuracy. Our system primarily relies on textual content of the news linked from breaking Tweet volumes to predict the direction of the stock price during the next hour. We also experimented with extracted sentiment as an additional feature to determine if it leads to a boost in the overall directional prediction accuracy. Unlike [10] and [29], our system did not experience a statistically significant boost in predictive accuracy as a result of including sentiment information alongside other textual content features. [10]’s accuracy is not comparable to ours since they are reporting the daily directional prediction accuracy for the Dow Jones Index Average (DJIA). Compared to predictive accuracies for four companies listed in [29], we have only one stock in common with their experiments, i.e. MSFT, where their system reported a daily directional predictive accuracy of 69% and our system reported an hourly directional accuracy of 82%.

4. Design of Experiments

In order to test the hypothesis that “information in news indicated by *breaking* Tweet volumes will lead to statistically significant boost in the directional prediction accuracy for the prices of the relevant stock symbols mentioned in such articles”, we designed five ex-

periments. In the first experiment we trained a classifier using *all* stock news articles for a 5 year period spanning 2010 to 2014 from NASDAQ news website. Figure 1 illustrates the system architecture that we used for the first experiment. For evaluation purposes we experimented with three different types of features extracted from text: 1-gram keywords, 2-gram phrases, and bi-polar sentiment (i.e. positive and negative) extracted from text. We grouped news hourly, and categorized each hourly collection as belonging to one of two categories: (1) those that led to an increased stock price or (2) those that led to a price reduction during the next hour. Next, we applied a feature selection method to reduce the features to relevant and non-redundant ones. The details of these steps are presented in Section 4.1. Finally we experimented with two types of text classifiers and evaluated their directional predictive accuracies using 10-fold cross validation. The results of the first experiment utilizing all news for all 30 company stocks are reported in Section 5.2.

In our second experiment, we evaluated the directional predictive accuracy of our classifier using *only* online articles collected during hourly breaking Tweet volume periods. Figure 2 illustrates the system architecture used for the second experiment. Steps involved in the second experiment were hourly profiling of the Tweet volumes mentioning a stock symbol, detection of breaking Tweet volume periods, collection of online news mentioned in Tweets during the volume breakout hours, feature extraction from news, and running the classifier to predict the stock price direction for the next hour by using the news content. We compared the accuracies of the classifiers in both first and second experiments to test the validity of our hypothesis. The de-

tails of steps involved in the design of the second experiment are explained in Section 4.2, and the results are presented in Section 5.3.

In the third experiment, we compared the performance of the breaking Tweet volumes based system on different time-frames to determine the time-frame yielding the best performance. The experimental results and evaluations for the third experiment are presented in Section 5.4. In Experiment-4 we developed a simulated trading system using the best performing time-frame and evaluated its performance. The experimental results for Experiment-4 are presented in Section 5.5.

Finally, in the fifth experiment, we developed a set of price momentum based trade exit rules to cut losing trades early and to allow the winning trades run longer. The results for Experiment-5 are presented in Section 5.6.

4.1. Experiment-1: Hourly Price Direction Prediction using Online News

Following is a detailed description of each step used in Experiment-1:

1. Hourly Stock Charts: We collected hourly stock price charts for all companies comprising the Dow Jones Index (DJI) using an API from ActiveTick³. For each trading hour the price direction was calculated based on the difference between hourly opening and closing prices according to the Formula 1, where d represents the trading date and h represents the trading hour:

$$Dir(d, h) = \begin{cases} 1 & \text{if } Open(d, h) \leq Close(d, h) \\ -1 & \text{otherwise} \end{cases} \quad (1)$$

2. Hourly News: We used Web Content Extractor⁴ to collect online news from NASDAQ website. We stored all metadata information related to news articles, such as, their titles, urls, dates, times, and sources. We fetched the news content and performed HTML stripping and content extraction using Boilerpipe⁵.

3. Feature Extraction for News:

- N-gram Features: R for Text Mining(TM)⁶ package was used to extract keyword features from the news corpus. First all whitespaces, stop words, numbers, punctuation were removed from the documents, then all terms were converted to lowercase and stemmed to their root forms. Next, features were recorded in a document-term matrix. For each stock symbol we created a pair of document-term matrices: one with 1-gram features and another with 2-gram features represented in a binary form. We used R.Matlab⁷ package to create Matlab format files for these matrices.

- Sentiment Features: To detect sentiment in news content we used SentiStrength library⁸. SentiStrength is a classifier that uses a predefined word list with expert-tagged sentiment polarity and strength judgments. Then it applies linguistic rules to detect sentiment in short text [28]. However, [20] showed that using general word lists for sentiment analysis of large financial text collections lead to misclassification of common words in the financial domain. So, alongside SentiStrength dictionary [20] we also used Loughran and McDonald Financial Sentiment Dictionary⁹ to compute sentiment in financial news. In order to get the sentiment of each document, we used OpenNLP¹⁰ Sentence Detector to extract its sentences mentioning a stock symbol, and then we applied the SentiStrength classifier on each sentence. For each stock symbol mentioned in the document, we determined the polarity of the sentences mentioning the symbol, and used the majority polarity (i.e. positive or negative) as the document sentiment for the stock.

4. Feature Selection: Feature selection in text mining reduces the features to only relevant and discriminative ones. We used Chi2 [18] feature selection algorithm from an online library¹¹. Chi2 is a two phase algorithm that automatically selects a proper critical value for a statistical χ^2 test to eliminate all irrelevant and redundant features [18].

