
AP[®] Research Academic Paper

Sample Student Responses and Scoring Commentary

Inside:

Sample B

- Scoring Guideline**
- Student Samples**
- Scoring Commentary**

AP[®] RESEARCH — ACADEMIC PAPER

2019 SCORING GUIDELINES

The Response...				
Score of 1 Report on Existing Knowledge	Score of 2 Report on Existing Knowledge with Simplistic Use of a Research Method	Score of 3 Ineffectual Argument for a New Understanding	Score of 4 Well-Supported, Articulate Argument Conveying a New Understanding	Score of 5 Rich Analysis of a New Understanding Addressing a Gap in the Research Base
Presents an overly broad topic of inquiry.	Presents a topic of inquiry with narrowing scope or focus, that is NOT carried through either in the method or in the overall line of reasoning.	Carries the focus or scope of a topic of inquiry through the method AND overall line of reasoning, even though the focus or scope might still be narrowing.	Focuses a topic of inquiry with clear and narrow parameters, which are addressed through the method and the conclusion.	Focuses a topic of inquiry with clear and narrow parameters, which are addressed through the method and the conclusion.
Situates a topic of inquiry within a single perspective derived from scholarly works OR through a variety of perspectives derived from mostly non-scholarly works.	Situates a topic of inquiry within a single perspective derived from scholarly works OR through a variety of perspectives derived from mostly non-scholarly works.	Situates a topic of inquiry within relevant scholarly works of varying perspectives, although connections to some works may be unclear.	Explicitly connects a topic of inquiry to relevant scholarly works of varying perspectives AND logically explains how the topic of inquiry addresses a gap.	Explicitly connects a topic of inquiry to relevant scholarly works of varying perspectives AND logically explains how the topic of inquiry addresses a gap.
Describes a search and report process.	Describes a nonreplicable research method OR provides an oversimplified description of a method, with questionable alignment to the purpose of the inquiry.	Describes a reasonably replicable research method, with questionable alignment to the purpose of the inquiry.	Logically defends the alignment of a detailed, replicable research method to the purpose of the inquiry.	Logically defends the alignment of a detailed, replicable research method to the purpose of the inquiry.
Summarizes or reports existing knowledge in the field of understanding pertaining to the topic of inquiry.	Summarizes or reports existing knowledge in the field of understanding pertaining to the topic of inquiry.	Conveys a new understanding or conclusion, with an underdeveloped line of reasoning OR insufficient evidence.	Supports a new understanding or conclusion through a logically organized line of reasoning AND sufficient evidence. The limitations and/or implications, if present, of the new understanding or conclusion are oversimplified.	Justifies a new understanding or conclusion through a logical progression of inquiry choices, sufficient evidence, explanation of the limitations of the conclusion, and an explanation of the implications to the community of practice.
Generally communicates the student’s ideas, although errors in grammar, discipline-specific style, and organization distract or confuse the reader.	Generally communicates the student’s ideas, although errors in grammar, discipline-specific style, and organization distract or confuse the reader.	Competently communicates the student’s ideas, although there may be some errors in grammar, discipline-specific style, and organization.	Competently communicates the student’s ideas, although there may be some errors in grammar, discipline-specific style, and organization.	Enhances the communication of the student’s ideas through organization, use of design elements, conventions of grammar, style, mechanics, and word precision, with few to no errors.
Cites AND/OR attributes sources (in bibliography/ works cited and/or in-text), with multiple errors and/or an inconsistent use of a discipline-specific style.	Cites AND/OR attributes sources (in bibliography/ works cited and/or in-text), with multiple errors and/or an inconsistent use of a discipline-specific style.	Cites AND attributes sources, using a discipline-specific style (in both bibliography/works cited AND in-text), with few errors or inconsistencies.	Cites AND attributes sources, with a consistent use of an appropriate discipline-specific style (in both bibliography/works cited AND in-text), with few to no errors.	Cites AND attributes sources, with a consistent use of an appropriate discipline-specific style (in both bibliography/works cited AND in-text), with few to no errors.

AP[®] RESEARCH 2019 SCORING COMMENTARY

Academic Paper

Overview

This performance task was intended to assess students' ability to conduct scholarly and responsible research and articulate an evidence-based argument that clearly communicates the conclusion, solution, or answer to their stated research question. More specifically, this performance task was intended to assess students' ability to:

- Generate a focused research question that is situated within or connected to a larger scholarly context or community;
- Explore relationships between and among multiple works representing multiple perspectives within the scholarly literature related to the topic of inquiry;
- Articulate what approach, method, or process they have chosen to use to address their research question, why they have chosen that approach to answering their question, and how they employed it;
- Develop and present their own argument, conclusion, or new understanding while acknowledging its limitations and discussing implications;
- Support their conclusion through the compilation, use, and synthesis of relevant and significant evidence generated by their research;
- Use organizational and design elements to effectively convey the paper's message;
- Consistently and accurately cite, attribute, and integrate the knowledge and work of others, while distinguishing between the student's voice and that of others;
- Generate a paper in which word choice and syntax enhance communication by adhering to established conventions of grammar, usage, and mechanics.

USING SENTIMENT ANALYSIS TO PREDICT GOOGLE STOCK PRICES

Using Sentiment Analysis to Predict Google Stock Prices

Word Count: 4434

USING SENTIMENT ANALYSIS TO PREDICT GOOGLE STOCK PRICES

Introduction

Stock markets play an active role in the modern-day economy. In the month of February in 2019, the NASDAQ Stock Market recorded an average of 12 million shares traded daily (NasdaqTrader 2019). Based on a survey conducted in 2016, it is estimated that 52% of Americans have money invested in the stock market (Jones and Saad 2016, 1). Despite the vast number of stock investors, they mostly strive for a common goal of profiting. Typically, traders desire to find and purchase stocks that will rise in terms of prices in the future. As time passes and their stocks' value rises, they can choose to sell it at a higher price than before and earn a profit. As a result, it is essential for traders to have the ability to foresee future stock trends. This skill of prediction enables the trader to select the stock with great potential to rise in value and acquire it at a low price.

