

# “A Stock-Purchasing Agent from Sentiment Analysis of Twitter”

Alvin Chyan  
[achyan@stanford.edu](mailto:achyan@stanford.edu)

Tim Hsieh  
[timhsieh@stanford.edu](mailto:timhsieh@stanford.edu)

Chris Lengerich  
[ctl51@stanford.edu](mailto:ctl51@stanford.edu)

## Abstract

In recent years, there has been a number of forays into stock market prediction via sentiment analysis of microblogs, most notably in Bollen et al.’s 2010 work “Twitter mood predicts the stock market.” However, their work leaves open a number of important questions, most practically, whether it is possible to make money using signals generated from Twitter. An agent which trades based on signals from Twitter is particularly interesting, as there are a number of features of real-life stock trading which make it substantially different from simply predicting a directional signal for a stock. While obtaining a reliable directional signal is certainly a *sine qua non* for successful trading, calibrating a trading strategy to incorporate risk attitudes as well as real-world parameters such as spread and transaction costs requires additional complexity in the model. In this paper, we attempt to answer those questions by obtaining a solid directional signal of the movement of the Dow Diamonds ETF based on market price changes and Twitter data and subsequently building a stock-trading agent to make simulated trades from those signals.

## Prior Work

There have been a number of prior works using sentiment analysis to predict stock data. Most notably, Bollen et al. used Twitter data to predict the direction of DJIA movement, achieving an accuracy of 87.6% using the self-developed Google Profile of Mood States (GPOMS) and a self-organizing Fuzzy Neural Network [2010]. They found that the “Calm” mood profile yields the best result for stock market prediction. In similar research, Oh and Sheng found that using a combination of manual and an automated “bag-of-words” technique yielded high predictive accuracy for a variety of stock tickers [2011]. Their SMO algorithm was able to achieve a weighted F-measure of 85.3%.

## Methods

### Overview

Our ultimate goal was to create a system that could maximize return on investments. The problem was broken into three components, each of which could be independently validated:

1. Tweet sentiment analysis - Accurately identify the sentiments of tweets, so that, if Bollen et al.’s predictions held true, the directional signal could be improved by the use of mood data
2. Stock movement prediction - Find a model to generate a prediction on the stock market movement given the previous days’ stock changes and the Twitter mood
3. Stock-purchasing agent - Accepting the output of Part 2 as input, determine a specific market order to place with the goal of maximizing investment returns.

The overall design of the system is presented in Figure 1.

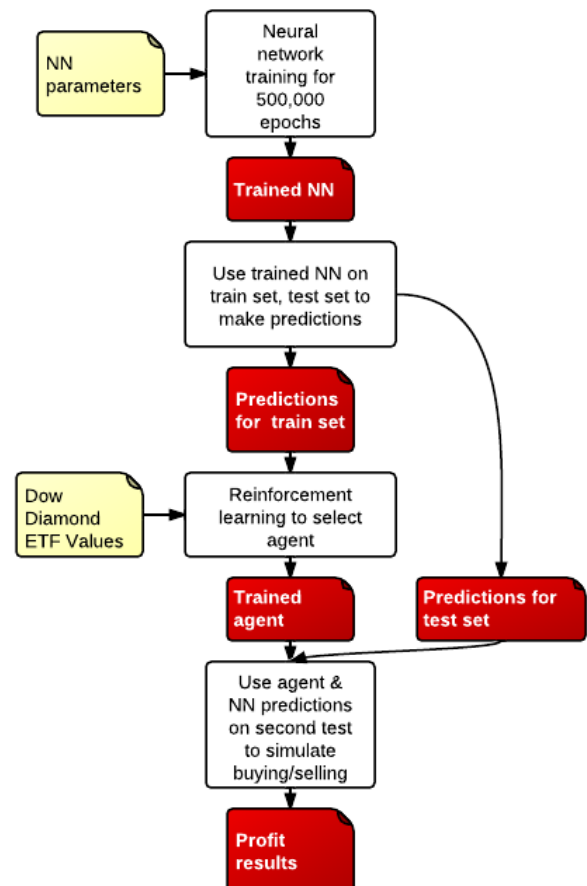


Figure 1: Overall System Design

## Data Sources

### Stock

The daily Dow Diamonds ETF (DIA) opening market price from June 12, 2009 to December 31, 2009 was retrieved from Yahoo! Finance.<sup>1</sup> Figure 2 presents the magnitude of the changes while Figure 3 below presents the normalized DIA in comparison with the “calm” score derived from our tweets. For the time period under consideration, the DIA rose on 60% of the days.<sup>2</sup>

