How much does Stock Prediction improve with Sentiment Analysis?

Frederico G. Monteiro, Diogo R. Ferreira

Instituto Superior Técnico (IST), Universidade de Lisboa, 1049-001 Lisboa, Portugal
{frederico.monteiro,diogo.ferreira}@tecnico.ulisboa.pt

Abstract. Financial markets, such as the stock exchange, are known to be extremely volatile and sensitive to news published in the media. Using sentiment analysis, as opposed to using time series alone, should provide a better indication for the prospects of a given financial asset. In this work, the main goal is to quantify the benefit that can be obtained by adding sentiment analysis to predict the up or down movement of stock returns. The approach makes use of several different deep learning models, from vanilla models that rely on market indicators only, to recurrent networks that incorporate news sentiment as well. Surprisingly, the results suggest that the added benefit of sentiment analysis is diminute, and a more significant improvement can be obtained by using sophisticated models with advanced learning mechanisms such as attention.

Keywords: Stock Prediction · Sentiment Analysis · Deep Learning · Attention

1 Introduction

Financial markets are part of the core structure of modern society, since the stability and prosperity of these markets have a great impact on economic development. Stock prediction, i.e. the prediction of the future value of a company stock, is an extremely difficult task because the extent to which the future price of a stock can be predicted from its past history has always been the source of much debate [1].

In essence, there are three major theories about stock behavior: the Random Walk Theory [1] endorsing that markets are completely arbitrary and stock prices are not predictable; the Elliott Wave Theory (EWT) [2] defending that trends follow repeating patterns and those patterns can be used to predict future movements; and the Efficient-Market Hypothesis (EMH) [3] advocating that stock prices are by themselves the result of all available information and it is impossible to derive sustained gains.

Financial data plays a key role in all of the above theories, and its analysis is divided in two categories. Technical Analysis (TA), supporting EWT, is frequently employed in short-term approaches, as it tries to find the correct timing to enter and exit the market; however, it is somewhat subjective in nature, since it depends of the interpretation of technical charts and data. The other category is Fundamental Analysis (FA), supporting EMH, an approach that focuses more on the long-term by analyzing the intrinsic value of information, but is also hard to formalize into rules, with the challenge being how to handle unstructured data in a systematic manner.

It is interesting to observe how the two approaches above can be mapped to two fundamentally different data sources. On one hand, there are market indicators that reflect the day-to-day evolution of stock prices and can be used for TA. On the other hand, there are media news and sentiments that ultimately contribute to shape the evolution of those stocks, and are important for FA. While TA can be addressed, for example, with Time Series Prediction (TSP), FA may require the use of additional data sources and techniques, e.g. Sentiment Analysis (SA), to fully understand how the current stock price synthesizes all the information available about a company.

But how much can one gain by incorporating SA over a pure technical analysis such as TSP? In principle, SA should be able to provide some explanation for the seemingly random evolution of a time series, so it should be possible to achieve an increase in the accuracy of any stock prediction model. The question that we address in this paper is how much of an improvement can one actually obtain by incorporating SA. Since recently there has been an increasing focus on addressing both TSP and SA with Deep Learning (DL) models, we also use DL to facilitate the development of prediction models that combine both components.

In particular, and contrary to our initial intuition, we find that the choice of model and learning mechnisms, namely recurrence and attention, play a far greater role in improving the prediction accuracy than the inclusion of SA, which provides only a minor improvement. Although this might be somewhat unexpected, it does provide some confidence that the continued development and increasing sophistication of DL models will be able to extract useful predictions from market and news data.

In the following, Section 2 provides an overview of related work, Section 3 presents and discusses our implemented models and results, and Section 4 concludes the paper.

2 Related Work

Numerous attempts have been made to predict the movement of stock prices. Since stocks are in a continuous shaping process due to new information [4], their associated data is usually complex, uncertain, incomplete, and vague [5], with neither stationary (mean, variance, and frequency change over time) nor linear properties [6]. Stock prediction is acknowledged to be one of the most challenging time series tasks [7], with the hardest part being to select relevant features and learning mechanisms.

In this section, we have a look at how other authors have approached stock prediction from two different perspectives, namely from the perspective of TSP and from the perspective of SA, which are the most relevant components for our purpose. We also briefly discuss how Deep Learning can be used to address both components.