Living in the age of the Internet, people are able to express their opinions with ease. Online news, public forums, and social media are examples of popular platforms available for people to communicate their thoughts. Facebook, an American online social media company, recorded 4 billion pieces of content posted daily in 2012 (Wilson et al. 2012, 203). Through online posts, users indirectly indicate their attitudes and views on certain events. This immense online content can be treated as data suggesting the mood of the public. The sentiment of online news attracts the attention of stock investors as it is directly related to the market. Traders read and interpret news articles related to their investments. The sentiment conveyed by the most up-to-date news will impact their decisions to buy or sell their stocks. Therefore, in logical terms, the general opinion of online reports has an impact on the stock market. By collecting the sentiment of news, analysts can derive a correlation between the public's mood with the stock market.

USING SENTIMENT ANALYSIS TO PREDICT GOOGLE STOCK PRICES

Literature Review

Numerous studies had been done to examine the accuracy and applications of sentiment analysis. Sentiment analysis can be defined as a computational method that extracts opinions by analyzing raw text data (Kechaou et al. 2011, 1032). There are innumerable real-world applications of sentiment analysis, such as labeling customer reviews, developing recommendation systems, and etc. Studies in these areas generally reported a high accuracy in analyzing the sentiment of text data. For instance, in 2002, Bo Pang, a graduate student studying computer science at Cornell University, tested the accuracy of machine-based sentiment classification in analyzing movie reviews. He compared the sentiment results from movie reviews on IMDb with the corresponding numerical ratings (Pang et al. 2002, 80). From his results, the Naïve Bayes classifier, a popular sentiment analysis technique using simple probabilities (Explained in detail in next paragraph), obtained a 78.7% accuracy of correctly labeling movie reviews as either positive, negative, or neutral.

One study done by Sarkis Agaian and Petter Kolm in 2017 focused on the accuracy of sentiment analysis in financial news. Agaian, a consultant at Capstone Investment Advisors, compared the accuracy of measuring the sentiment of business news using support vector machine, maximum entropy, and Naïve Bayes classifiers (Agaian and Kolm 2017, 3). In detail, the Naïve Bayes classifier relies on Bayes' Theorem of Probability, which explains the probability of an event based on the conditions that could be associated with the event. Maximum entropy finds and determines a data group by considering the most extreme scenario of the dataset. The support vector machine classifies two clusters of data (In this case, whether the word is positive or negative) by finding the best-fit curve that can cut between them. Despite using various machine learning algorithms, Agaian's results revealed an average classification accuracy of around 75%. Through this study, Agaian concluded that using sentiment classification in analyzing financial articles is promising.

USING SENTIMENT ANALYSIS TO PREDICT GOOGLE STOCK PRICES

Although a general consensus could be drawn that sentiment analysis can measure the mood of public opinions with high accuracy, there are not many studies focusing on the applications of sentiment analysis in predicting the stock market. Out of the studies that did focus on forecasting financial markets through sentiment analysis, the results generally revealed a low correlation between public sentiment and the stock market.

A research study was conducted to explore the influence of news reports on stocks using sentiment analysis in 2013. In his paper, XiaoDong Li, a graduate student at the City University of Hong Kong computer science department, and his colleagues examined the accuracy of using sentiment analysis to predict the Hang Seng Index (HSI). The Hang Seng Index, containing the top 50 companies in Hong Kong, is regarded as the main measure of market performance in the region. As for business articles, Li extracted news from FINET, an archive comprised of articles relating to both individual companies and the Hong Kong market from January 2003 to March 2008 (Li et al. 2014, 16). A dictionary-based method of sentiment analysis is then applied to the news articles. Li utilized both the Harvard IV-4 sentiment dictionary (HVD) and the Loughran-McDonald financial sentiment dictionary (LMD) in his study. Both sentiment dictionaries were compiled manually. The HVD contains over 10,000 words with 15 dimensions to each word while the LMD consists of more than 3,911 words with 6 dimensions. These dimensions include positive/negative connotation, cognitive orientation, motivation, and etc. By cross-referencing the news articles with the dictionaries, Li produced a data chart tallying articles that fit under specific dimensions. Essentially, he converted each article's sentiment into numerical values. Moreover, each article had a time stamp and a tag indicating its relevance to certain companies. This allowed Li to match the news sentiment data with changes in the Hang Seng Index of a certain day. With this dataset, Li extrapolated a correlation to predict changes in the stock price index based on the latest news articles. He found that both dictionary-based sentiment analysis

USING SENTIMENT ANALYSIS TO PREDICT GOOGLE STOCK PRICES

techniques had around the same accuracy of 57% in predicting the rise or fall of the Hang Seng Index. Considering that a random binary guess would yield a 50% accuracy, Li concluded that his model would be 7% better than a random guess.