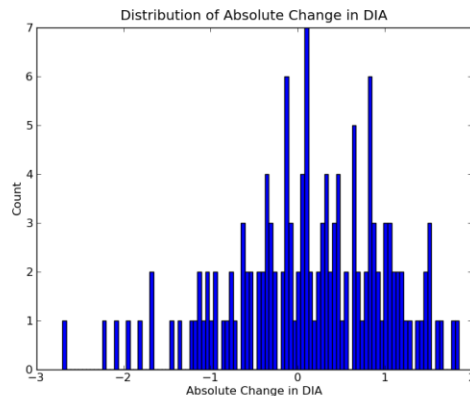


Figure 2: Distribution of Changes in the DIA

### Twitter

A full set of all tweets on Twitter from June 12, 2009 to December 31, 2009 were kindly provided by Stanford’s SNAP laboratory. This encompassed some 476 million tweets [Yang and Leskovec 2011].

### Stock Preprocessing

The direction of change in market opening price across days was scored as a 1/-1 binary classification.<sup>3</sup> This formed the labels for the supervised learning portion of Part II. The normalized market opening price is presented in Figure 3 on the following page.

## Part I: Extracting Mood Score for Each Day

### Creating a Dictionary of “Calm” Words

In Bollen et al., it was observed that only the “calm” sentiment score made a significant contribution to improving the directional accuracy of the DJIA movement signal [2010]. Therefore, a dictionary of “calm” words was compiled and used for a bag-of-words sentiment scoring on data from Twitter. The kernel of the calm dictionary was a base list of words which decrease the score of the “tense” mood state on the Profile of Mood States (POMS). This kernel of words was expanded to include additional words with high statistical correlation to the base words using Google’s 1T 5-gram corpus (some words were manually filtered to avoid words with overloaded sentiment meaning). This resulted in a 22-word dictionary.

### Assigning Mood Scores to Each Day

The following preprocessing steps were used to filter out spam and convert each tweet into a standard form suitable for bag-of-words sentiment analysis. These steps are similar to those documented by Bollen et al. [2009]:

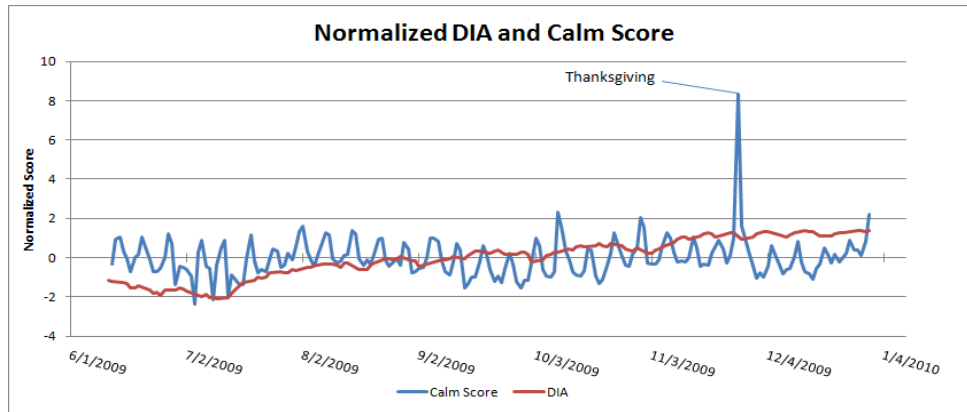
1. Remove all tweets containing “http” and “www” to filter out spam.
2. Keep only tweets which have a high likelihood of reflecting emotional states, such as those containing “I feel”, “I am”, “makes me”, etc.
3. Tokenize each tweet.
4. Remove all non-alpha numeric characters.
5. Convert all letters to lower case from upper case.
6. Porter Stem each word. This step involves a standard pre-processing technique for text, which strips English words of their inflections. For example, “happiness” and “happy” would be converted into the same stem word.

<sup>1</sup> Data for July 15th was missing from the data set provided. Accommodations for this were made in the pre-processing steps.

<sup>2</sup> Note that, as an exchange-traded fund, the DIA does not map perfectly on the rise and fall of the DJIA, which rose 58.5% of the period under consideration. This is a small difference, however, it reflects the fact that the DIA receives influences from other factors besides simply the basket of its incoming stocks.

<sup>3</sup> Note that this indicates that stock changes over the weekends (during which time the markets were not open) had 1 trading day and 2 non-trading days to occur whereas stock changes over a weekday had only 1 trading day to occur. This is a more difficult problem for prediction, however, it is a realistic representation of the decision-making that a trading agent would face. A similar pattern occurred for holidays.

