

BIROn - Birkbeck Institutional Research Online

Mudinas, Andrius and Zhang, Dell and Levene, Mark (2018) Market trend prediction using sentiment analysis: lessons learned and paths forward. In: Workshop on Issues of Sentiment Discovery and Opinion Mining (WISDOM'18), 10 Aug 2018, London, UK.

Downloaded from: <http://eprints.bbk.ac.uk/23962/>

Usage Guidelines:

Please refer to usage guidelines at <http://eprints.bbk.ac.uk/policies.html>

or alternatively

contact lib-eprints@bbk.ac.uk.

Market Trend Prediction using Sentiment Analysis: Lessons Learned and Paths Forward

Anonymous Authors

ABSTRACT

Financial market forecasting is one of the most attractive practical applications of sentiment analysis. In this paper, we investigate the potential of using sentiment *attitudes* (positive vs negative) and also sentiment *emotions* (joy, sadness, etc.) extracted from financial news or tweets to help predict stock price movements. Our extensive experiments using the *Granger-causality* test have revealed that (i) in general sentiment attitudes do not seem to Granger-cause stock price changes; and (ii) while on some specific occasions sentiment emotions do seem to Granger-cause stock price changes, the exhibited pattern is not universal and must be looked at on a case by case basis. Furthermore, it has been observed that at least for certain stocks, integrating sentiment emotions as additional features into the machine learning based market trend prediction model could improve its accuracy.

ACM Reference Format:

Anonymous Authors. 2018. Market Trend Prediction using Sentiment Analysis: Lessons Learned and Paths Forward. In *Proceedings of the 7th KDD Workshop on Issues of Sentiment Discovery and Opinion Mining (WISDOM'18)*. ACM, New York, NY, USA, 10 pages. <https://doi.org/10.1145/nnnnnnn.nnnnnnn>

1 INTRODUCTION

In recent years, a whole industry has been formed around financial market sentiment detection [37, 38]. Traditional financial news/data providers, notably Thomson Reuters and Bloomberg, have started providing commercial sentiment analysis services. As a result, new financial platforms, such as StockTwits¹ which offers sentiment analysis tools, have also emerged. Nowadays, many investment banks and hedge funds are trying to exploit the sentiments of investors to help make better predictions about the financial market. Some of the most prominent financial institutions (including DE Shaw, Two Sigma, and Renaissance Technologies) have been reported to utilise sentiment signals, in addition to structured transactional data (like past prices, historical earnings, and dividends), in their sophisticated machine learning models for algorithmic trading.

In this paper, we aim to re-examine the application of sentiment analysis in the financial domain. Specifically, we try to answer the following research question.

Can market sentiment really help to predict stock price movements?

¹<http://stocktwits.com>

Permission to make digital or hard copies of part or all of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for third-party components of this work must be honored. For all other uses, contact the owner/author(s).

WISDOM'18, August 2018, London, UK

© 2018 Copyright held by the owner/author(s).

ACM ISBN 978-x-xxxx-xxxx-x/YY/MM.

<https://doi.org/10.1145/nnnnnnn.nnnnnnn>

Although our intuition and experience both tell us that sentiment and price are correlated, it is not clear which is the cause and which is the effect. Furthermore, we also have little idea of what exact types of sentiment are really relevant.

Before we embark on our investigation, it is necessary for us to clarify that in this paper, we use the term “sentiment” to describe all kinds of affective states [23, 35] and we draw a distinction between sentiment *attitudes* and sentiment *emotions* following the typology proposed by Scherer [27]. By attitude, we mean the narrow sense of sentiment (as in most research papers on sentiment analysis) – whether people are positive or negative about something. By emotion, we mean the eight “basic emotions” in four opposing pairs – joy-sadness, anger-fear, trust-disgust, and anticipation-surprise, as identified by Plutchik [24].

The rest of this paper is organised as follows. In Section 2, we review related research. In Section 3, we describe the text data sources from which sentiment signals have been extracted: Financial Times (FT) news articles, Reddit WorldNews Channel (RWNC) headlines, and Twitter messages (tweets). In Section 4, we conduct the *Granger-causality* test [11] to find out whether sentiment attitudes and sentiment emotions cause stock price changes, or is it actually the other way around. In Section 5, we carry out extensive experiments to see if a strong baseline model that utilises fifteen technical indicators for market trend prediction can be further enhanced by adding sentiment attitude and/or sentiment emotion features. Finally, in Section 6, we give some concluding remarks and discuss future directions.

The source code for our implemented system is open to the research community².

2 RELATED WORK

The ability to predict price movements on the financial market would offer a lucrative competitive edge over other market participants. Therefore, it is not surprising that this topic has attracted much attention from both academic researchers and industrial practitioners. According to the *efficient market hypothesis* (EMH), it is impossible to “beat the market”, since stock market efficiency always causes existing share prices to incorporate and reflect all relevant market information. However, many people have challenged this claim and declared that it is possible to predict price movements with more than 50% accuracy [13, 26].

A variety of technical approaches to market trend prediction have been proposed in the research literature, ranging from Autoregressive Integrated Moving Average (ARIMA) [22, 34] to ensemble methods [26]. Huang et al. [13] in their work demonstrated the superiority of Support Vector Machines (SVM) in forecasting weekly movement directions of the NIKKEI 225 index, and Lin et al. [16] managed to achieve 70% accuracy by combining decision trees and neural networks. Recent advances in *deep learning* have brought a

²<https://github.com/AndMu/Market-Wisdom>

new wave of methods [5, 10] to this field. In particular, the Long-Short Term Memory (LSTM) recurrent neural network has been shown to be very effective.

Numerous studies have been carried out to understand the intricate relationship between sentiment and price on the financial market. Wang et al. [33] investigated the correlation between stock performance and user sentiment extracted from StockTwits and SeekingAlpha³. Ding et al. [9] proposed a deep learning method for event-driven stock market prediction and achieved nearly 6% improvements on S&P 500 index prediction. Arias et al. [1] investigated whether information extracted from Twitter can improve time series prediction, and found that indeed it could help predict the trend of volatility indices (e.g., VXO, VIX) and historic volatilities of stocks. Bollen et al. [2] in their research identified that some emotion dimensions, extracted from Twitter messages, can be good market trend predictors. Similar to our approach, Deng et al. [7] combined technical analysis with sentiment analysis. However, they only used a limited set of technical indicators together with a generic lexicon-based sentiment analysis model, and attempted to predict future prices using simple regression models.