In 2000, another study used a different approach from Li et al. to investigate this topic. Kenneth L. Fisher, the founder of Fisher Investments, and Meir Statman discovered that there was a negative relationship between investor sentiment and stock price changes (2000, 16). To obtain data about investor sentiment, Fisher surveyed three groups of stock investors: Wall Street strategists, investment news writers, and small individual investors. Respectively, these groups represented the large, medium, and small investors in the stock market. Unlike Li's method to obtain sentiment data, Fisher conducted surveys and questionnaires on each investor group. For instance, Merrill Lynch, an American wealth management group, conducted and provided the surveys on around 15 to 20 Wall Street strategists since September 1985. In the form of questionnaires, each survey measured how *bullish* its subjects are. In stock market terms, to be *bullish* is to have a high inclination of purchasing stocks. Fisher selected data surveyed monthly from September 1985 to July 1998. He then compared each group's sentiment values with movements in the S&P 500 index according to trading days. Three scatter plots (One for each investor group) visualized the dataset and displayed the correlation of the data. Though all three groups had a negative correlation between investor mood and stock price changes, only the correlations for individual investors and Wall Street strategists were statistically significant at the 1 percent level with an adjusted R-squared value of 0.05 and 0.03 respectively. Although the R-squared values suggest that investor sentiment justifies only 3 to 5 percent of S&P 500 returns, Fisher explained that the information could be useful to stock traders. According to Clarke et al. in another study about correlations between the stock market and information, an R-squared value of 0.09 gives stock traders a 5.9 percent higher accuracy in forecasting expected stock

USING SENTIMENT ANALYSIS TO PREDICT GOOGLE STOCK PRICES

return (1989, 31). Essentially, Clarke's research explained the validity and implications of Fisher's results, concluding that even a small R-squared value could give a strategic advantage to investors.

Although the aforementioned studies indicated the potential of using sentiment analysis to estimate changes in stock price indexes, there remains a significant research gap in predicting individual company stocks. In these studies mentioned above, they tended to relate sentiment with stock indexes. In my opinion, focusing on entire indexes, which typically includes a number of company stocks, could possibly be the reason why existing research yielded comparatively low correlations. As asserted by Malcolm Baker, a professor at the Harvard Business School, certain stocks are more subject to changes in public sentiment than others (Baker and Wurgler 2007, 131). For instance, the S&P 500 is an American stock market index containing 500 companies. When there is a notable change in public sentiment, some stocks within the S&P 500 may be impacted significantly while others are less affected by it. The correlation between sentiment and the S&P 500 will be less significant as only parts of the index are affected by the general view of the public.

In my experiment, I hypothesized that a higher level of correlation could be determined from investor sentiment and stocks by looking at only one company. This led to the essential question of my research: how do social media news regarding a specific company impact its stock market prices? To test my hypothesis, I centered my study on the stock return of Google based on online articles about the company. Using Python programming, I scrapped articles about Google from a news website and performed sentiment analysis on each piece. By comparing this data with the stock price change, a best-fit relationship could be derived, possibly leading to higher accuracy in predicting individual stock prices based on public sentiment.

USING SENTIMENT ANALYSIS TO PREDICT GOOGLE STOCK PRICES

Methodology

Google was suitable to be the subject of this study as it fulfilled several requirements. Google, or officially Alphabet Inc., is one of the most valuable companies in the world with an estimated market valuation of 848 billion dollars as of April 2019 (Yahoo! Finance). The company's stock, GOOG, receives an average trade volume of 1,380,563. Google is also one of the most recognized brands in the world (Casella 2017). The company's focus matches the characteristics of stocks more subject to sentiment as shown in a study done by Baker and Wurgler. From their experiment, they concluded that well-recognized companies following the modern business trend are more affected by changes in public sentiment (Baker and Wurgler 2007, 131). Google's renowned reputation in Internet-related services thus fits the criteria well.

Downloading Online Articles

As for extracting news sentiment of a company, Bhardwaj et al. described a reasonable method to obtain news archives using Python programming. By using Beautiful Soup, a Python package for scrapping online content, Bhardwaj et al. were able to fetch Indian news articles on the Internet (Bhardwaj et al. 2015, 89). Bhardwaj's procedure suited my study because it was manageable for a high schooler like me to implement and fulfilled the purpose of retrieving online news content.

For my experiment, I extracted articles about Google from Techmeme, a technology news aggregator. In order to ensure multiple news sources were taken into account, I selected Techmeme as my data provider since it compiles its content from dozens of other online news sites, including the New York Times, TechCrunch, and etc. Similar to Bhardwaj's method, I utilized Beautiful Soup and Selenium, another Python package for automated web

USING SENTIMENT ANALYSIS TO PREDICT GOOGLE STOCK PRICES

browsing, to extract news articles onto my computer. My algorithm in Python worked as follows:

1. Use Selenium to open web browser with the address of `https://www.techmeme.com/search/query?q=google&wm=false` (Search query of “Google” on Techmeme)
2. For each article on the page, use BeautifulSoup to parse article’s text, title, and publishing time.
3. Write each article’s text, title, and publishing time onto a local CSV (Comma-separated values) file on my computer. Each article would be a single row with 3 columns.
4. Use Selenium to flip to next page of results.
5. Loop the above steps for 100 times.

Each results page from Techmeme displays 10 articles at a time. By flipping the page 100 times and recording each article, the above algorithm would theoretically retrieve 1,000 most recent articles about Google. However, nearly half of the articles had alternate formats and layouts, causing it to be inaccessible for BeautifulSoup to parse. As a result, the above procedure yielded 573 articles about Google from Techmeme. The retrieved articles had publishing dates ranging from June 19th, 2018 to March 20th, 2019 with an average of 3.12 articles per day.