⁶<http://cran.r-project.org/web/packages/tm/index.html>

⁷<http://cran.r-project.org/web/packages/R.matlab/index.html>

⁸<http://sentistrength.wlv.ac.uk/>

⁹http://www3.nd.edu/~mcdonald/Word_Lists.html

¹⁰<https://opennlp.apache.org/>

¹¹<http://featureselection.asu.edu/software.php>

³<http://www.activetick.com/>

⁴<http://www.newprosoft.com/>

⁵<http://code.google.com/p/boilerpipe/>

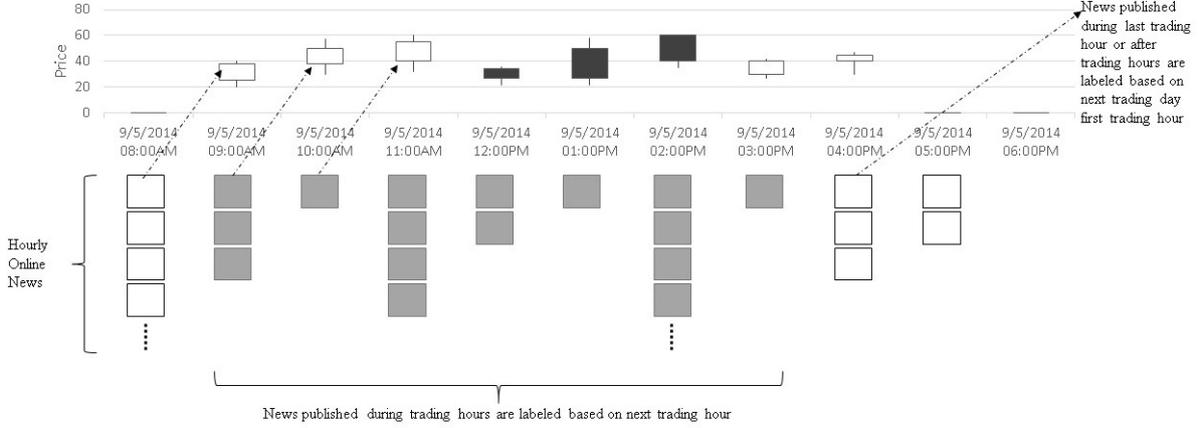


Fig. 3. An Illustration of News Labeling

5. News Labeling: Figure 3 shows an illustration of the news labeling step. We used the stock price direction of the following hour to determine the directionality of the hourly collections of news articles. In order to align the news article hours with the stock chart hours we had to standardize their time zones. Formula 2 is used to label the news articles where d represents the publishing date, and h represents the publishing hour.

$$Label(d, h) = Dir(d, Next(h)) \quad (2)$$

In this paper we initially assume that the effect of published news articles will be reflected on the stock price direction during the next hour. Later we relax this assumption and evaluate the system for all time periods varying between 5 mins up to 4 hours to determine the optimum timeframe. Formula 2 applies to all news articles published during official trading hours which start at 9AM and end at 3PM EST.

For articles that are published during the last trading hour, or after trading hours, or during holidays and weekends we assume that their effect will be seen on the direction of the first trading hour on the next trading day. For this case Formula 3 is used to label the directionality of corresponding news articles.

$$Label(d, h) = Dir(Next(d), First(h)) \quad (3)$$

6. News Classifier: We formulate price direction prediction problem as a classification problem in

a general structured sparse learning framework [19]. In particular, the logistical regression formulation presented below fits this application, since it is a dichotomous classification problem (e.g. upwards vs. downwards price correction). In Formula 4, a_i is the vector representation of the news during the i^{th} hour, w_i is the weight assigned to the i^{th} document ($w_i=1/m$ by default), and $A=[a_1, a_2, \dots, a_m]$ is the document n-gram matrix, y_i is the directionality of each hour based upon the stock price action of the next hour, and the unknown x_j , the j-th element of x , is the weight for each n-gram feature. $\lambda > 0$ is a regularization parameter that controls the sparsity of the solution, $|x|_1 = \sum |x_i|$ is 1-norm of the x vector. We used the SLEP [19] sparse learning package that utilizes gradient descent approach to solve the above convex and non-smooth optimization problem. The n-grams with non-zero values on the sparse x vector yield the discriminant factors for classifying a news collection as leading to upwards or downwards price correction. n-grams with positive polarity correspond to upward directional indicators, and those with negative polarity correspond to downward directional indicators.

$$\min_x \sum_{i=1}^n w_i \log(1 + \exp(1 + y_i(x^t a_i + c))) + \lambda |x| \quad (4)$$

Alternatively, we also utilized a SVM classifier

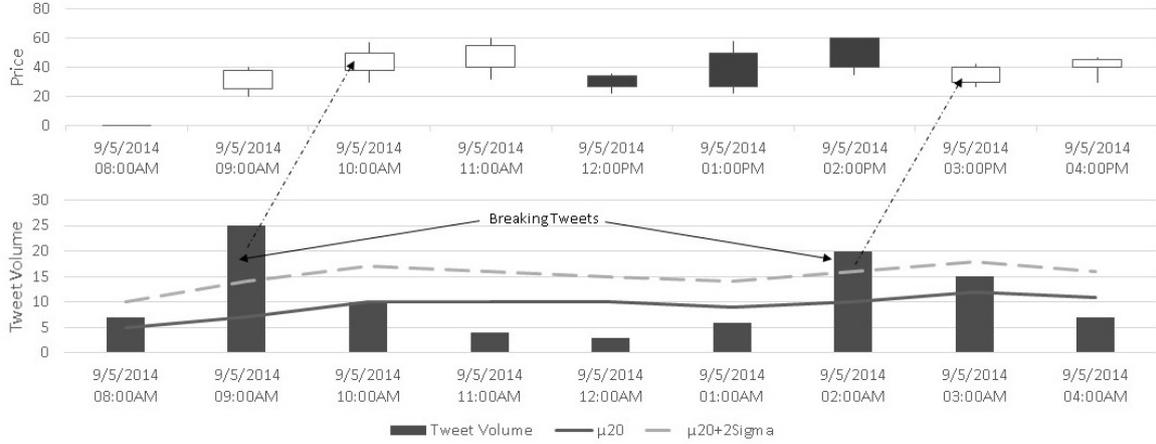


Fig. 4. An Illustration of Breaking Tweets

using LIBSVM¹² library in our experiments.

- 10-Fold Cross Validation: We performed a total of 8 experiments for each stock symbol where we experimented: (1) with SVM and sparse logistic regression classifiers, (2) with 1-gram and 2-gram features, and (3) with and without extracted sentiment features. After the training phase of the classifier, we validated the accuracies using 10-fold cross validation. The evaluation results for Experiment-1 are presented in Section 5.2.

4.2. Experiment-2: Hourly Price Direction Prediction using Breaking News

We selected the classifier with the best performance emerging from Experiment-1 to use in Experiment-2. Experiment-2 was designed to test if online news indicated by *breaking* Tweet volumes would lead to a statistically significant boost in the directional prediction accuracy for the prices of the relevant stock symbols mentioned in the news. The system architecture in Figure 2 shows the steps used in the design of this experiment. Following is a description of each step:

1. Twitter Stock Symbol Feed: Twitter streaming API was used to collect Tweets related to companies listed in the Dow Jones Index (DJI). In order to collect relevant Tweets we used a keyword filter made from the stock symbols, either prefixed by a dollar sign (\$) or prefixed by "NYSE:" or "NASDAQ:". For example, the key-

word filter for Microsoft Corporation are \$MSFT and NYSE:MSFT. For each matching Tweet we stored the stock symbol, Tweet text, date, time, and the set of URLs mentioned in the Tweet. If the Tweet text contained more than one stock symbol then we stored the same Tweet information for each stock symbol.