Twitter market sentiment analysis is also related to the problem of *stance detection* (SD) [28]. As defined by Mohammad et al. [18], a typical sentiment detection system classifies the text into positive, negative or neutral categories, while in SD the task is to detect the text that is favourable or unfavourable to a specific given target. Most of the existing research on SD is focused on the area of politics [14, 15, 30]. Financial market participants also often express strong stances towards particular stocks (which can be divided into the so-called “bulls” and “bears”). However, there are non-trivial differences between political sentiment and market sentiment, as the financial market is usually more cyclical and dynamic, has different sentiment drivers, and can be impacted by various external factors (e.g., company performances and geopolitical events). Moreover, market sentiment extracted from news articles rather than social media would exhibit different characteristics: the former is less about the authors’ stance but more about the facts and the interpretation of events in a significantly richer context.

The sentiment analysis system used in our experiments is a publicly-available one called *pSenti*⁴ [20, 21] which is equipped with a pre-compiled built-in financial-domain sentiment lexicon. When carrying out sentiment analysis experiments, we followed the procedure outlined by Mudinas et al. [21] in which domain-specific word embeddings would be constructed and domain-specific sentiment lexicons would be induced. Our experiments on several publicly available datasets have confirmed that consistent with the reported results in the original paper, this approach could achieve around 80% sentiment classification accuracy for Twitter messages and slightly higher accuracy for longer texts such as Financial Times (FT) news articles.

3 DATA

To obtain relevant sentiment signals, we have collected three Financial Times (FT)⁵ datasets covering different time periods (see Table 1). In addition, we have obtained a large set of historical news

Source	From	To	Count
Financial Times I	2011-04-01	2011-12-25	11978
Financial Times II	2014-04-01	2014-10-26	9731
Financial Times III	2014-10-26	2015-03-08	6037
Reddit	2008-06-08	2016-07-01	76600
Twitter	2014-05-01	2015-02-01	1145784

Table 1: Financial market datasets used in our experiments.

headlines from Reddit’s WorldNews Channel (RWNC): for each date in the time period we picked the top 25 headlines ranked by Reddit users’ votes. Moreover, we have also gathered from Twitter a large collection of financial tweets, which contain in their text one or more “cashtags”. A cashtag is simply a ‘\$’ sign followed by a stock symbol (ticker). For example, the cashtag for the company Apple Inc., whose ticker is AAPL on the stock market, would be \$AAPL. Here we have only collected the tweets mentioning stocks from the S&P 500 index.

For the stock price data, we have used the *end of day* (EOD) adjusted close price according to the Dow Jones Industrial Average (DJIA) index. In our experiments, we have focused on several representative companies, Apple (AAPL), Google (GOOGL), Hewlett-Packard (HPQ), and JPMorgan Chase & Co. (JPM), as well as a couple of the most liquid FX currency pairs, EUR/USD and GBP/USD. All such financial market data were acquired from public datasets published by Quandl⁶, Kaggle⁷, and Bloomberg⁸.

4 CAUSALITY

To verify whether market sentiments can indeed be useful for predicting stock price movements, we started the investigation with *Granger-causality* test [11] which is a time series data-driven method for identifying causality based on a statistical hypothesis test that determines whether one time series is instrumental in forecasting the other. Granger-causality test has been widely accepted in econometrics as a technique to discover causality in time series data. In the sense of Granger-causality, x is a cause of y if it is instrumental in forecasting y , where ‘instrumental’ means that x can be used to increase the accuracy of y ’s prediction compared with considering only the past values of y itself. Essentially, Granger-causality test is a *null hypothesis significance test* (NHST): the null hypothesis is that the lagged x -values do not explain the variation in y . If the p -value given by the test is less than 0.10, we would be able to reject the null hypothesis and claim that x indeed Granger-causes y .

Through our experiments, we try to find the answers to two questions: *does market sentiment cause changes in stock price*, and conversely, *does stock price cause changes in market sentiment*.

4.1 Time Series

Before performing causality tests, it is necessary to ensure that both time series are *stationary*, because otherwise the results can lead to spurious causality [12]. The stationarity check is typically done by analysing the autocorrelation (ACF) and partial autocorrelation

³<https://seekingalpha.com/>

⁴<https://github.com/AndMu/Wikiled.Sentiment>

⁵<http://www.ft.com>

⁶<https://www.quandl.com>

⁷<https://www.kaggle.com>

⁸<https://www.bloomberg.com/>

Stock	Model	Lag	Attitude ⇒Price	Price⇒ Attitude
S&P 500	Standard	1	0.1929	0.1105
		2	0.2611	0.0780
	Temporal	1	0.2689	0.0495
		2	0.1692	0.0940
APPL	Standard	1	0.7351	0.4253
		2	0.9117	0.6426
	Temporal	1	0.9478	0.6725
		2	0.9715	0.8245
GOOGL	Standard	1	0.5285	0.4035
		2	0.8075	0.0418
	Temporal	1	0.6920	0.5388
		2	0.8516	0.0422
HPQ	Standard	1	0.1534	0.3996
		2	0.1877	0.5322
	Temporal	1	0.4069	0.0836
		2	0.5097	0.1180
JPM	Standard	1	0.8991	0.0461
		2	0.9963	0.0435
	Temporal	1	0.9437	0.1204
		2	0.7722	0.2720

Table 2: Sentiment attitudes Granger-causality on the FT I dataset.

(PACF) functions, and performing the Ljung-Box [17] or the augmented Dickey-Fuller (ADF) [8] t -statistic tests. Using the above methods, all our time series were verified to be stationary. Fig. 1 shows the stationarity check results of the FT-I and DJIA market datasets.

The next step in our investigation into the relationship between sentiment and price time series is to look at their *cross-correlation function* (CCF). Although “correlation does not imply causation”, it is frequently used as a test to discover possible causal relationship from data. In Fig. 2 we present the cross-correlation analysis results for the S&P 500 index on all three FT news and RWNC headlines datasets. In the first dataset (see Fig. 2a) we have a strong CCF between sentiment attitudes and stock prices, in the second (see Fig. 2b) the CCF is significant in the lower left and upper right quadrants. However, in the FT III (see Fig. 2c) dataset, the CCF is not significant (below the confidence threshold). This seems to suggest that the relationship between market sentiments and stock prices can be quite complex and may exist only in some of the time periods. It is unsurprising that the financial market exhibited different behaviours in different time periods. As shown in Fig. 3, from 2011 to mid-2013 we had a volatile market without a clear trend, whereas from 2013 until 2015 we saw a strong *bull* run with continual rising prices. Then we calculated the CCF between the sentiment attitudes found in RWNC headlines and the index prices for a longer time span from 2008 to 2016 (see Fig. 2d), but still could not detect any long-term correlation.