Extracting Stock Price

From Yahoo! Finance, I downloaded the recent stock prices of Alphabet Inc. (GOOG). With the news sentiment of Google as the independent variable in my experiment, I used the percent change in stock prices as the dependent variable, which is given by

$$\text{Percent Change} = \frac{\text{Closing Price} - \text{Previous Closing Price}}{\text{Previous Closing Price}}.$$

USING SENTIMENT ANALYSIS TO PREDICT GOOGLE STOCK PRICES

As opposed to other similar studies, I used the percent change between closing prices across trading days, also known as the close-to-close return. Li et al. measured the percent change between the opening price and the closing price in one trading day, or the open-to-close return, in their study. They justified their means since it resolved the issue of non-trading day gaps, where close-to-close return behaves differently over weekends and holidays (Li et al. 2014, 17). However, in my opinion, the open-to-close method only measures the change happening in the market hours. Due to pre-market and after-market trading, the opening price does not necessarily equal to the previous trading day's closing price. News published before the opening time of the market could potentially impact the market's opening price. Thus, I chose the close-to-close return method as it accounted for news not published during the market hours.

Applying Sentiment Analysis

To apply sentiment analysis on each article, I used another Python library, TextBlob. Rather than other text analysis methods, I chose TextBlob because it was feasible for a high school student like me to use. TextBlob is also more efficient in performing sentiment analysis due to its relatively simple algorithm. By inputting raw text, TextBlob is able to return the polarity and subjectivity of the text. The polarity score is a value between -1.0 and 1.0 with -1.0 being very negative and 1.0 being very positive. The subjectivity score is a value within the range of 0.0 and 1.0 where 0.0 is very objective and 1.0 is very subjective. For instance, the sentence "TextBlob is amazing" would yield a polarity of 0.6 and a subjectivity of 0.9.

TextBlob's underlying concept to derive these values is based on Bayes' theorem in conditional probability. Bayes' theorem describes the probability of an outcome occurring given a linked precondition. TextBlob takes into account the likelihood of a phrase being

USING SENTIMENT ANALYSIS TO PREDICT GOOGLE STOCK PRICES

positive or negative based on other information related to it. An example would be to give a high value in the subjectivity score of a phrase with an exclamation mark since exclamation marks typically meant an outburst of emotion. TextBlob adds or subtracts the probability of the text being positive/negative or objective/subjective based on the many characteristics of the passage. After adding the values together, the outcome with the highest probability is declared to be the correct answer. This method, known as the Naïve Bayes classifier, consists of only counting and multiplication, enabling a quick and efficient runtime. Despite the simplistic algorithm, the accuracy of using a Naïve Bayes classifier for sentiment analysis is similar to other text classifiers (Agaian and Kolm 2017, 3).

To match the derived sentiment values from the articles with the stock trading days, I averaged all of the articles' polarity and subjectivity in the same time interval. NASDAQ, the stock exchange handling Alphabet Inc. (GOOG), operates its market regularly from 9:30 AM to 4:00 PM (EST). For my experiment, each time interval is defined as 4:00 PM to the next trading day's 4:00 PM (EST). All in all, the procedures I took to apply TextBlob on the article text and match the sentiment data with the corresponding trading times can be described as the following:

1. Starting with the most recent article, use python to read the article text from the CSV file.
2. Apply TextBlob to retrieve polarity and subjectivity value of 1 article.
3. Repeat Steps 1-2 for every article within the same day and add up the polarity and subjectivity values.
4. When the next article is not of the same day, divide the polarity and subjectivity values with the number of articles in this day (This will derive the average polarity and subjectivity of the articles in that day). Write the starting time, ending time, number of articles in that day, average polarity, average subjectivity, and stock price

USING SENTIMENT ANALYSIS TO PREDICT GOOGLE STOCK PRICES

percent change onto a separate CSV file. Each day would be a single row with 6 columns.

5. Loop the above steps until the last article.

Start Time	End Time	No. of Articles	Avg. Polarity	Avg. Subjectivity	Percent Change
03/19/2019	03/20/2019	4	0.097	0.41	2.1
03/18/2019	03/19/2019	8	0.074	0.42	1.2
03/15/2019	03/18/2019	1	0.089	0.38	-0.017
03/14/2019	03/15/2019	3	0.1	0.41	-0.092
03/13/2019	03/14/2019	3	0.024	0.41	-0.65
03/12/2019	03/13/2019	6	0.13	0.43	0.01
03/11/2019	03/12/2019	2	0.13	0.43	1.5
03/08/2019	03/11/2019	5	0.09	0.41	2.9
03/07/2019	03/08/2019	1	0.14	0.47	-0.086
03/06/2019	03/07/2019	1	0.17	0.47	-1.3
03/05/2019	03/06/2019	8	0.12	0.42	-0.36
03/04/2019	03/05/2019	3	0.091	0.36	1.2
03/01/2019	03/04/2019	2	0.14	0.36	0.6
02/28/2019	03/01/2019	1	0.15	0.42	1.9
02/27/2019	02/28/2019	2	0.11	0.49	0.35
02/26/2019	02/27/2019	1	0.16	0.41	0.083

Sample data

The resulting dataset contains 175 time intervals between trading days spanning from June 19th, 2018 to March 20th, 2019. Each row represents a time interval and contains the sentiment values from the articles published within that interval. The percent change is the close-to-close stock price change. All values were rounded to 2 significant digits. A further analysis of this dataset is explored in the results section of this paper.