2. Hourly Tweet Volume Profiling: We utilize an experimentally determined fixed sized 20 hours sliding window, to compute a running average $\mu[20]$ and standard deviation σ for the hourly volumes of Tweets that mention a stock symbol.
3. Hourly Tweet Volume Breakouts: We identify *breakout* signals within hourly time-series of Tweet volumes for each stock symbol using Formula 5.

$$Breakout = \begin{cases} True & \text{if } N(d, h) \geq \mu[20](d, h) + 2\sigma(d, h) \\ False & \text{otherwise} \end{cases} \quad (5)$$

In Formula 5, N represents Tweet volume on specific date d and hour h . $\mu[20]$ is 20-hour simple moving average applied on Tweets volume. $\mu[20](d, h) + 2\sigma(d, h)$ represents the upper band for simple moving average - a 20-hour moving average plus 2-times its standard deviation. If the volume of hourly Tweets N exceeds the upper band value, this would indicate a volume breakout. Otherwise the hourly Tweet volume is classified as non-breaking. In Figure 4, the dotted ar-

¹²<http://www.csie.ntu.edu.tw/~cjlin/libsvm/>

Table 2
Tweets Breaking News Grouping and Prediction Time-Frames

News grouping Time-frames	Price Direction Labeling Time-frames				
	4h	1h	30m	15m	5m
4h	✓	✓	✓	✓	✓
1h		✓	✓	✓	✓
30m			✓	✓	✓
15m				✓	✓

rows shows two instances of Tweet volume break-outs at 9/5/2014 at 9AM and 9/5/2014 at 2PM, where the corresponding articles from these hours will be used to predict the price directions of the mentioned stocks at the following hours.

4. Breaking News: In this step the news content of URLs found in the Tweets during the breaking hours are downloaded and their textual contents are extracted using the following steps:
 - (a) For each breaking hour for a specific symbol we fetch the URLs found in Tweets. If the URLs are mentioned in their short forms, before fetching them, they are expanded to their long forms.
 - (b) We fetch the URL links content and perform content extraction using the JSoup HTML parser¹³.
5. News Classifier: After fetching hourly breaking news and extracting their 1-gram features we utilize a logistic regression classifier to predict the price direction for the next hour.
6. Evaluation: The predictive accuracies of the news classifier for the price direction following the breaking Tweet volume hours are presented in Section 5.3.

4.3. Experiment-3: Comparison Between Different Time-Frames For Price Direction Prediction using Breaking News

In Experiment-3 instead of using the 1 hour time-frame for news grouping and the price direction labeling, we evaluated the prediction performance of the news content based classifier using all possible time-frame combinations (i.e. 4 hours, 1 hour, 30 minutes, 15 minutes, and 5 minutes) for both news grouping and news price effects labeling. Table 2 lists the evaluated time-frame combinations. The goal of this experiment is to identify the best time-frame combina-

tion that should be used for (i) news grouping and (ii) news effect labeling which yields the highest directional stock price prediction accuracy. The findings of Experiment-3 are presented in Section 5.4.

4.4. Experiment-4: Tweets Volume Breakout based Trading System

We selected the classifier and the best time-frames as determined by the previous experiment in Experiment-4. Experiment-4 is designed to test if prediction with the online news linked from *breaking* Tweet volumes would lead to a higher performance directional stock price prediction system. Figure 5 shows the flowchart of the proposed design for Experiment-4. The following are the trade entry and exit rules that we used for the trading system:

1. ENTER a trade at the beginning of the next trade period, if there was a Tweet volume breakout for Tweets matching a stock's symbol in the preceding time-frame.
2. The trade DIRECTION (i.e. buy or sell) is determined by the news classifier applied to the content linked from the Tweets matching a stock symbol during the Tweet volume breakout period.
3. EXIT the trade at the end of the next time-frame period.
4. Evaluate the performance of the resulting trades (i.e. a win or a loss) and the corresponding amounts according to the real stock price movements during the evaluation period.

In order to evaluate the performance of the above trading system we performed a backtesting of the system on real data, recorded the outcomes of its trades accounting for wins and losses, as follows:

- Percentage of winning trades (Won%)
- Percentage of long positions (Long Won%)
- Percentage of short positions (Short Won%)

¹³<http://jsoup.org/>

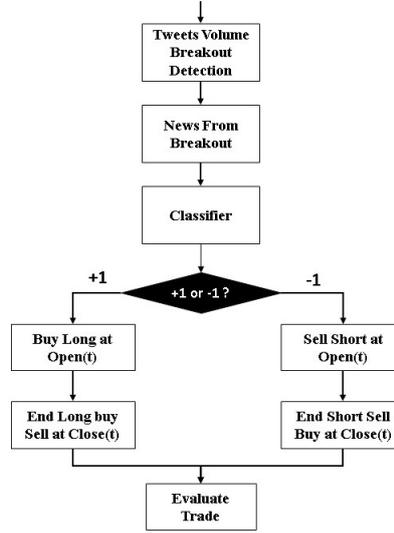


Fig. 5. A Flowchart of Tweets Breakout Trading Strategy

- Return On Investment (ROI%) according to Equation 6,

$$ROI = \frac{Gain - Cost}{Cost} \quad (6)$$

ROI is a performance measure used to evaluate the efficiency of an investment or to compare the efficiency of a number of different investments [5]. The higher the value of ROI% the better.

- Maximum Drawdown (MDD%), is the maximum loss from a peak to a trough of a portfolio, before a new peak is attained. It is an indicator of downside risk over a specified time period [3]. The lower the value of MDD% the better for investment.
- Return Over Maximum Drawdown (RoMaD) shown in Equation 7,

$$RoMaD = \frac{ROI}{MDD} \quad (7)$$

RoMad is a risk-adjusted return metric. It enables investors to ask the question: Are they willing to accept an occasional drawdown of X% in order to generate an average return of Y%? [6]. For example, an investment instrument with a ROI of 20% and MDD of 10% (RoMAD = 2.0) would be considered as more attractive than another with a ROI of 30% and a MDD of 20% (RoMAD = 1.5).

The experimental results for the Tweet volume breakout based trading system are presented in Section 5.5.

4.5. Experiment-5: Price Momentum Based Trade EXIT Rules

In this experiment we develop a set of price momentum based trade exit rules to cut losing trades early and to allow the winning trades run longer. We apply these price momentum based EXIT rules to the trading system developed in Experiment-4, and compare their performance in Section 5.6.

The rules are based on the Squeeze Momentum Indicator [7], which is a derivative of John Carter's "TTM Squeeze" volatility indicator [11]. This indicator has been used to detect periods while the market is quiet (i.e. squeeze) and the periods while the market is volatile (i.e. price breakouts). Squeeze Momentum Indicator is comprised of three components:

1. Bollinger Bands [2].

$$\begin{aligned} UpperBollingerBand &= \mu[20] + 2\sigma, \\ LowerBollingerBand &= \mu[20] - 2\sigma, \\ MiddleBollingerBand &= \mu[20] \end{aligned} \quad (8)$$

where $\mu[20]$ is the average of the closing prices for the previous 20 time-periods and σ is their standard deviation.