After preprocessing, the selected tweets for a single day were selected and scored using a bag-of-words method. For each occurrence of a word in the “calm dictionary” in one day’s tweets, the “calm score” for that day would be incremented by 1. After calculating the total calm score for each day, the total score for each day was then normalized for volume by dividing the score by the number of sentiment-containing tweets on that day. The normalized DIA and calm score for each day over the period of our data are presented below:



**Figure 3: Normalized DIA and Calm Score**

The calm score plot shows weekly cycles not seen by Bollen et al., which roughly map to low calm scores during the weekdays and higher calm scores during the weekend. This result by itself, in fact, is similar to the findings of Kamvar and Harris, who suggested that the “relaxed” public sentiment increased over the weekends [2011]. Moreover, the calm score is highly correlated with national events, such as Thanksgiving. This indicates that the calm score indeed reflects aspects of the public mood.

The lack of obvious correlation between calm score and DIA value suggests that, if a relationship is indeed present, it may be non-linear. As a result, it may not be feasible to isolate particular words as individually predictive of the DIA value.<sup>4</sup>

## Part II: Stock Price Direction Prediction

Based on the non-linearities of the correlation between the DIA score and the calm score, a neural network was selected as a modeling tool, as neural networks have sufficient power to handle arbitrarily complicated non-linear functions. The pre-processed time series were converted to into the following two lagged time series<sup>5</sup> :

1. T = Each entry contained the previous 3 calendar days of Twitter calm score
2. D = Each entry contained the previous 3 days of change in DIA opening price

Taking these inputs separately and in combination, a neural network was trained to predict the directional change of the stock price over each day. After experimenting with sample sizes as small as 45 days (with ~60% accuracy), it was found that the largest training set possible (n=117) would produce optimal results. Training on 117 days for 500,000 epochs and testing from December 2nd to December 30th on a 6-5-4-1 neural network with sigmoid activation functions, the following accuracies were obtained:

Input Data	Test Set Accuracy (n=20)
D	75%
D,T	80%

**Table 1: NN Directional Prediction Accuracy**

<sup>4</sup> A Naive Bayes multinomial event model was also tested with an expanded vocabulary list, however, the test accuracy for this model was under 55%. This also suggests that the relationship between calm words and sentiment may not follow a linear relationship

<sup>5</sup> Note that because the markets are not open on the weekend, this involved making a prediction for the opening price change from Monday->Tuesday by using the changes from the previous Wednesday->Thursday, Thursday->Friday and Friday->Monday, while the Twitter calm scores from Friday, Saturday and Sunday would be used. We did not investigate the effect that this alignment causes in particular, however, this area would a worthwhile topic of future research.

It is worth noting that that the difference in accuracy of these two test results in a 20-point data set only corresponds to one extra correct classification. However, as noted in Part III, the output from the neural network trained with calm score has a different distribution from the output from the neural network not trained with calm score.

### Part III: Reinforcement Learning Agent

The neural network output a series of signals, with magnitudes that roughly corresponded to the confidence of the network in its prediction. These signals, then, were taken as input to a policy search reinforcement learning agent. The following actions, states and rewards were defined for the RL agent. States are defined as:

$$s = [\text{neural network signal}, \text{number of shares}, \text{amount of cash}, \text{market value of stock}]$$

Taking  $\pi_\theta, 0 \leq \pi_\theta \leq 1$ , as the fraction of the portfolio to be invested in stock, the calculated policy is defined as:

$$\pi_\theta(s) = \frac{1}{1 + e^{-\theta T s}}$$

with  $s'_0=1$ ,  $s'_1$  = the directional signal from the neural network. The neural network output should already encompass the information needed to predict a market change, so the other state parameters are excluded when calculating a policy to help reach convergence of  $\theta$  faster. A sigmoid function was used as it tends to make more aggressive decisions, which would likely receive higher rewards when using a signal that is correct 80% of the time.


The corresponding action was to buy at the ask price or sell at the bid price to reach the desired percentage of investment in stocks (starting from the current percentage of stocks). It was assumed that the agent faced a spread of 0.3%<sup>6</sup> as well as a transaction cost of \$7 per trade<sup>7</sup>. Finally, the reward after taking an action was defined as the difference in the total portfolio value before and after the action is taken, with the stock valued the bid price for the respective day. Through this reinforcement learning algorithm, the RL agent learned the optimal amount of risk to take given its signals.