Predicting Stock Returns using Machine Learning

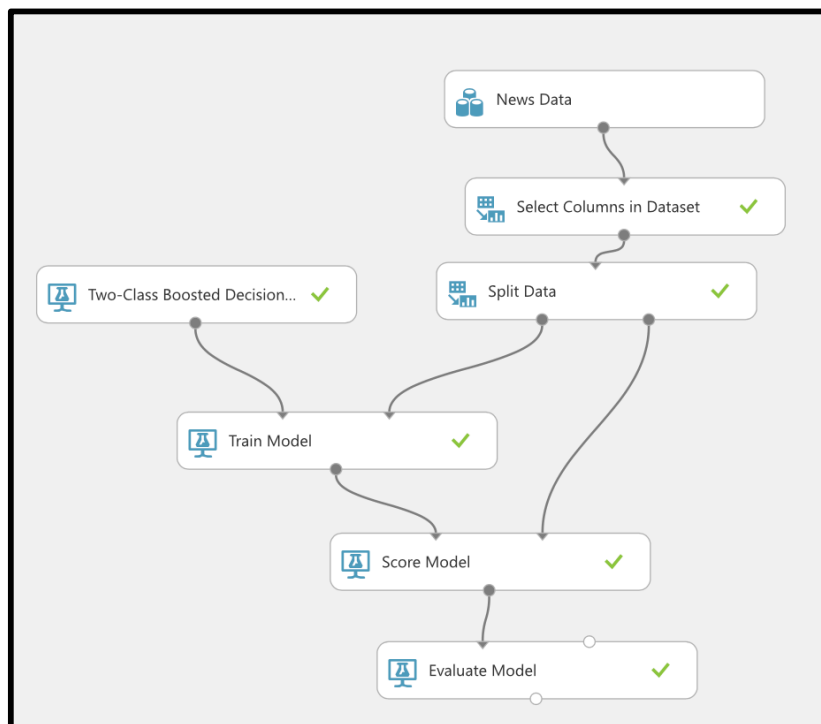
After gathering the news sentiment, I uploaded the compiled data onto Microsoft Azure Machine Learning Studio, a free-to-use cloud computing service for machine learning. Azure was suited for me as it provides a simple drag-and-drop user interface to run machine learning. To predict stock prices from the news sentiment, I employed a supervised machine learning algorithm. According to Andrew Ng, the founder of Landing AI and an adjunct

USING SENTIMENT ANALYSIS TO PREDICT GOOGLE STOCK PRICES

professor of computer science at Stanford University, supervised learning is essentially training the machine with a dataset and asking the machine to then predict values (2012).

This fits my goal of predicting future stock prices based on news sentiment.

There are two general types of supervised machine learning algorithms: regression and classification. Regression models can predict a continuous output of values while classification models can only predict discrete valued outputs (Ng 2012). Both methods would serve the purpose of my experiment as I could predict future price returns using a regression model or I could predict whether the future price return would increase or decrease using a classification model. Ultimately, I chose to implement a classification algorithm due to its higher accuracy and efficiency as compared to regression (Ng 2012). My study's trained machine thus predicted whether the stock return of Google will increase or decrease.



A basic diagram of the machine learning procedure


USING SENTIMENT ANALYSIS TO PREDICT GOOGLE STOCK PRICES

After uploading the dataset onto Azure, I filtered the data by selecting the columns with the number of articles, polarity, and subjectivity as the independent variables and the stock return percent change as the dependent variable. In order to give the machine the ability to predict the stock price change, I needed to train it with a training set. As a result, I split 80% of the data into a training set and the rest 20% as the final test set. I then selected the two-class boosted decision tree classifier as the algorithm for the machine to learn from the training set with 140 rows. Essentially, the machine would learn the pattern in producing binary outputs based on given inputs. In my experiment, the binary output was whether the stock price increased or decreased while the inputs were the number of articles, average polarity, and average subjectivity. The end result is a trained machine. That trained machine, or trained model, was then applied on the test set with the remaining 35 rows and predicted whether the stock price increased or decreased based on the aforementioned independent variables. Lastly, Azure evaluated the accuracy of the trained model in predicting the values of the test set.