Most of the real-life automated trading systems need to make **BUY** or **SELL** decisions for the given stocks. Therefore, from the trading perspective, the actual price of a stock is less crucial, and the profit relies on the price changes (often measured in percentages). Similar to the previous experiments on the cross-correlation

between sentiment attitudes and stock prices (see Fig. 2), additional experiments on the cross-correlation between sentiments and stock price changes were performed (see Fig. 4). Contrary to the previous results, we found correlations in all the FT news articles (see Figs. 4a to 4c) and the RWNC headlines (see Fig. 4d). This suggests that the percentage changes of stock prices would have higher predictability than stock prices themselves. However, the correlations are often only present with a substantial lag. Therefore it is still valid that sentiment attitudes are unlikely to be useful in market trend prediction.

4.2 Experimental Setup

To further analyse the relationship, we performed a set of *Granger-causality* tests for the S&P 500 index and four selected stocks using the FT I dataset, in which we detected the strongest correlation. The lags in the causality tests were set to just one or two days, considering that the financial market usually reacts to relevant news events almost instantaneously.

The sentiment analysis was performed using both a standard model and an enhanced temporal model. In the latter, we associate the sentiments with the corresponding temporal orientations by labelling each sentence with one of the four temporal-categories (past/present/future/unknown) and calculate the sentiment strength accordingly. Thus, the $pSenti$ sentiment outputs would be generalised from a single dimension into one of those four temporal-categories. The determination of temporal orientation could be done by two different methods. One method is to identify temporal expressions using the *SuTime* temporal tagger [4]. *SuTime* is a rule-based tagger built on regular expression patterns to recognise and normalise temporal expressions in the form of TIMEX3 in English text. TIMEX3 is part of the TimeML annotation language [25] for marking up events, times, and their temporal relations in documents. It recognises both absolute time (such as “January 12, 2000”) and relative time (such as “next month”). The relative time can be then transformed into the corresponding absolute time using the underlying document creation time. For example, in the sentence “I hope next year they will release an improved version”, we would identify “improved” as a positive sentiment attitude with “next year” as its time point, so the text would lead to a positive *future* sentiment feature. Another method, which we also employed in the temporal model, is to use the tense of the sentence as the clue to determine which temporal category should be assigned to the sentiment found in the sentence. Intuitively, only the sentiments about the present and the future value of the stock would have significant impacts on its price. Therefore, we would filter out all sentiment scores with the *past* tag.

In each of our causality test experiments, two competing hypotheses would be examined: market sentiments cause stock price changes and vice versa.

4.3 Experimental Results

The results obtained from the experiments (see Table 2) show a mixed picture. In all the experiments, we failed to discover any sign that sentiment attitudes Granger-cause stock price changes, which would suggest that in general sentiment attitudes probably cannot be useful for the prediction of stock price movements. However, in many cases, we found that the opposite was true — stock price

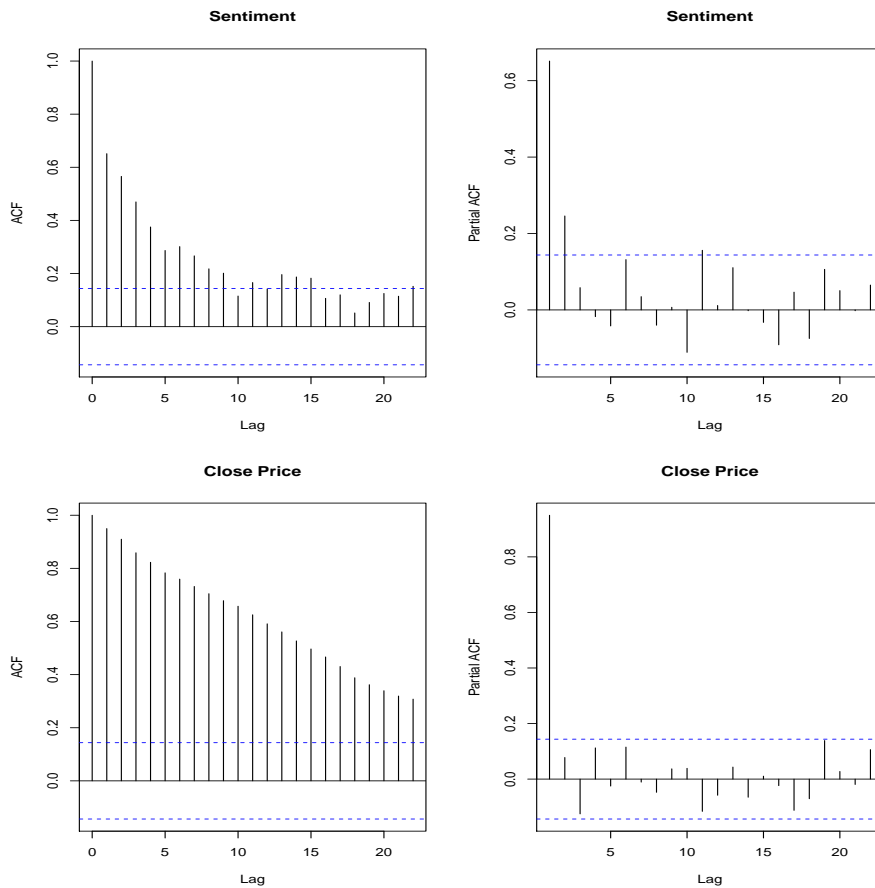


Figure 1: Stationary analysis for DJIA close prices on the FT I dataset.

changes Granger-cause sentiment attitudes in the news, with the strongest causality found using the temporal sentiment analysis model. The individual stocks also produced mixed results, with each company behaving differently. For the Apple stock, we failed to detect any causality. For the Google stock, we identified that the prices would Granger-cause sentiment attitudes, but only with a two-day lag. For the HP stock, we detected causality only in temporal sentiment and only with a one-day lag. For the JPM stock, we found causality using standard sentiment, but it was absent using temporal sentiment. It is difficult to draw general conclusions from such varying results. According to the Granger-causality test with a one-day or two-day lag, sentiment attitudes do not seem to be useful for predicting stock price movements. However, the opposite seems to be true – the sentiment attitudes should be predicted using stock price movements. It is still possible that the Granger-causality from sentiment attitudes to stock price changes is present at a finer time granularity (e.g., minutes), but we are unable to perform such an analysis using our current datasets.