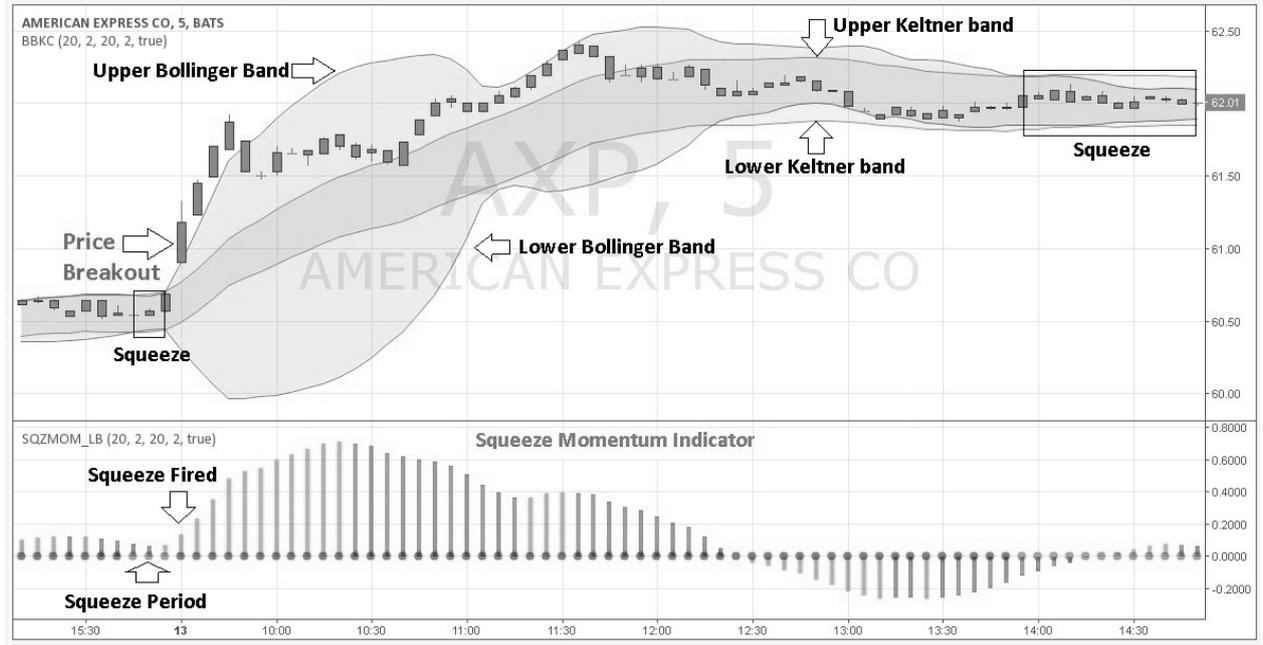


Fig. 6. Squeeze Momentum Indicator for AXP Stock symbol on 5 min Chart
Created using <https://www.tradingview.com> Charts

tion.

2. Keltner Channels [4].

$$\begin{aligned} UpperChanelLine &= \mu[20] + 2ATR(10), \\ LowerChanelLine &= \mu[20] - 2ATR(10) \end{aligned} \quad (9)$$

where $ATR[1]$ is defined as follows:

$$\begin{aligned} ATR(t) &= \frac{ATR(t-1) \times (n-1) + TR}{n} \\ TR &= \max[(high - low), \text{abs}(high, close_{prev}), \\ &\quad \text{abs}(low - close_{prev})] \end{aligned}$$

where t is the current time, $n=10$, and true range TR is the largest of either the most recent period's high minus the most recent period's low, the absolute value of the most recent period's high minus the previous close, or the absolute value of the most recent period's low minus the previous close.

3. Momentum Indicator [7].

$$\begin{aligned} Momentum &= close[0] - \mu[\mu[high[high, 20], \\ &\quad low[low, 20]], \mu[20]] \end{aligned} \quad (10)$$

where Momentum is the difference between the current close values to the average of the average between highest high of the previous 20 time periods and lowest low of the previous 20 time periods, to the average of the closing prices for the previous 20 time-periods.

Figure 6 illustrates the Squeeze Momentum Indicator components. The Squeeze Momentum Indicator is used as follows: It signals a red dot when the Bollinger Bands are inside of the Keltner Channel, hence the market is said to be in a *squeeze*. Otherwise, it signals a green dot signaling that the market is *volatile* (i.e. price breakout). In order to determine the direction of the volatility. We inspect the sign of the Momentum Indicator. If it is positive, then the price momentum is in the upward direction, otherwise it is in the downward direction.

The momentum based trade EXIT rules are defined as follows:

- Cut Losses Early (CLE) EXIT Rule is defined in Section 4.5.1.
- Conservative Let the Winners Run (ConsLWR) EXIT Rule is defined in Section 4.5.2.
- Aggressive Let the Winners Run (AggLWR) EXIT Rule is defined in Section 4.5.3
- Cut Losses Early (CLE) + Conservative Let the Winners Run (ConsLWR) EXIT Rule is defined in Section 4.5.4
- Cut Losses Early (CLE) + Aggressive Let the Winners Run (AggLWR) EXIT Rule is defined in Section 4.5.5

4.5.1. Trading with Cut Losses Early (CLE) EXIT Rule

This rule applies during the initial fixed time-period of the trade, where we track the price action on the 5 minute chart to cut the losses early if the price is *volatile* and the Momentum Indicator points to a direction that is opposite to that of the trade's direction.

1. Stock price is volatile and not in a squeeze period i.e. TTM squeeze indicator should be off (i.e. green)
2. The momentum indicator for the previous pair of bars are rising in an opposite direction.
 - Long EXIT: Momentum indicator for the previous pair of bars are both negative and declining.
 - Short EXIT: Momentum indicator for the previous pair of bars are positive and rising.

4.5.2. Trading with the Conservative Let the Winners Run (ConsLWR) EXIT Rule

In this strategy a trade is allowed to run, past its fixed time period, if it is in profit at the end of its fixed time-period and while the following conditions are true on the 5 minute price chart:

- Long/Buy Continuation: Momentum indicator for the previous pair of bars are rising.
- Short/Sell Continuation: Momentum indicator for the previous pair of bars are declining.

The trade is exited, using the 5 minute chart when one of the following conditions are met:

- Long/Buy EXIT: Momentum indicator for the previous pair of bars are declining.
- Short/Sell EXIT: Momentum indicator for the previous pair of bars are rising.

4.5.3. Trading with the Aggressive Let the Winners Run (AggLWR) EXIT Rule

In this strategy a trade is allowed to run, past its fixed time period, if it is in profit at the end of its fixed time-period and while the following conditions are true on the 5 minute price chart:

- Long/Buy Continuation: Momentum indicator for the previous pair of bars are positive.
- Short/Sell Continuation: Momentum indicator for the previous pair of bars are negative.