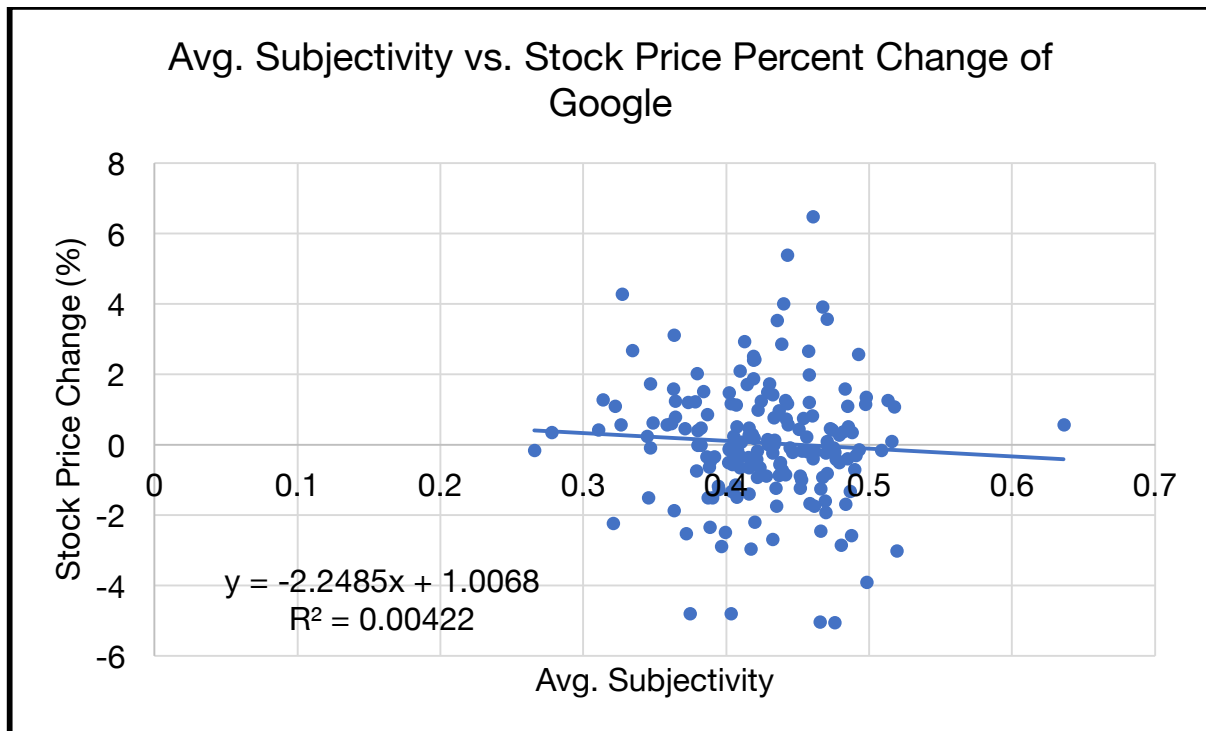
Results

A preliminary examination of the dataset revealed that the correlation between news sentiment and the stock return of Google is low. By graphing the values on Microsoft Excel, there was a negative sloped relationship of -4.30 between the average polarity and stock return. However, the R-squared value of 0.01 meant that the data was not statistically significant at the 0.05 level.

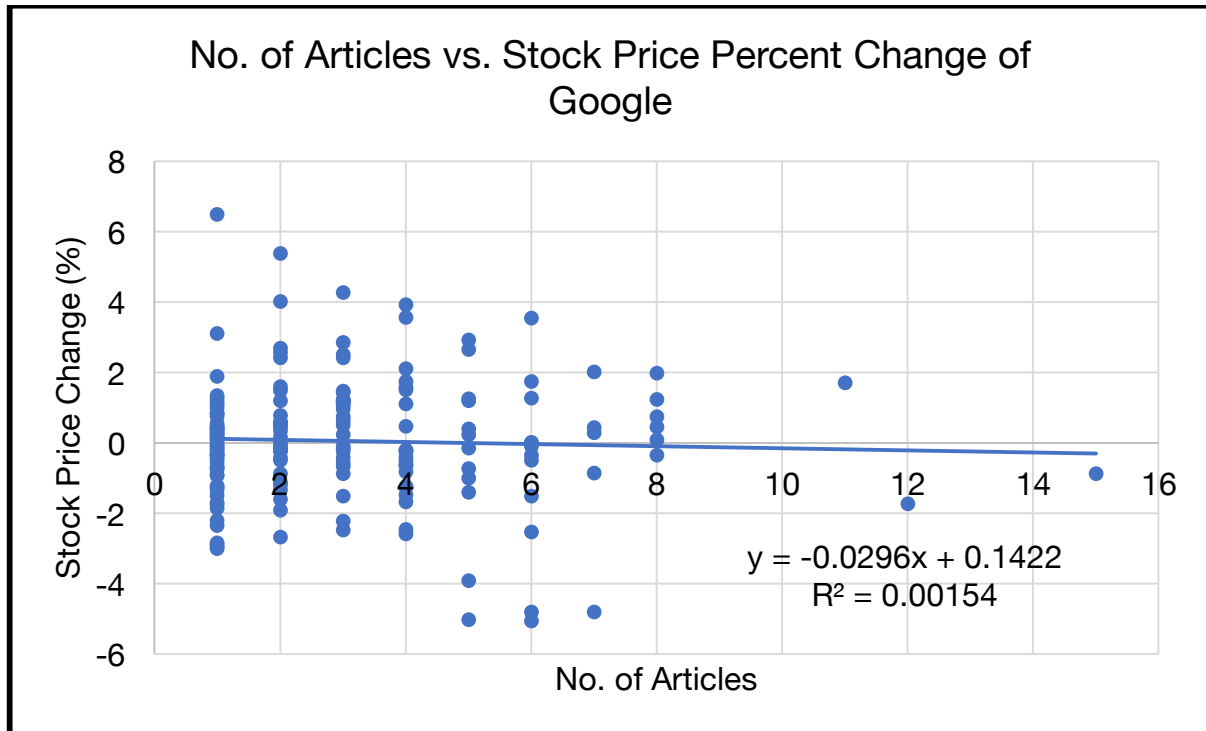
USING SENTIMENT ANALYSIS TO PREDICT GOOGLE STOCK PRICES



When switching the independent variable to the average subjectivity and number of articles, a similar negative sloped trend line could be seen with a smaller R-squared value, causing the data to be statistically insignificant as well.



USING SENTIMENT ANALYSIS TO PREDICT GOOGLE STOCK PRICES



After uploading this dataset to Azure, the resulting prediction accuracy from machine learning was low as well. A sample of the predicted data by the machine is shown below. All values were rounded to 2 significant values.

No. of Articles	Avg. Polarity	Avg. Subjectivity	Change	Scored Labels	Scored Probabilities
2	0.073	0.47	Decrease	Increase	0.67
6	0.13	0.37	Decrease	Increase	1
5	0.12	0.5	Decrease	Increase	0.55
4	0.13	0.39	Decrease	Decrease	0.00018
7	0.12	0.4	Decrease	Decrease	0.12
4	0.085	0.47	Decrease	Decrease	0.0014
2	0.12	0.37	Increase	Increase	0.82
2	0.025	0.27	Decrease	Increase	1
4	0.082	0.41	Decrease	Decrease	0.03
2	0.091	0.42	Increase	Increase	0.9
2	0.08	0.42	Decrease	Decrease	0.05
3	0.12	0.44	Increase	Decrease	0.41
8	0.12	0.42	Decrease	Increase	0.98
2	0.024	0.44	Increase	Increase	0.8
2	0.11	0.47	Decrease	Decrease	0.0086
1	0.12	0.39	Increase	Decrease	0.0084
2	0.12	0.44	Decrease	Decrease	0.00086
2	0.31	0.64	Increase	Decrease	0.11