2.1 Time Series Prediction

Stock prediction has been often approached with statistical techniques, using models such as AR, MA, ARMA and ARIMA [8]. Typically, these models make predictions about a single stock and generally assume that it has linear properties. The main point here is that these models lack the capability to perceive non-linear behavior and dynamics between stocks [9].

Machine Learning (ML) has brought substantial advances in forecasting multiple stocks and understanding the hidden relationships between them. The most commonly used methods used are Support Vector Machines (SVM) [10], Random Forests [11], Bayesian Networks [12] and, more recently, Deep Learning (DL) [13].

The frequent use of DL and Neural Networks (NN) for stock prediction is justified by their ability to generalize the prediction to other assets and the relative ease when dealing with complex and noisy data. These models outperform statistical techniques [14, 15] and other ML models as well [11, 16].

Another alternative is Reinforcement Learning (RL), where an agent automates trades and reduces costs by systematizing TA as a set of rules to anticipate future price shifts [17]. This has been shown to achieve satisfactory results [18]. However, due to high volatility and noise, there might be significant differences between the patterns observed during training and testing, causing the agent to under-perform.

Table 1 presents a summary of the methods employed, the stock indexes and the data sources used by different authors. This survey leads us to conclude that *Yahoo Finance* and *Thomson Reuters* have been the most popular data sources; other sources used are *Bloomberg*, *Google Finance* and public datasets. On the other hand, the most common indexes and markets are the *American* (*S&P*, *NASDAQ*, *DJIA* and *NYSE*) and the *Asian* (*NSE*, *Nikkei*, *KOSPI* and *CSI*); others include *EuroStoxx50* and *Ibovespa*.

Table 1. Techniques and data sources for stock prediction based on time series.

Technique	Index	Yahoo Finance	Thomson Reuters	Other	Unidentified
Stats	American	[19–22]	[23, 24]		[25]
	Asian		[24, 26]		[13, 27–29]
ML	American	[12, 22, 30–34]	[11, 23, 24]	[16, 35, 36]	[25, 37, 38]
	Asian	[12, 39]	[24, 26]	[40]	[28, 38, 41–43]
	Other	[12]		[44]	
NN	American	[22, 31, 32, 34, 45–47]	[11, 23, 48, 49]	[16, 50, 51]	[25, 37, 38]
	Asian	[46]	[49]	[40]	[13, 27–29, 38, 41, 42, 52]
RL	American	[18, 53, 54]	[54]		
	Asian	[54]	[54]		
	Other	[54]	[54]		

Regarding the time frame for training and evaluation, the shortest period found in these works was one month [30] and the longest was 66 years [37]. It is worth noting that statistical techniques tend to use shorter time frames to avoid noise and dimensionality issues [19–22], while ML techniques tend to use longer time frames to properly train the models [11, 23, 37, 49]. Also, the directional movement of stocks (up or down) is the most common prediction goal, and arguably the most important one [35]. Only two articles tried to predict the actual stock price [36, 47].

2.2 Sentiment Analysis

While SA can be regarded as a subfield of Natural Language Processing (NLP), its fast growth in recent years [55] has been propelled by numerous applications that transform human-generated information (news, tweets, etc.) about any topic (companies, products, politics, etc.) into a sentiment signal (usually positive or negative).

In stock markets, participants take actions (to buy, hold or sell stocks) that are defined and affected by what they read and by what those surrounding them read and share, including the opinions of sources they trust, which are also influenced by market news. From this point of view, the emotional sentiment upon a particular stock or company has become a fundamental part of stock prediction [56,57], and many authors have included this additional component in their prediction models.

Table 2 provides an overview of the related works using SA. The data sources are quite diversified, with the most popular being *News* and *Twitter*, which can be explained by the fact that they provide a simple way to access and collect data. In addition, *News* have the advantage of being created by accredited authors and newspapers [58], while *Twitter* can be used to more easily perceive sentiment, even though tweets are not suited for longer pieces or fully justified opinions on a topic [59].

Table 2. Techniques and data sources for stock prediction based on sentiment analysis.