The trade is exited, using the 5 minute chart when one of the following conditions are met:

- Long EXIT: Momentum indicator for the previous bar is negative (i.e. opposite direction)
- Short EXIT: Momentum indicator for the previous bar is positive (i.e. opposite direction)

4.5.4. CLE + ConsLWR Trading Strategy

This strategy combines the CLE exit rule during the initial time-frame with the ConsLWR rule following the initial time-frame.

4.5.5. CLE + AggLWR Trading Strategy

This strategy combines the CLE exit rule during the initial time-frame with the AggLWR rule following the initial time-frame.

Experimental results of back-testing of the Tweets volume breakout based trading system with price momentum based trade EXIT rules are presented in Section 5.6.

5. Experimental Results and Evaluations

5.1. Experimental Data

We collected online news articles and stock price charts related to 30 stock symbols in Dow Jones Index for the period between October, 2009 and September, 2014. The total number of news articles collected from the NASDAQ website is 53,641. We also collected Tweets matching stock symbols for the period between March, 2014 and September, 2014. The total number of Tweets matching 30 stock symbols is 780,139. Table 3 shows the number of news articles, total number of collected Tweets, and the number of hourly breakout periods for each symbol.

5.2. Experiment-1: Hourly Price Direction Prediction using Online News

We executed the steps described in Figure 1 on data sets collected for the 30 Dow Jones Index companies.

Table 3

Counts of Collected News Articles, Tweets, and Calculated Hourly Breaking Tweets for 30 Dow Jones Stock Symbols

Stock Symbol	News	Tweets	Hourly Breaking Tweets	Avg URLs per Hourly Breaking Tweets
\$AXP	1614	15251	155	10
\$BA	2006	19041	207	15
\$CAT	1842	19303	188	18
\$CSCO	1984	26611	285	19
\$CVX	2168	15897	202	13
\$DD	1553	11218	63	11
\$DIS	1870	25014	199	17
\$GE	2260	31336	274	21
\$GS	1878	52888	246	40
\$HD	1743	18459	164	13
\$IBM	2188	80412	248	50
\$INTC	2157	30724	273	23
\$JNJ	2232	18236	224	16
\$JPM	1543	33658	274	27
\$KO	1899	22688	208	21
\$MCD	1879	21419	228	17
\$MMM	1183	11438	122	9
\$MRK	1573	28882	199	39
\$MSFT	1733	59469	271	54
\$NKE	1080	18206	196	18
\$PFE	1841	50859	292	57
\$PG	1781	15097	192	12
\$T	1784	30128	263	24
\$TRV	968	8858	52	7
\$UNH	1133	11224	144	10
\$UTX	1278	10872	68	9
\$V	1683	19174	190	13
\$VZ	2194	20896	252	16
\$WMT	2216	30448	236	27
\$XOM	2378	22433	248	15

In order to identify the best set of text features and the best classifier we performed several experiments detailed in Section 4.1. We run a total of 8 experiments for each stock symbol: (1) with SVM and sparse logistic regression classifiers, (2) with 1-gram and 2-gram keyword features, and (3) with and without extracted sentiment for documents. After the training phase of the classifier, we validated the accuracies using 10-fold cross validation. The results for the first experiment are presented in Table 4. The bold numbers on each row indicate the experimental setup which led to the best accuracy for each stock symbol. Figure 7 also shows the whisker plot for Experiment-1 results.

The evaluations show that 1-gram features led to higher overall accuracies compared to 2-gram features for both SVM and LogisticR experiments. Also, the experimental setup with the LogisticR classifier using 1-Gram features, where the sentiment features were excluded, led to maximal accuracies in 19 out of 30 instances. The second best experimental setup that achieved the maximal accuracies was also with the LogisticR classifier with 1-gram features integrated with

the sentiment feature. Hence, in order to determine the utility of extracted sentiment features we formulated the following hypotheses and applied the non-parametric sign test [17] at confidence level 95% to test if sentiment features would yield a statistically significant boost in the overall prediction accuracies:

1. Null Hypothesis (h0): 1-gram LogisticR classifier without sentiment features accuracies' mean = 1-gram LogisticR classifier with sentiment accuracies' mean, indicating that they are at the same level of performance.
2. Alternative Hypothesis (h1): 1-gram LogisticR classifier without sentiment features accuracies' mean \neq 1-gram LogisticR classifier with sentiment features accuracies' mean, indicating that they are not at the same level of performance.

The p-value of the sign test to compare 1-gram LogisticR classifier without sentiment features with 1-gram LogisticR classifier with sentiment features at significance level 0.05 equals to 0.1221, which leads to the acceptance of the null hypothesis h0 and the rejection of the alternative hypothesis h1, concluding that us-

Table 4
Accuracy Results of Experiment-1: Hourly Price Direction Prediction using Online News

Classification Method	SVM		LogisticR		SVM		LogisticR	
Feature Representation	1-Gram				2-Gram			
Sentiment Features	No	Yes	No	Yes	No	Yes	No	Yes
\$AXP	71.44%	70.19%	75.47%	74.28%	67.72%	67.66%	70.82%	70.26%
\$BA	68.70%	69.39%	72.64%	72.98%	66.60%	66.50%	68.29%	68.29%
\$CAT	72.21%	72.47%	76.49%	75.95%	65.15%	65.63%	66.18%	67.21%
\$CSCO	70.36%	71.17%	72.83%	72.68%	67.19%	66.58%	70.51%	70.11%
\$CVX	69.00%	69.01%	75.19%	75.19%	65.96%	64.62%	65.82%	65.27%
\$DD	66.39%	66.19%	71.73%	72.18%	65.29%	64.97%	68.19%	68.12%
\$DIS	72.03%	72.41%	74.76%	75.40%	66.04%	65.99%	68.82%	68.29%
\$GE	70.31%	69.78%	63.36%	63.19%	64.96%	64.56%	58.76%	58.19%
\$GS	71.19%	71.14%	74.60%	74.07%	66.35%	65.55%	67.62%	68.00%
\$HD	70.74%	71.95%	73.84%	74.76%	67.18%	67.64%	67.98%	68.15%
\$IBM	70.47%	70.48%	73.99%	73.81%	67.46%	66.68%	66.69%	67.37%
\$INTC	69.17%	69.49%	73.39%	73.85%	68.20%	69.36%	69.87%	70.01%
\$JNJ	70.39%	70.75%	72.72%	72.18%	64.87%	64.20%	66.22%	66.62%
\$JPM	72.13%	72.01%	76.22%	76.54%	66.37%	65.58%	68.38%	67.20%
\$KO	72.88%	73.83%	76.20%	76.04%	70.67%	69.88%	70.93%	70.62%
\$MCD	71.42%	71.64%	74.30%	74.93%	68.02%	67.91%	70.62%	70.73%
\$MMM	73.96%	74.22%	77.85%	77.85%	72.87%	72.44%	74.22%	74.72%
\$MRK	72.60%	72.54%	76.10%	75.27%	68.91%	68.34%	70.95%	71.27%
\$MSFT	73.92%	73.98%	77.85%	77.84%	68.72%	68.27%	71.61%	71.33%
\$NKE	73.24%	72.31%	79.17%	78.80%	72.13%	71.67%	74.54%	73.80%
\$PFE	72.95%	72.03%	76.32%	75.83%	70.07%	70.18%	72.03%	71.86%
\$PG	71.48%	72.26%	75.91%	76.47%	67.94%	68.16%	67.83%	68.72%
\$T	70.90%	70.90%	75.67%	75.67%	63.34%	64.57%	62.95%	65.14%
\$TRV	69.83%	70.98%	75.42%	75.11%	69.52%	67.66%	71.79%	72.00%
\$UNH	72.11%	72.64%	76.08%	75.82%	71.31%	71.32%	74.76%	74.24%
\$UTX	70.89%	71.76%	75.90%	75.67%	67.30%	68.86%	70.51%	70.89%
\$V	71.90%	73.50%	75.22%	76.30%	68.15%	67.32%	69.87%	68.87%
\$VZ	70.14%	68.96%	72.56%	71.83%	65.36%	64.31%	66.46%	66.22%
\$WMT	70.31%	69.90%	65.97%	65.57%	63.81%	64.53%	64.62%	63.99%
\$XOM	69.89%	69.72%	72.75%	71.70%	66.48%	66.53%	67.70%	68.04%