Sample Data of Predicted Output from Trained Model

USING SENTIMENT ANALYSIS TO PREDICT GOOGLE STOCK PRICES

Unlike the previous scatter graphs, the machine learning algorithm takes into account all 3 features (Number of articles, average polarity, and average subjectivity). The Scored Labels column is the prediction made by the machine of whether the stock price will increase or decrease. The Scored Probabilities column shows the probability of the stock increasing as predicted by the machine. By comparing the predictions with the actual outcome in the Change column, the machine yielded a 51.4% accuracy. Considering that a random guess would be 50% accurate, this prediction model using news sentiment is 1.4% better than a random guess.

Discussion

Implications

The scatter graphs from the dataset implicated a negative relationship between Google's news sentiment and stock return. The slope value of -4.30 implied that a 0.1 increase in news polarity (as calculated by TextBlob) would mean a decrease of 0.43% in stock prices. However, despite the negative slope, there was little to no statistical correlation between Google's news sentiment and stock return since the data revealed a statistically insignificant R-squared value of 0.01. From a logical standpoint, the data's negative slope also appeared counterintuitive. Typically, one would expect an increase in public sentiment to increase stock prices as well. As a result, from both a statistical and logical perspective, no correlation could be determined from investor sentiment and stocks by looking at Google.

The prediction model therefore could not be trained to provide an accurate prediction. With an accuracy of 51.4%, my model focusing on an individual company stock performed worse than other similar studies measuring entire stock indexes. Li et al.'s research, for instance, produced a 57% accuracy in predicting the Hang Seng Index (HSI) (2014, 16). Therefore, a conclusion could be drawn that focusing on the sentiment and stock returns of

USING SENTIMENT ANALYSIS TO PREDICT GOOGLE STOCK PRICES

individual companies, specifically Google, does not improve the accuracy. Since this experiment focused on Google only, the same conclusion cannot be drawn for other companies. New research in this area is thus advised to measure stock indexes or other companies different from Google.

Limitations

Apart from the rejection of my hypothesis, there were limitations in my study worth mentioning. Firstly, the sentiment analysis done on the news articles could potentially be less accurate than expected. Compared to online opinions, news articles tend to convey a more objectivity stance with facts. It would be harder for a machine to judge the polarity of the text. As most previous studies about sentiment analysis were done on opinion-based text, TextBlob's accuracy thus could be lower than expected in this study.

Another possibility is that the general sentiment of the public plays a more important factor in influencing the stock price than the mood of the news regarding a company. In my experiment, I only considered the sentiment of the news about Google. However, the business cycle of the overall economy could skew the effects of news sentiment. Mian et al. support this notion by stating that “stock price sensitivity to good earnings news is higher during high sentiment periods than during periods of low sentiment” (2012, 1357).

I did not address these limitations earlier because it occurred to me late into the research process. Due to the complexity of the coding aspect in my research mythology, I spent much of my research time on that. There were no forewarnings of the low correlation from the data as well. These possible mistakes came to me only after the results came out. Despite these potential limitations, one would expect to see at least some correlation between news sentiment and stock prices. Nevertheless, it is important to consider the statistical insignificance of my results as it could have been caused by these discrepancies. Future

USING SENTIMENT ANALYSIS TO PREDICT GOOGLE STOCK PRICES

research in this field could look at other sentiment analysis algorithms suited for analyzing less biased business news. Taking into account the broad economic state of the country could also lead to better results.

Future Studies

In conclusion, more research should be conducted to address these limitations in my experiment. As seen from the low correlation of my results, news sentiment cannot fully explain the changes in Google's stock prices. Future directions in this area can measure individual companies other than Google to confirm whether this study's conclusions can be applied to other companies as well. External variables other than public sentiment can also be included in future studies to better predict the stock market.

I plan to continue this study by taking into consideration the possible limitations of this experiment. To avoid repetition with research done before, I will still focus on measuring the stock prices of individual companies. However, this time, I will select multiple other companies of varying types as my test subjects. By using different company types, I can analyze which type of company stocks is best predicted using sentiment analysis. Secondly, I will include the general state of the economy as one of the independent variables for my research. Since the stock price responds differently to the news at different economic phases, including a quantitative indicator of the economic state could perhaps derive a better correlation between news sentiment and the stock price. The US Federal Reserve's discount rate, or the "interest rate that the Federal Reserve charges on loans", is an example indicator of whether the United States' economy is in a recessionary or expansionary phase (Ray and Anderson 2015, 262). Lastly, I will compare other sentiment analysis algorithms to validate which one is best at analyzing non-biased business text. By refining the process of this

USING SENTIMENT ANALYSIS TO PREDICT GOOGLE STOCK PRICES

research, we would then be one step closer to truly finding out whether social media news impacts the stock market.