### Results

The RL algorithm was trained until convergence on the period from July 20, 2009 through December 1, 2009 and tested from December 2, 2009 through December 30, 2009, starting with \$100,000 in cash. The results are presented in Table 2 and the learning curve (for the training set), is presented in Figure 4:

	Theoretical Maximum <sup>8</sup>	Buy and hold <sup>9</sup>	Buy \$5,000 of stock each day, sell on last day	Training on NN Predictions with NN Input = D	Training on NN Predictions with NN Input = D and T
Profit	\$4337	\$114	\$1344	\$2493	\$1329

**Table 2: Profit Under Agent vs. Baseline Strategies**



**Figure 4: Profit Achieved By RL Agent over Training Epochs**

<sup>6</sup> This spread is typical of highly liquid stocks like the DIA

<sup>7</sup> This is typical of online brokerages

<sup>8</sup> The agent, given a perfect signal, goes all in each time, trading only at the opening market bell each day

<sup>9</sup> Go all in the first day, sell all on the last day

## Analysis

Interestingly, even though the accuracy of the directional prediction is improved by adding the sentiment analysis from Twitter, the RL agent trading on the associated signal makes more money without the mood score. While this may be caused by a number of factors, it may be explained by the plotting the output of the neural network versus the actual change that day, as in Figure 5 and Figure 6:




Figure 5: NN Output Without Calm Score

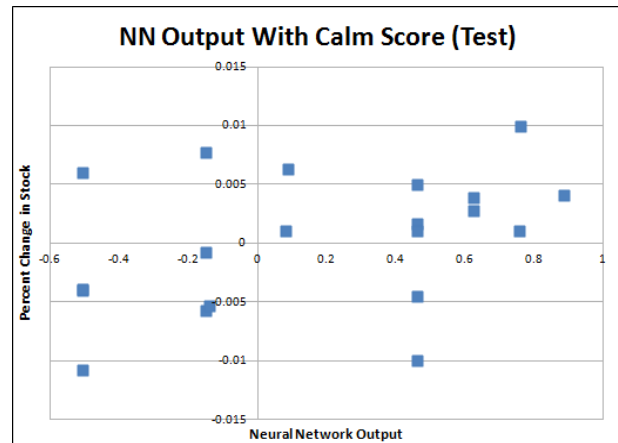


Figure 6: NN Output With Calm Score

For both training runs, the  $\theta^*$  selected sets the decision boundary near zero. Therefore, if a point lies in the first or third quadrant, then this means the directional signal is correct. As shown in the graph above, the directional signal makes five incorrect decisions without the calm score and four with the calm score. However, the incorrect decisions with the calm score have greater magnitude of change. This may contribute to the differences in their respective profits.

Regardless of the ordering of the networks trained with or without mood score, the agent using either of the directional signals makes at least as much money as the largest baseline (\$2493 and \$1329 versus \$1344, respectively). This is a strong indication that the reinforcement learning algorithm worked correctly. Moreover, the agent trained on stock only makes approximately \$1000 more than the highest baseline, indicating that it outperformed some simple investment strategies.

## Conclusion

Twitter calm mood improved the prediction accuracy of the neural network by approximately 5% over the test period. However, this did not appear to increase the total profit of the agent over the test period. To improve performance, it may be desirable to train the system holistically, rather than simply optimizing each step in sequence. This is due to the fact that the NN output used by the agent in RL training had less accuracy than the output used in RL testing. As suggested by the distribution of the NN output, it may also be that an agent with a weaker directional signal but stronger confidence in its correct predictions would be more desirable than the existing agent. Finally, it may be desirable to further vary the RL policy function, or include more parameters than simply one parameter in the RL algorithm. These steps may be considered on the path towards making a better stock-purchasing agent from sentiment analysis of Twitter.

## References

- Bollen, J., Mao, H. and Zeng, X.-J. 2010. Twitter mood predicts the stock market. *Journal of Computational Science* 2(1):1–8
- Bollen, J., Mao, H. and Pepe, A. 2009. Modeling Public Mood and Emotion: Twitter Sentiment and Socio-Economic Phenomena. In *Proceedings of the Fifth International AAAI Conference on Weblogs and Social Media*.
- Kamvar, S.D. and Harris, J. 2011. *We Feel Fine* and Searching the Emotional Web. *ACM International Conference on Web Search and Data Mining (WSDM '11)*.
- Oh, C. and Sheng, O.R.L. Investigating Predictive Power of Stock Microblog Sentiment in Forecasting Future Stock Price Directional Movement. In *ICIS 2011 Proceedings*, Shanghai, China, Dec. 2011.
- Yang, J. and Leskovec, J. 2011. Temporal Variation in Online Media. *ACM International Conference on Web Search and Data Mining (WSDM '11)*.