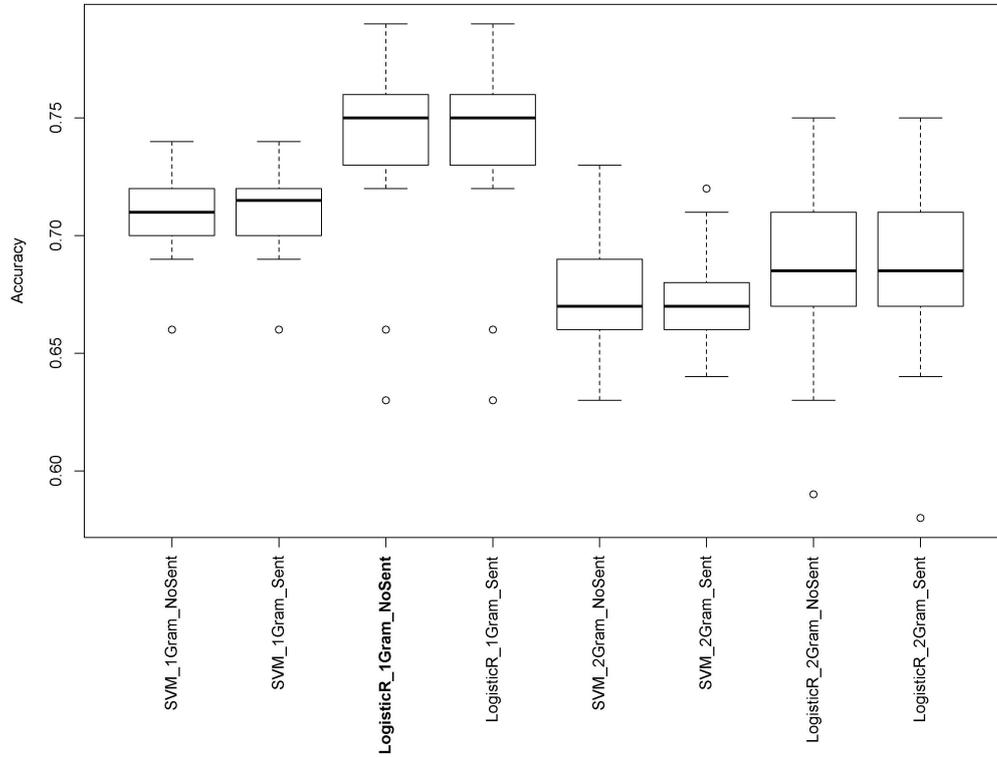


Fig. 7. Whisker Plot of Experiment-1 Accuracy

Table 5

Accuracy Results of Experiment 2: Hourly Price Direction Prediction using Breaking News

Stock Symbol	Experiment-1 Accuracy	Experiment-2 Accuracy	Stock Symbol	Experiment-1 Accuracy	Experiment-2 Accuracy
\$AXP	67.84%	69.66%	\$MCD	67.80%	65.30%
\$BA	68.20%	79.00%	\$MMM	70.60%	76.50%
\$CAT	65.37%	68.39%	\$MRK	69.30%	64.80%
\$CSCO	67.70%	69.55%	\$MSFT	72.60%	81.80%
\$CVX	66.60%	65.22%	\$NKE	70.30%	71.30%
\$DD	67.93%	66.67%	\$PFE	70.70%	59.50%
\$DIS	65.60%	64.75%	\$PG	70.80%	72.70%
\$GE	66.60%	66.67%	\$T	75.80%	82.90%
\$GS	68.30%	71.60%	\$TRV	65.40%	75.00%
\$HD	68.70%	75.00%	\$UNH	68.40%	75.20%
\$IBM	66.10%	80.80%	\$UTX	67.80%	58.60%
\$INTC	69.60%	67.30%	\$V	71.20%	71.00%
\$JNJ	64.40%	69.20%	\$VZ	65.30%	72.80%
\$JPM	67.50%	74.00%	\$WMT	66.60%	71.00%
\$KO	67.90%	68.60%	\$XOM	65.20%	67.30%

Table 6

Accuracy Results of Experiment-3: Performance Comparison Between Different Time-frames for News Grouping and Direction Labeling

News grouping Time-frames	Price Direction Labeling Time-frames				
	4h	1h	30m	15m	5m
4h	75%	71%	74%	74%	74%
1h	–	71%	72%	73%	71%
30m	–	–	68%	68%	68%
15m	–	–	–	65%	66%

ing sentiment would *not* yield a statistically significant boost in the overall prediction accuracy in this setup.

5.3. Experiment-2: Hourly Price Direction Prediction using Breaking News

In Experiment-2 we applied steps outlined in Figure 2 to 30 stock symbols in Dow Jones Index using breaking news periods only as trade triggers. Table 5 shows that Experiment-2 led to a boost in predictive accuracies for 70% of the stock symbols (i.e. 21 out of 30 cases). In order to prove that Experiment-2 yields a statistically significant boost in prediction accuracy compared to Experiment-1 we applied sign test at confidence level 95%. We formulated the following hypotheses:

1. Null Hypothesis (h0): Experiment-1 accuracies mean = Experiment-2 accuracies mean, indicating that they are at the same level of performance.
2. Alternative Hypothesis (h1): Experiment-1 accuracies mean \neq Experiment-2 accuracies mean, indicating that they are not at the same level of performance.