Bollen et al. [2] attempted to predict the behaviour of the stock market by measuring the sentiment emotion of people on Twitter and identified that some of the emotion dimensions have predictive power. To verify their findings, we employed a similar model based on Plutchik’s emotion dimensions extracted using the NRC

Emotion	Lag	Standard		Temporal	
		Emotion ⇒Price	Price⇒ Emotion	Emotion ⇒Price	Price⇒ Emotion
anger	1	0.3815	0.6299	0.2555	0.4155
	2	0.3402	0.9153	0.3097	0.6886
anticipation	1	0.5320	0.2650	0.9216	0.9389
	2	0.4989	0.5765	0.4930	0.7173
disgust	1	0.6668	0.0688	0.2482	0.2031
	2	0.7166	0.3118	0.1160	0.2852
fear	1	0.5821	0.1255	0.8698	0.0591
	2	0.8934	0.2601	0.9888	0.1604
joy	1	0.6972	0.5549	0.3521	0.1530
	2	0.5567	0.8451	0.4045	0.4089
sadness	1	0.3885	0.1067	0.0258	0.1019
	2	0.6166	0.2027	0.0983	0.1423
surprise	1	0.5866	0.7022	0.3830	0.2315
	2	0.9802	0.8414	0.8445	0.3838
trust	1	0.9983	0.6892	0.9490	0.1124
	2	0.5534	0.8523	0.9586	0.2239

Table 3: Sentiment emotions Granger-causality: S&P 500.




Figure 2: The cross-correlation between sentiment attitudes and S&P 500 prices.



Figure 3: The market close price changes (%).

sentiment lexicon [19] and *pSenti*. In the S&P 500 index analysis (see Table 3), we found that only sadness could Granger-cause stock price changes, which is different from the results of Bollen et al. [2]. Such a discrepancy might be explained by the fact that Bollen et al. [2] used different emotion dimensions, lexicons, and a

different time period in their analysis. An interesting finding which we have obtained from the experimental results is that some individual stocks like HP (see Table 6) and JPM (see Table 7) have significantly more emotion dimensions with predictive power than others. It could be seen in those cases that some emotion dimension

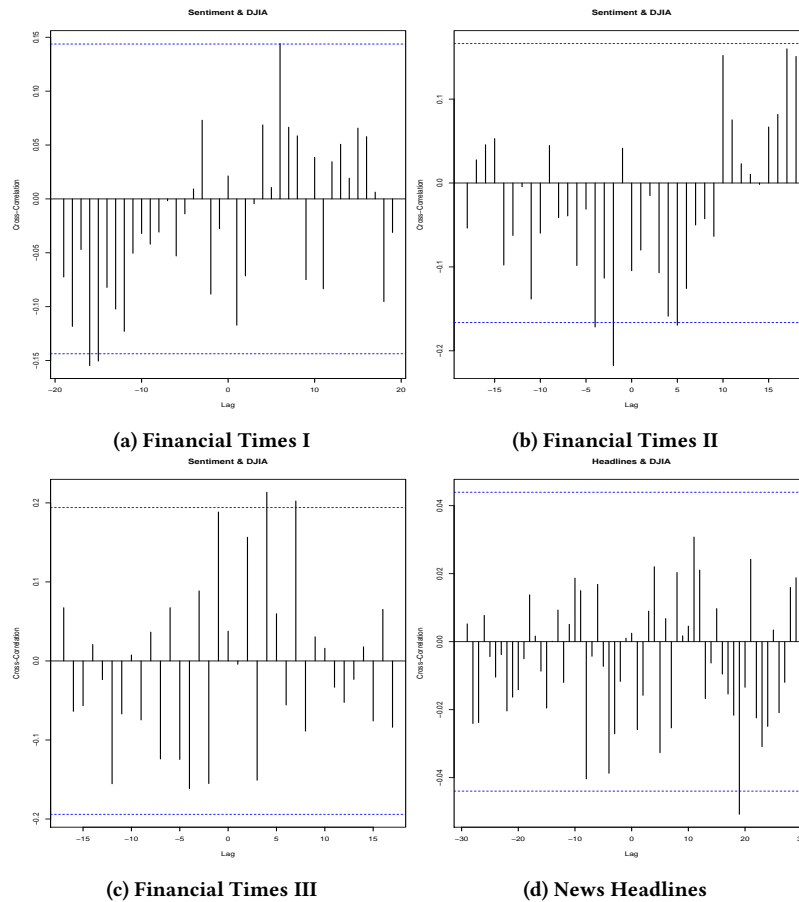


Figure 4: The cross-correlation between sentiment attitudes and S&P 500 price changes.

other than sadness, including surprise, fear, joy and trust, also demonstrated predictive power. On both Google (see Table 5) and Apple (see Table 4) stock price data, we failed to find any emotion causality on their stock price. These results indicate that even if in some cases there is substantial Granger-causality from sentiment emotions to stock price changes, it is not a general pattern, and should be looked at on a case by case basis. To find out why that is happening, it would be necessary to perform further investigation, which is beyond the scope of this paper.

5 PREDICTION

The causality analysis in Section 4 has revealed that in some cases sentiment emotions could be good indicators for stock price changes. In the next set of experiments, we would like to investigate how sentiment attitudes and/or sentiment emotions could be exploited in a machine learning model for market trend prediction to improve its accuracy.

Basically, there are two types of stock market analysis: fundamental and technical. The former evaluates a stock based on its corresponding company' business performance, whereas the latter evaluates a stock based on its volume and price on the financial market as measured by a number of so-called *technical indicators*.

Emotion	Lag	Standard		Temporal	
		Emotion ⇒ Price	Price ⇒ Emotion	Emotion ⇒ Price	Price ⇒ Emotion
anger	1	0.9452	0.3512	0.6490	0.2352
	2	0.9851	0.4367	0.7703	0.1461
anticipation	1	0.5237	0.8272	0.3032	0.1245
	2	0.6368	0.3331	0.595	0.1518
disgust	1	0.2412	0.4128	0.1376	0.9851
	2	0.5154	0.5877	0.3392	0.3130
fear	1	0.3717	0.0867	0.2727	0.5577
	2	0.5609	0.1698	0.4139	0.1114
joy	1	0.3301	0.9946	0.7916	0.2580
	2	0.6657	0.5264	0.9843	0.4312
sadness	1	0.2217	0.8139	0.2280	0.6620
	2	0.1669	0.4266	0.1710	0.3245
surprise	1	0.9413	0.1960	0.1083	0.6093
	2	0.9733	0.2433	0.2033	0.2609
trust	1	0.5663	0.8439	0.3219	0.3539
	2	0.8760	0.5520	0.4473	0.2608

Table 4: Sentiment emotions Granger-causality: APPL.