USING SENTIMENT ANALYSIS TO PREDICT GOOGLE STOCK PRICES

References

- Agaian, Sarkis and Petter Kolm. 2017. "Financial Sentiment Analysis Using Machine Learning Techniques." *International Journal of Investment Management and Financial Innovations* 3(1): 1-9. Accessed March 26, 2019.
<https://pdfs.semanticscholar.org/100f/ef8fb36d71b096a57d9f463b3ad11c86930e.pdf>.
- Baker, Malcolm and Jeffrey Wurgler. 2007. "Investor Sentiment in the Stock Market." *The Journal of Economic Perspectives*, Vol. 21, No. 2: 129-151. Accessed February 20, 2019. <https://www.jstor.org/stable/30033721>.
- Bhardwaj, Aditya and Yogendra Narayan and Vanraj and Pawan and Maitreyee Dutta. 2015. "Sentiment Analysis for Indian Stock Market Prediction Using Sensex and Nifty." *Procedia Computer Science* 70: 85-91. Accessed February 13, 2019.
<https://www.sciencedirect.com/science/article/pii/S187705091503207X>.
- Casella, Brett. 2017. "The World's 21 Most Recognized Brand Logos of All Time." Accessed February 26, 2019. <https://www.impactbnd.com/blog/most-recognized-brand-logos-identities>.
- Clarke, Roger G. and Michael T. Fitzgerald and Phillip Berent and Meir Statman. 1989. "Market Timing with Imperfect Information." *Financial Analysts Journal*, Vol. 45, No. 6: 27-36. Accessed April 4, 2019. <https://www.jstor.org/stable/4479275>.
- Fisher, Kenneth L. and Meir Statman. 2000. "Investor Sentiment and Stock Returns." *Financial Analysts Journal*, Vol. 56, No. 2: 16-23. Accessed March 27, 2019.
<https://www.jstor.org/stable/4480229>.
- Jones, Jeff and Lydia Saad. 2016. "Gallup Poll Social Series: Economy and personal Finance." *Gallup*. Accessed March 26, 2019.
http://www.gallup.com/file/poll/191045/Half_Americans_Owns_Stock_Record_Low

USING SENTIMENT ANALYSIS TO PREDICT GOOGLE STOCK PRICES

_160420.pdf?utm_source=link_news9&utm_campaign=item_190883&utm_medium=copy.

Kechaou, Zied and Mohamed Ben Ammar and Adel.M Alimi. 2011. "Improving e-learning with sentiment analysis of users' opinions." *Global Engineering Education Conference*: 1032-1038. Accessed March 26, 2019.

https://www.researchgate.net/profile/Adel_Alimi2/publication/224238696_Improving_e-learning_with_sentiment_analysis_of_users'_opinions/links/02e7e53790c325cf81000000.pdf?inViewer=true&pdfJsDownload=true&disableCoverPage=true&origin=publication_detail.

Li, Xiaodong and Haoran Xie and Li Chen and Jianping Wang and Xiaotie Deng. 2014.

"News impact on stock price return via sentiment analysis." *Knowledge-Based Systems (Elsevier)* (69): 14-23. <https://doi.org/10.1016/j.knosys.2014.04.022>.

Mian, G. Mujtaba and Srinivasan Sankaraguruswamy. 2012. "Investor Sentiment and Stock Market Response to Earnings News." *The Accounting Review*, Vol. 87, No. 4: 1357-1384. Accessed February 27, 2019. <https://www.jstor.org/stable/23246280>.

NasdaqTrader.com. 2019. "Nasdaq Daily Market Summary." Accessed March 26, 2019.

<http://www.nasdaqtrader.com/Trader.aspx?id=DailyMarketSummary>.

Ng, Andrew. "Gradient Descent For Linear Regression." *Coursera*, July, 2012. Video, 10:00.

<https://www.coursera.org/learn/machine-learning/lecture/kCvQc/gradient-descent-for-linear-regression>.

Pang, Bo and Lillian Lee and Shivakumar Vaithyanathan. 2002. "Thumbs up? Sentiment Classification using Machine Learning Techniques." *Empirical Methods in Natural Language Processing*: 79-86. Accessed March 26, 2019.

<http://www.cs.cornell.edu/home/llee/papers/sentiment.pdf>.

Ray, Margaret and David Anderson. 2015. *Krugman's Economics for AP*. New York: Worth

USING SENTIMENT ANALYSIS TO PREDICT GOOGLE STOCK PRICES

Publishers.

Wilson, Robert E. and Samuel D. Gosling and Lindsay T. Graham. 2012. "A Review of Facebook Research in the Social Sciences." *Perspectives on Psychological Science*, Vol. 7, No. 3: 203-220. Accessed March 27, 2019.
<https://www.jstor.org/stable/pdf/41613559.pdf>.

Yahoo! Finance. 2019. "Alphabet Inc. (GOOG)." Accessed January 16, 2019.
<https://finance.yahoo.com/quote/GOOG/>.

AP[®] RESEARCH 2019 SCORING COMMENTARY

Academic Paper

Note: Student samples are quoted verbatim and may contain spelling and grammatical errors.

Sample: B

Score: 5

This paper earned a 5 because it provides an excellent and comprehensive lit review; it is clear how this project fits into the research that has already been done (pages 3–5). The paper poses a narrow and focused question: “how do social media news regarding a specific company impact its stock market prices? To test my hypothesis, I centered my study on the stock return of Google based on online articles about the company.” The response defends the choices made (e.g., page 7: “of stocks more subject to sentiment as shown in a study done by Baker and Wurgler. From their experiment, they concluded that well-recognized companies following the modern business trend are more affected by changes in public sentiment (Baker and Wurgler 2007, 131). Google’s renowned reputation in Internet-related services thus fits the criteria well”). The paper describes the process of data collection in detail (pages 8–12) and discusses their new understanding (which is aligned with their data) with an acknowledgment of limitations of the data sources (page 16). The paper makes clear the limitations and offers suggestions for how to modify the research in the future (pages 17–18).

This paper did not earn a score of 4 because the language and organization of the paper elegantly communicate the ideas therein. The new understanding is justified and situated in comparison to prior studies (page 16). The paper incorporates the implications for future research and the community of practice on page 18.