The p-value of the sign test to compare Experiment-1 with Experiment-2 at significance level 0.05 equals to 0.0357, which leads to the rejection of the null hypothesis h0 and the accepting of the alternative hypothesis h1 thus confirming that using 1-gram based LogisticR classifier with *breaking news* yields a *statistically significant boost* in directional prediction accuracy for 30 DJI stocks compared to using the same classifier with *all* of the stock news every hour.

5.4. Experiment-3: Comparison Between Different Time-frames For Price Direction Prediction using Breaking News

Fourteen pairs of time-frame combinations were tested in Experiment-3 where we applied the steps outlined in Figure 2 to each of the 30 stocks in Dow Jones Index using varying news grouping and price direction

labeling periods. Table 6 shows that in this experiment 4h4h time-frame combination yields the best average predictive accuracy for the price direction. This experiment indicates that (i) the 4 hour time period is the best time-period for detecting Tweet volume breakouts, and (ii) the 4 hour time-period is also the best time-frame to label and predict the trend direction following a Tweet volume breakout session.

5.5. Experiment-4: Tweets Volume Breakout based Trading System

We performed a simulated financial evaluation of the proposed trading system by back-testing its trades and accounting for its return on investment (ROI%) for a period of 6 months, between March 2014 and September 2014. In this simulation it is assumed that no commissions or fees are charged for each trade. Since the 4h4h time-frame yields the best accuracy results from Experiment-3 the system was tested using these time-frames for its trade entries and exists. The system entered a trade following a Tweet volume breakout session for Tweets matching a stock's symbol, with a trade fired in the direction (e.g. a long/buy or short/sell trade) predicted by our classifier based on the content that was collected by following the links from the Tweets during the breakout period, for a fixed duration of 4 hours. For each company, we measured the percentage of winning trades (Won%), percentage of long positions (Long Won%), percentage of short positions (Short Won%), return on investment (ROI%), maximum drawdown (MDD%), and risk adjusted return over maximum drawdown (RoMad).

Table 7 shows the results of the simulated back-testing evaluations. The system was profitable overall on its recommended trades with each stock symbol. Since each stock has a different stock price, we performed simulated trading using a diversified portfolio based on equal exposure to risk or gains from each stock in order to calculate the total and monthly average return on investment (ROI%). The simulated

Table 7
Experiment-4 Financial Evaluation of Tweets Breakout Trading System

Stock Symbol	Won%	Short Won%	Long Won%	ROI%	MDD	RoMaD
\$AXP	66%	100.00%	61.76%	6.90%	4.17	1.65
\$BA	65%	60.61%	68.75%	13.94%	3.52	3.96
\$CAT	69%	63.64%	72.00%	6.06%	2.66	2.28
\$CSCO	77%	74.19%	78.72%	15.10%	1.47	10.27
\$CVX	82%	85.00%	80.60%	13.67%	1.69	8.09
\$DD	76%	88.24%	71.70%	7.70%	2.39	3.22
\$DIS	78%	100.00%	75.00%	6.42%	3.99	1.61
\$GE	78%	83.87%	75.00%	14.96%	2.65	5.65
\$GS	76%	67.50%	83.67%	18.45%	1.52	12.14
\$HD	79%	100.00%	76.06%	20.20%	1.07	18.88
\$IBM	79%	83.33%	78.38%	10.46%	1.48	7.07
\$INTC	81%	100.00%	77.14%	31.90%	1.19	26.81
\$JNJ	77%	82.61%	74.55%	8.61%	2.60	3.31
\$JPM	77%	67.39%	89.47%	16.96%	2.83	5.99
\$KO	77%	86.67%	74.58%	12.46%	1.11	11.23
\$MCD	62%	60.42%	64.44%	0.50%	2.39	0.21
\$MMM	67%	88.89%	63.89%	0.65%	2.97	0.22
\$MRK	80%	85.00%	77.78%	31.41%	2.38	13.20
\$MSFT	80%	95.83%	72.88%	12.56%	2.64	4.76
\$NKE	75%	50.00%	75.36%	11.22%	1.61	6.97
\$PFE	61%	56.52%	68.57%	12.12%	2.97	4.08
\$PG	76%	82.14%	72.22%	14.12%	0.91	15.52
\$T	86%	95.00%	83.33%	20.13%	0.75	26.84
\$TRV	62%	64.71%	60.42%	1.08%	2.04	0.53
\$UNH	73%	85.71%	70.31%	22.28%	1.70	13.11
\$UTX	67%	76.19%	63.46%	3.15%	5.79	0.54
\$V	77%	94.12%	71.93%	12.24%	2.51	4.88
\$VZ	87%	80.00%	90.16%	26.46%	1.17	22.62
\$WMT	73%	68.97%	75.00%	13.55%	1.85	7.32
\$XOM	80%	81.08%	78.72%	20.09%	0.65	30.91
Average	74.73 %	80.25 %	74.20%	13.51 %	2.22	6.09

trades show that trading with the system during the 6 months period results in a winning ratio of 74% for its long/buy directional trades and 80% winning ratio for its short/sell directional trades. Trading with an equally diversified portfolio yields a total (ROI%) of 14% for 6 months, indicating an average monthly (ROI%) of 2.22% and RoMaD value of 6.09. The highest total (ROI%) achieved was 31.9% with Intel corpora-

tion(\$INTC), and the lowest total (ROI%) was 0.50% with Mcdonald's (\$MCD), the highest RoMaD value was achieved by trading AT&T (\$T) equals 26.84 and the lowest RoMaD was 0.21 achieved by trading Mcdonald's (\$MCD).

Table 8
Comparison Between Different Trading Strategy Results

	Won%	Short Won%	Long Won%	ROI%	MDD%	RoMaD
Tweets breakout	75%	80%	74%	14%	2%	7.00
Tweets breakout + CLE	79%	83%	77%	16%	2%	8.00
Tweets breakout + ConsLWR	74%	79%	73%	14%	3%	4.66
Tweets breakout + AggLWR	72%	78%	72%	16%	6%	2.66
Tweets breakout + CLE + ConsLWR	78%	83%	77%	17%	2%	8.5
Tweets breakout + CLE + AggLWR	78%	83%	77%	19%	2%	9.5

5.6. Experiment-5: Tweets Breakout Stock Trading System with Price Momentum Based Trade EXIT Rules

We performed a simulated financial evaluation of the Tweet volume breakout based stock trading system with the price momentum based trade EXIT rules as defined in Section 4.5. Table 8 shows a comparison of the average results of the proposed trading strategies based on the Won%, Short Won%, Long Won%, ROI%, MDD, and RoMaD. The results show that Tweets Breakout+CLE+AggLWR trading strategy yields the best risk adjusted return metric value (Ro-MaD) of 9.5 - meaning that, this system would yield 9.5% return for an occasional drawdown risk of 1%, or 95% returns for an occasional drawdown risk of 10%, essentially almost doubling the initial investment in 6 months. The second best trading strategy was obtained by the Tweets Breakout + CLE + ConsLWR with Ro-MaD value of 8.5.