Emotion	Lag	Standard		Temporal	
		Emotion ⇒Price	Price⇒ Emotion	Emotion ⇒Price	Price⇒ Emotion
anger	1	0.2706	0.4460	0.4420	0.1530
	2	0.1709	0.7454	0.2457	0.2677
anticipation	1	0.1137	0.1951	0.2720	0.4348
	2	0.1487	0.4839	0.3363	0.7986
disgust	1	0.4459	0.3250	0.4000	0.5865
	2	0.7031	0.4294	0.6608	0.8880
fear	1	0.3362	0.1020	0.2874	0.1211
	2	0.2763	0.0757	0.3011	0.1765
joy	1	0.4718	0.0417	0.8350	0.0959
	2	0.7855	0.0998	0.9755	0.2282
sadness	1	0.4184	0.1316	0.3917	0.1782
	2	0.6599	0.1236	0.5286	0.3844
surprise	1	0.6551	0.0606	0.6869	0.0755
	2	0.7604	0.1166	0.6626	0.2156
trust	1	0.5008	0.0727	0.7541	0.0680
	2	0.5991	0.1302	0.8334	0.1052

Table 5: Sentiment emotions Granger-causality: GOOGL.

Emotion	Lag	Standard		Temporal	
		Emotion ⇒Price	Price⇒ Emotion	Emotion ⇒Price	Price⇒ Emotion
anger	1	0.1788	0.0488	0.1796	0.2349
	2	0.4903	0.1223	0.3155	0.3713
anticipation	1	0.3893	0.1360	0.1389	0.5989
	2	0.7729	0.3145	0.3389	0.4732
disgust	1	0.2168	0.1260	0.2208	0.2267
	2	0.5297	0.2913	0.3611	0.1637
fear	1	0.0298	0.1565	0.0214	0.2072
	2	0.1173	0.2495	0.0210	0.1187
joy	1	0.3417	0.2169	0.1073	0.9574
	2	0.8544	0.3905	0.3293	0.9086
sadness	1	0.6079	0.3038	0.3985	0.5781
	2	0.9297	0.5856	0.4495	0.4194
surprise	1	0.1351	0.0303	0.0498	0.1145
	2	0.4296	0.0593	0.0850	0.0461
trust	1	0.0991	0.2218	0.0458	0.6664
	2	0.1232	0.2066	0.0165	0.6645

Table 7: Sentiment emotions Granger-causality: JPM

Emotion	Lag	Standard		Temporal	
		Emotion ⇒Price	Price⇒ Emotion	Emotion ⇒Price	Price⇒ Emotion
anger	1	0.2129	0.8639	0.1300	0.9466
	2	0.4084	0.9521	0.2234	0.9689
anticipation	1	0.0757	0.6288	0.1316	0.7853
	2	0.2279	0.9059	0.3371	0.8986
disgust	1	0.4868	0.8126	0.2001	0.4536
	2	0.3803	0.9353	0.2252	0.6722
fear	1	0.2679	0.4841	0.1214	0.8193
	2	0.5361	0.4741	0.2371	0.9319
joy	1	0.0399	0.8186	0.0902	0.6261
	2	0.1255	0.8945	0.2410	0.7273
Sadness	1	0.0106	0.8669	0.0110	0.9208
	2	0.0416	0.9365	0.0388	0.8456
surprise	1	0.0217	0.6825	0.0010	0.3890
	2	0.0759	0.7830	0.0064	0.3034
trust	1	0.0447	0.8620	0.0766	0.7693
	2	0.1340	0.9034	0.2158	0.7948

Table 6: Sentiment emotions Granger-causality: HPQ.

Both types of analysis generate trading signals, which would be monitored by human traders or automated trading systems who use that information to execute trades. In our experiments, only technical analysis has been utilised. It is likely that incorporating fundamental analysis and employing more technical indicators would improve the predictive model's performance. However, our research objective is not to create the optimal market trend prediction system, but to analyse and understand the predictive power of sentiments on the financial market. For this purpose, a baseline model with several common technical indicators should be good enough.

5.1 Baseline

We first built a baseline machine learning model to predict stock price changes with a number of selected technical indicators, and then tried to incorporate additional sentiment-based features, i.e., sentiment attitudes and sentiment emotions.

In order to construct a decent baseline model, we made use of ten common technical indicators which led to a total of fifteen features as follows.

- **Moving Averages (MA).** A moving average is frequently defined as a support or resistance level. Many basic trading strategies are centred around breaking support and resistance levels. In a rising market, a 50-day, 100-day or 200-day moving average may act as a support level, and in a falling market as resistance. We calculated 50-day, 100-day and 200-day moving averages and included each of them as a feature.
- **Williams %R.** This indicator was proposed by Larry Williams to detect when a stock was overbought or oversold. It tells us how the current price compares to the highest price over the past period (10 days).
- **Momentum (MOM).** This indicator measures how the price changed over the last N trading days. We used two momentum-based features, one-day momentum and five-day momentum.
- **Relative Strength Index (RSI).** This is yet another indicator to find overbought and oversold stocks. It compares the magnitude of gains and losses over a specified period. We used the most common 14 days period.
- **Moving Average Convergence Divergence (MACD).** This is one of the most effective momentum indicators, which shows the relationship between two moving averages. It generates three features: MACD, signal, and histogram values.
- **Bollinger Bands** is one of the most widely used technical indicators. It was developed and introduced in the 1980s by the famous

technical trader John Bollinger. It represents two standard deviations away from a simple moving average, and can thus help price pattern recognition.

- Commodity Channel Index (CCI) is another a momentum indicator introduced by Donald Lambert in 1980. This indicator can help to identify a new trend or warn of extreme conditions by detecting overbought and oversold stocks. Its normal movement is in the range from -100 to +100, so going beyond this range is considered a BUY/SELL signal.
- Average Directional Index (ADX) is a non-directional indicator which quantifies the price trend strength using values from 0 to 100. It is useful for identifying strong price trends.
- Triple Exponential Moving Average (TEMA) was developed by Patrick Mulloy and first published in 1994. It serves as a trend indicator, and in contrast to moving averages it does not have the associated lag.
- Average True Range (ATR) is a non-directional volatility indicator developed by Wilder [36]. The stocks and indexes with higher volatility typically have higher ATR.

The features were all normalised to zero mean and unit variance in advance.

In our context, the machine learning model is just a binary classifier that generates two kinds of signals: **BUY** (+1) and **SELL** (-1). It aims to predict whether or not the stock's price, n days in the future, will be higher (+1) or lower (-1) than today's price. In the preliminary experiments, we tried to find out which machine learning algorithm would perform best and how far into the future the model would be able to predict.

Following the research literature in this area [5, 10, 13], we evaluated the two most popular machine learning approaches to market trend prediction, SVM (with the RBF kernel) and LSTM recurrent neural network. Each dataset was randomly divided into two sets: 2/3 for training and 1/3 for testing. The parameters of the SVM and LSTM algorithms were set via grid search on the training set. The final LSTM model consists of a single LSTM layer with 400 units and utilises a drop-out rate of 0.5 [29, 32].

It is common for such market trend prediction models to use a time lag of a few days and by doing so avoid short-term price volatility [6]. In our experiments, we tried both three-day and five day lags. Similar to the previous studies by Cao and Tay [3] and Thomason [31], using five-day lags was found to be optimal.

The preliminary experimental results as shown in Table 8 indicate that SVM outperformed LSTM on all the datasets. The F_1 scores suggest that LSTM often favoured the positive class over the negative class and produced unbalanced results. The reason could be that the size of the dataset is relatively small: there are 670 data points in the analysed time period 2011-2015. Contrary to LSTM, SVM always yielded balanced and stable results.

In the end, SVM with a five-day lag was selected as the baseline model which produced a reasonable accuracy of around 70% and similar F_1 scores for both classes.

5.2 Using Sentiment Signals in News

In the next set of experiments, we evaluated the predictive power of sentiments extracted from financial news articles/headlines. The

time granularity here is a single day, i.e., all sentiment-based features (including both attitudes and emotions) would be aggregated by calculating their daily averages. If there was no sentiment information available on that day, the value zero would be assigned to the corresponding sentiment features.

The proposed new model consists of the same technical indicator features as in the baseline plus nine additional sentiment-based features:

- Sentiment attitudes. The average daily sentiment attitudes, extracted using $pSenti$, with values in the range from -1 to +1.
- Sentiment emotions in eight categories: anger, anticipation, disgust, fear, joy, sadness, surprise, and trust, with values being the normalised occurrence frequency.

Sentiment attitudes and emotions were extracted from the FT news articles and the RWNC headlines in the time period from 2011 to 2015. The experimental results as shown in Table 9 indicate that incorporating sentiment attitudes and sentiment emotions from the headlines actually had a negative impact on the predictive model's performance. This is consistent with the previous section, in which no correlation or causality link was established between headlines sentiments and stock prices. It might be explained by the fact that headlines are very short text snippets and therefore provide little chance for us to reliably detect sentiment attitudes and sentiment emotions. The sentiments extracted from FT news articles painted a quite different picture. The sentiment enriched model outperformed the baseline model in half of the scenarios: it demonstrated slightly better results for DJIA, JPM, and EUR/USD, but slightly worse results for AAPL and GBP/USD. These experimental results are consistent with the previous section and confirm again that for some stocks sentiment emotions could be used to improve the baseline model for market trend prediction.

5.3 Using Sentiment Signals in Tweets

In the last set of experiments, we created the enriched model based on sentiment attitudes and sentiment emotions extracted from financial tweets. The time period of the Twitter messages dataset is significantly shorter: only from 2014 to 2015. Consequently, the experiments were performed on a shorter time period with only 275 data points. In this time period, almost all stock prices were continually rising (see Fig. 3). Such a so-called bull run makes it even more difficult to assess a predictive model's performance, as any basic strategy like buy and hold would be a winning strategy.

Let us consider three different scenarios. In the first scenario ("all+attitude+emotion"), both sentiment attitudes and sentiment emotions were extracted from all financial tweets and used as additional features. This allowed us to verify how useful sentiment information is for market trend prediction. In the second scenario ("all+emotion"), only those eight sentiment emotions were used as additional features. This provided an opportunity to validate the usefulness of sentiment emotions alone. For the last scenario ("filtering+emotion"), only the Twitter messages (tweets) mentioning the company of our interest were utilised to extract sentiment emotions as additional features.

The experimental results as shown in Table 10 indicate that most of the time, the baseline model would actually outperform the expanded model with sentiment attitudes, sentiment emotions, or

Type	Method	3-day ahead			5-day ahead		
		Acc	F_1^{up}	F_1^{down}	Acc	F_1^{up}	F_1^{down}
DJIA	SVM	0.616	0.738	0.282	0.700	0.754	0.615
	LSTM	0.559	0.706	0.120	0.585	0.728	0.127
AAPL	SVM	0.577	0.676	0.391	0.685	0.723	0.634
	LSTM	0.547	0.693	0.138	0.521	0.641	0.282
JPM	SVM	0.677	0.747	0.552	0.673	0.733	0.578
	LSTM	0.541	0.665	0.269	0.573	0.676	0.373
EUR/USD	SVM	0.642	0.607	0.672	0.671	0.620	0.710
	LSTM	0.509	0.423	0.572	0.563	0.370	0.665
GBP/USD	SVM	0.610	0.589	0.630	0.714	0.705	0.723
	LSTM	0.500	0.604	0.323	0.633	0.646	0.618

Table 8: Market trend prediction using main technical indicators – the baseline model.

Type	Baseline			Financial Times			Reddit Headlines		
	Acc	F_1^{up}	F_1^{down}	Acc	F_1^{up}	F_1^{down}	Acc	F_1^{up}	F_1^{down}
DJIA	0.700	0.754	0.615	0.706	0.752	0.639	0.618	0.716	0.417
AAPL	0.685	0.723	0.634	0.652	0.723	0.531	0.624	0.700	0.496
JPM	0.673	0.733	0.578	0.679	0.739	0.583	0.615	0.713	0.415
EUR/USD	0.641	0.715	0.518	0.653	0.691	0.605	0.638	0.684	0.578
GBP/USD	0.714	0.705	0.723	0.711	0.708	0.715	0.615	0.625	0.605

Table 9: Market trend prediction using FT news articles and RWNC headlines (2011-2015).