6. Conclusion and Future Work

In this paper we start with a system to predict the hourly stock price direction based on the textual analysis of news articles' content mentioning a stock symbol. First, we show that using LogisticR classifier with 1-gram keyword features leads to the best overall directional prediction accuracy based on news articles. Next, we show that using extracted document-level sentiment features do not yield to a statistically significant boost in directional predictive accuracies in the presence of other 1-gram features. Then, we proceed to show that information in articles indicated by breaking Tweet volumes leads to a statistically significant boost in the hourly directional prediction accuracies for the prices of DJI stocks mentioned in these articles. We experiment with all time-frame combinations and identify the 4h time period as the best time-period

for detecting Tweet volume breakouts, and it is also as the best time-frame for the price-charts to label and predict the trend direction following a Tweet volume breakout session. Finally, we develop price momentum based trade exit rules to cut losing trades early and to allow the winning trades run longer. We show that the Tweet volume breakout based trading system with the momentum based exit rules not only improves the winning accuracy and the return on investment, but it also lowers the maximum drawdown and achieves the highest overall return over maximum drawdown. Our future work includes developing a real-time distributed trading system to monitor the Tweeter streams of different categories of stocks (i.e. large cap, mid cap and small cap) and trade with the their breaking volumes. We also plan to develop online learning methods to maintain the currency of the predictive models.

References

- [1] Average true range. https://en.wikipedia.org/wiki/Average_true_range. Accessed: 2016-05-02.
- [2] Bollinger bands. https://en.wikipedia.org/wiki/Bollinger_Bands. Accessed: 2016-04-29.
- [3] Drawdown (economics). [https://en.wikipedia.org/wiki/Drawdown_\(economics\)](https://en.wikipedia.org/wiki/Drawdown_(economics)). Accessed: 2016-05-01.
- [4] Keltner channels. http://stockcharts.com/school/doku.php?id=chart_school:technical_indicators:keltner_channels. Accessed: 2016-04-29.
- [5] Return on investment - roi. <http://www.investopedia.com/terms/r/returnoninvestment.asp>. Accessed: 2016-05-01.
- [6] Return over maximum drawdown (romad). <http://www.investopedia.com/terms/r/return-over-maximum-drawdown-romad.asp>. Accessed: 2016-05-01.
- [7] Squeeze momentum indicator. <https://www.tradingview.com/script/nqQ1DT5a-Squeeze-Momentum-Indicator-LazyBear>. Accessed: 2016-04-29.
- [8] H. Achrekar, A. Gandhe, R. Lazarus, S. Yu, and B. Liu. Twitter improves seasonal influenza prediction. In *HEALTHINF*

- 2012 - *Proceedings of the International Conference on Health Informatics*, pages 61–70, 2012.
- [9] S. Asur and B. Huberman. Predicting the future with social media. In *Proceedings of the 2010 IEEE/WIC/ACM International Conference on Web Intelligence and Intelligent Agent Technology*, volume 01, pages 492–499. IEEE Computer Society, 2010.
- [10] J. Bollen, H. Mao, and X. Zeng. Twitter mood predicts the stock market. 2(1):1–8, 2011.
- [11] J. F. Carter. *Mastering the Trade: Proven Techniques for Profiting from Intraday and Swing Trading Setups*. American Media International, 2007.
- [12] R. Compton, C. Lee, J. Xu, L. Artieda-Moncada, T.-C. Lu, L. Silva, and M. Macy. Using publicly visible social media to build detailed forecasts of civil unrest. 3(1), 2014.
- [13] J. Gong and S. Sun. A new approach of stock price prediction based on logistic regression model. In *2009 International Conference on New Trends in Information and Service Science NISS '09.*, pages 1366–1371, 2009.
- [14] M. Hagenau, L. Michael, and D. Neumann. Automated news reading: Stock price prediction based on financial news using context-capturing features. 55(3):685–697, 2013.
- [15] M. Y. Kaya and M. E. Karsligil. Stock price prediction using financial news articles. In *2010 Second IEEE International Conference on Information and Financial Engineering (ICIFE)*, pages 478–482. IEEE, 2010.
- [16] S. Lauren and S. D. Harlili. Stock trend prediction using simple moving average supported by news classification. In *2014 International Conference of Advanced Informatics: Concept, Theory and Application (ICAICTA)*, pages 135–139. IEEE, 2014.
- [17] E. L. Lehmann. *Nonparametrics :statistical methods based on ranks*. Holden-Day, San Francisco.
- [18] H. Liu and R. Setiono. Feature selection via discretization. 9(4):642–645, 1997.
- [19] J. Liu, S. Ji, and J. Ye. Slep: Sparse learning with efficient projections. 6:491, 2009.
- [20] T. Loughran and B. McDonald. When is a liability not a liability? textual analysis, dictionaries, and 10⁴Årks. 66(1):35–65, 2011.
- [21] Y. Mao, W. Wei, and B. Wang. Twitter volume spikes: analysis and application in stock trading. In *Proceedings of the seventh Workshop on Social Network Mining and Analysis*, page 4. ACM, 2013.
- [22] Y. Mao, W. Wei, B. Wang, and B. Liu. Correlating s&p 500 stocks with twitter data. In *Proceedings of the First ACM International Workshop on Hot Topics on Interdisciplinary Social Networks Research*, pages 69–72. ACM, 2012.
- [23] A. K. Nassirtoussi, S. Aghabozorgi, T. Y. Wah, and D. C. L. Ngo. Text mining of news-headlines for forex market prediction: A multi-layer dimension reduction algorithm with semantics and sentiment. 42(1):306–324, 2015.
- [24] J. Patel. Predicting stock market index using fusion of machine learning techniques. 42(4):2162–2172, 2015.
- [25] F. Pukelsheim. The three sigma rule. 48(2):88–91, 1994.
- [26] S. S. Roy, D. Mittal, A. Basu, and A. Abraham. Stock market forecasting using lasso linear regression model. In *Afro-European Conference for Industrial Advancement*, pages 371–381. Springer, 2015.
- [27] R. Schumaker and H. Chen. Textual analysis of stock market prediction using financial news articles. page 185, 2006.
- [28] M. Thelwall, K. Buckley, and G. Paltoglou. Sentiment strength detection for the social web. 63(1):163–173, 2012.
- [29] T.-T. Vu, S. Chang, Q. T. Ha, and N. Collier. An experiment in integrating sentiment features for tech stock prediction in twitter. 2012.
- [30] X. Wang, D. Brown, and M. Gerber. Spatio-temporal modeling of criminal incidents using geographic, demographic, and twitter-derived information. In *International Conference on Intelligence and Security Informatics*, Lecture Notes in Computer Science. IEEE Press, 2012.