Type	baseline			all+attitude+emotion			all+emotion			filtering+emotion		
	Acc	F_1^{up}	F_1^{down}	Acc	F_1^{up}	F_1^{down}	Acc	F_1^{up}	F_1^{down}	Acc	F_1^{up}	F_1^{down}
DJIA	0.810	0.854	0.727	0.810	0.846	0.750	0.778	0.829	0.682	-	-	-
AAPL	0.889	0.918	0.829	0.810	0.860	0.700	0.794	0.847	0.683	0.794	0.831	0.735
JPM	0.746	0.800	0.652	0.730	0.779	0.653	0.746	0.789	0.680	0.778	0.829	0.682
GBP/USD	0.708	0.387	0.808	0.662	0.389	0.766	0.631	0.294	0.750	-	-	-
EUR/USD	0.685	0.627	0.727	0.685	0.627	0.727	0.692	0.626	0.739	-	-	-

Table 10: Market trend prediction using financial tweets from Twitter (01/04/2014 – 01/04/2015).

both as additional features. Only for the JPM stock we could see noticeable performance improvements in the “filtering+emotion” setting. Once again these results are consistent with the causality analysis in Section 4 and the market trend prediction experiments using financial news in Section 5.2 – the JPM stock demonstrated that integrating sentiment emotions has the potential to enhance the baseline model. Our results have also confirmed that sentiment attitudes on their own are probably not very useful for market trend prediction, but at least for some particular stocks sentiment emotions could be exploited to improve machine learning models like LSTM to get better market trend prediction.

Our findings are mostly in line with other researchers’ results [2]. However, there are still many questions remaining unanswered in this area.

6 CONCLUSIONS

In this paper, we have empirically re-examined the feasibility of applying sentiment analysis to make market trend predictions.

Our experiments investigated the causal relationship between sentiment attitude/emotion signals and stock price movements

using various sentiment signal sources and different time periods. The experimental results indicate that the interaction between sentiment and price is complex and dynamic: while some stocks in some time periods exhibited strong cross-correlation, it was absent in other cases. We have discovered that in general sentiment attitudes do not seem to have any Granger-causality with respect to stock prices but sentiment emotions do seem to Granger-cause price changes for some particular stocks (e.g., JPM). Furthermore, we have attempted to incorporate sentiment signals into machine learning models for market trend prediction. Specifically, we have compared two popular machine learning approaches and finally selected SVM (with the RBF kernel) as the baseline which was trained using fifteen technical indicators and on average achieved 70% accuracy for five-day market trend prediction. The baseline model was then expanded using sentiment attitudes and sentiment emotions extracted from financial news or financial tweets as additional features. In some scenarios, the proposed model outperformed the baseline model and demonstrated that sentiment emotions could be employed to help predict stock price movements though sentiment attitudes could not. The sentiment emotions extracted from Financial Times

news articles yielded better performances than those extracted from Reddit news headlines.

An important research question for future work is how to identify the stocks whose price changes are indeed predictable using sentiment emotions. Although the Granger causality test on the historical data could find the stocks that were predictable in the past, there is no guarantee that they will continue to be predictable in the future. It is very possible that a more sophisticated classifier for this purpose could be developed.

ACKNOWLEDGMENTS

The Titan X Pascal GPU used for this research was kindly donated by the NVIDIA Corporation. We thank the reviewers for their constructive and helpful comments. We also gratefully acknowledge the support of Geek.AI for this work.

REFERENCES

- [1] Marta Arias, Argimiro Arratia, and Ramon Xuriguera. 2014. Forecasting with Twitter Data. *ACM Trans. Intell. Syst. Technol.* 5, 1, Article 8 (Jan. 2014), 24 pages. <https://doi.org/10.1145/2542182.2542190>
- [2] Johan Bollen, Huina Mao, and Xiao-Jun Zeng. 2010. Twitter mood predicts the stock market. *CoRR abs/1010.3003* (2010).
- [3] Li-Juan Cao and Francis Eng Hock Tay. 2003. Support vector machine with adaptive parameters in financial time series forecasting. *IEEE Transactions on neural networks* 14, 6 (2003), 1506–1518.
- [4] Angel X. Chang and Christopher Manning. 2012. SUTime: A library for recognizing and normalizing time expressions. In *Proceedings of the Eight International Conference on Language Resources and Evaluation (LREC'12)* (23–25), Nicoletta Calzolari (Conference Chair), Khalid Choukri, Thierry Declerck, Mehmet Ugur Dogan, Bente Maegaard, Joseph Mariani, Jan Odijk, and Stelios Piperidis (Eds.). European Language Resources Association (ELRA), Istanbul, Turkey.
- [5] Kai Chen, Yi Zhou, and Fangyan Dai. 2015. A LSTM-based method for stock returns prediction: A case study of China stock market. In *Big Data (Big Data), 2015 IEEE International Conference on*. IEEE, 2823–2824.
- [6] Shom Prasad Das and Sudarsan Padiy. 2012. Support vector machines for prediction of futures prices in Indian stock market. *International Journal of Computer Applications* 41, 3 (2012).
- [7] Shangkun Deng, Takashi Mitsubuchi, Kei Shioda, Tatsuro Shimada, and Akito Sakurai. 2011. Combining technical analysis with sentiment analysis for stock price prediction. In *Dependable, Autonomic and Secure Computing (DASC), 2011 IEEE Ninth International Conference on*. IEEE, 800–807.
- [8] David A Dickey and Wayne A Fuller. 1979. Distribution of the estimators for autoregressive time series with a unit root. *Journal of the American statistical association* 74, 366a (1979), 427–431.
- [9] Xiao Ding, Yue Zhang, Ting Liu, and Junwen Duan. 2015. Deep Learning for Event-driven Stock Prediction. In *Proceedings of the 24th International Conference on Artificial Intelligence (IJCAI'15)*. AAAI Press, 2327–2333. <http://dl.acm.org/citation.cfm?id=2832415.2832572>
- [10] Qiyuan Gao. 2016. *Stock market forecasting using recurrent neural network*. Ph.D. Dissertation. University of Missouri–Columbia.
- [11] C. W. J. Granger. 1969. Investigating Causal Relations by Econometric Models and Cross-spectral Methods. *Econometrica* 37, 3 (Aug. 1969), 424–438. <https://doi.org/10.2307/1912791>
- [12] Zonglu He and Koichi Maekawa. 2001. On spurious Granger causality. *Economics Letters* 73, 3 (2001), 307–313.
- [13] Wei Huang, Yoshiteru Nakamori, and Shou-Yang Wang. 2005. Forecasting stock market movement direction with support vector machine. *Computers & Operations Research* 32, 10 (2005), 2513–2522.
- [14] Mirko Lai, Delia Irazú Hernández Farias, Viviana Patti, and Paolo Rosso. 2016. Friends and Enemies of Clinton and Trump: Using Context for Detecting Stance in Political Tweets. In *Mexican International Conference on Artificial Intelligence*. Springer, 155–168.
- [15] Mirko Lai, Viviana Patti, Giancarlo Ruffo, and Paolo Rosso. 2018. Stance Evolution and Twitter Interactions in an Italian Political Debate. In *International Conference on Applications of Natural Language to Information Systems*. Springer, 15–27.
- [16] Xiaowei Lin, Zehong Yang, and Yixu Song. 2009. Short-term stock price prediction based on echo state networks. *Expert systems with applications* 36, 3 (2009), 7313–7317.
- [17] Greta M Ljung and George EP Box. 1978. On a measure of lack of fit in time series models. *Biometrika* 65, 2 (1978), 297–303.
- [18] Saif Mohammad, Svetlana Kiritchenko, Parinaz Sobhani, Xiaodan Zhu, and Colin Cherry. 2016. SemEval-2016 Task 6: Detecting Stance in Tweets. In *Proceedings of the 10th International Workshop on Semantic Evaluation (SemEval-2016)*. 31–41.
- [19] Saif M. Mohammad and Peter D. Turney. 2013. Crowdsourcing a Word-Emotion Association Lexicon. 29, 3 (2013), 436–465.
- [20] Andrius Mudinas, Dell Zhang, and Mark Levene. 2012. Combining Lexicon and Learning Based Approaches for Concept-level Sentiment Analysis. In *Proceedings of the 1st International Workshop on Issues of Sentiment Discovery and Opinion Mining (WISDOM@KDD)*. Beijing, China, 5:1–5:8.
- [21] Andrius Mudinas, Dell Zhang, and Mark Levene. 2018. Bootstrap Domain-Specific Sentiment Classifiers from Unlabeled Corpora. *Transactions of the Association for Computational Linguistics (TACL)* 6 (2018), 269–285.
- [22] Ping-Feng Pai and Chih-Sheng Lin. 2005. A hybrid ARIMA and support vector machines model in stock price forecasting. *Omega* 33, 6 (2005), 497–505.
- [23] Rosalind Wright Picard. 1995. *Affective Computing*. techreport 321. MIT Media Lab.
- [24] R. Plutchik. 1980. A General Psychoevolutionary Theory of Emotion. In *Emotion: Theory, Research, and Experience*, R. Plutchik and H. Kellerman (Eds.). Vol. 1. Academic Press, 189–217.
- [25] James Pustejovsky, Patrick Hanks, Roser Sauri, Andrew See, Robert Gaizauskas, Andrea Setzer, Dragomir Radev, Beth Sundheim, David Day, Lisa Ferro, et al. 2003. The Timebank Corpus. In *Corpus Linguistics*, Vol. 2003. Lancaster, UK, 40.
- [26] Bo Qian and Khaled Rasheed. 2007. Stock Market Prediction with Multiple Classifiers. *Applied Intelligence* 26, 1 (2007), 25–33.
- [27] Klaus R Scherer. 2000. Psychological Models of Emotion. In *The Neuropsychology of Emotion*, Joan C. Borod (Ed.). Oxford University Press (OUP), 137–162.
- [28] Parinaz Sobhani, Saif Mohammad, and Svetlana Kiritchenko. 2016. Detecting Stance in Tweets and Analyzing Its Interaction with Sentiment. In *Proceedings of the 5th Joint Conference on Lexical and Computational Semantics*. 159–169.
- [29] Nitish Srivastava, Geoffrey E. Hinton, Alex Krizhevsky, Ilya Sutskever, and Ruslan Salakhutdinov. 2014. Dropout: A Simple Way to Prevent Neural Networks from Overfitting. *Journal of Machine Learning Research (JMLR)* 15, 1 (2014), 1929–1958.
- [30] Mariona Taulé, M Antonia Martí, Francisco M Rangel, Paolo Rosso, Cristina Bosco, Viviana Patti, et al. 2017. Overview of the Task on Stance and Gender Detection in Tweets on Catalan Independence at IberEval 2017. In *2nd Workshop on Evaluation of Human Language Technologies for Iberian Languages, IberEval 2017*, Vol. 1881. CEUR-WS, 157–177.
- [31] Mark Thomason. 1999. The practitioner methods and tool. *Journal of Computational Intelligence in Finance* 7, 3 (1999), 36–45.
- [32] Stefan Wager, Sida I. Wang, and Percy Liang. 2013. Dropout Training as Adaptive Regularization. In *Advances in Neural Information Processing Systems 26: Annual Conference on Neural Information Processing Systems (NIPS)*. Lake Tahoe, NV, USA, 351–359.
- [33] Gang Wang, Tianyi Wang, Bolun Wang, Divya Sambasivan, Zengbin Zhang, Haitao Zheng, and Ben Y Zhao. 2015. Crowds on wall street: Extracting value from collaborative investing platforms. In *Proceedings of the 18th ACM Conference on Computer Supported Cooperative Work & Social Computing*. ACM, 17–30.
- [34] Jung-Hua Wang and Jia-Yann Leu. 1996. Stock market trend prediction using ARIMA-based neural networks. In *Neural Networks, 1996., IEEE International Conference on*, Vol. 4. IEEE, 2160–2165.
- [35] Janyce Wiebe, Rebecca F. Bruce, and Thomas P. O'Hara. 1999. Development and Use of a Gold-Standard Data Set for Subjectivity Classifications. In *Proceedings of the 27th Annual Meeting of the Association for Computational Linguistics (ACL)*. University of Maryland, College Park, MD, USA.
- [36] J Welles Wilder. 1978. *New concepts in technical trading systems*. Trend Research.
- [37] Frank Z. Xing, Erik Cambria, and Roy E. Welsch. 2018. Intelligent Bayesian Asset Allocation via Market Sentiment Views. *IEEE Computational Intelligence Magazine* (2018).
- [38] Frank Z. Xing, Erik Cambria, and Roy E. Welsch. 2018. Natural Language based Financial Forecasting: A Survey. *Artificial Intelligence Review* 50, 1 (2018), 49–73.