Model \ Source	**News**	**Social Media**	**Twitter**	**Blog & Forums**	**Financial Reports**	**None**
Stats	[30,42,60]	[26,61]	[19,20,22]			[13,25,27,62–67], [14,15,24,28,29]
KNN	[60,68]	[68]		[68]		[66]
Naïve Bayes	[68]	[68]	[69]	[68]	[58]	[12]
MLP	[42,70]		[70]			[25,37,38,41]
Random Forest	[34]	[71]	[35,69]		[58]	[16,23,24,43,66]
SVM	[31,36,39,72], [32,34,68,73]	[26,68,71]	[22,69,72]	[40,68]	[33]	[10,16,66,67,74,75], [24,28,40,44,76,77]
NN	[31,34,60]	[71]	[22,47]			[13,16,23,49,52,63,78,79], [25,27,64–67,74,75], [11,14,15,28,46,48]
CNN	[32,45,50,70,80–82]		[70,80,81]			[28,29,37,38,41,65]
RNN	[32,42]					[28,29,37,38,49]
LSTM	[42,50,51,72,80,83]	[61]	[72,80]	[40]		[23,27–29,38,40,41,49,84]
RL						[18,53,54,85,86]

Regarding sentiment polarity, in its simplest form it involves just two classes: positive and negative. A simple extension is ternary polarity by adding a neutral class [87]. More fined-grained SA can be achieved by categorizing flavors of feelings (such as *fear*, *joy*, *interest*) and emotions (such as *pleasantness*, *attention*, *sensitivity*, *aptitude*) [88], as well as their level of intensity [56]. Some authors [89] have also used a continuous value in a certain range, e.g. $[-1,+1]$.

An interesting conclusion from our survey is that private information causes small or insignificant changes in stock price [90]; in contrast, public information can cause major changes [91]. Therefore, it can be concluded that prediction accuracy is primarily influenced by public information. However, with private information it is still possible to improve the accuracy a bit further.

Another important conclusion is that markets tend to overreact when presented with negative news [92], and it is easier to predict downward trends than upward trends [93].

In general, pessimism tends to increase the trading volume and lead to predictions of negative returns [94], but it tends to disappear within one week.

Since each work in Table 2 uses a different dataset, the prediction models cannot be compared directly, even if their evaluation results are available. However, it is fair to say that prediction accuracy tends to increase from the top row to the bottom rows. In general, it is not easy to integrate SA features into statistical models, KNN, Naïve Bayes, MLP, Random Forests and SVM due to sparsity issues. Hence, they achieve lower accuracy than DL models. On the other hand, a Long Short-Term Memory (LSTM) incorporating SA registered the highest directional accuracy [40]. For us, this is not surprising since RNNs are especially appropriate for NLP and SA tasks.

2.3 Deep Learning

In recent years, as computer hardware allowed for training larger, deeper and more sophisticated neural networks, Deep Learning has become the state-of-the art approach in many fields, including stock prediction.

A Neural Network (NN) is composed of nodes and activation functions. The simplest architecture is the Feed-Forward Neural Network (FNN), which is based on the Multi-Layer Perceptron (MLP), and its name derives from the fact that information flows one-way, from input to output, across a set of densely-connected layers. A different model is the Convolutional Neural Network (CNN), where each layer performs a computation in the form of a sliding window over its input. CNNs are very popular in image processing where the input is 2D, but they are also very successful at processing 1D input such as text sentences [95] and time series [96].

For sequence processing, Recurrent Neural Networks (RNNs) [97] are very appropriate, since they keep an internal state and have a feedback loop to use that internal state as an additional input at each time step. For stock prediction, this means that it is possible to reuse information from the past to predict the future. However, simple RNNs are unable to perform well when a long-term context is required [98]. The LSTM architecture [99] improves on long-term dependencies by keeping a cell state with the possibility of carrying it across multiple time steps. More recently, slightly different versions have been proposed, such as the Gated Recurrent Unit (GRU) [100]. However, the LSTM can still provide state-of-the-art results [41].

Although significant improvements have been obtained with LSTM models, there are still limitations such as the vanishing gradient problem [101]. The use of attention mechanisms mitigates this constraint [102] by selectively retrieving the most relevant information from hidden states [103]. The underlying rationale is to increase the focus on important parts of the input instead of just trying to remember it afterwards [104]. Results suggest that the use of attention can capture long-term dependencies and outperform stand-alone LTSMs [51].

One way to use attention is to attach this mechanism to an LSTM or, alternatively, to develop a model similar to the Transformer [105] which relies exclusively on attention, without recurrence. However, for sequential data it is useful to keep track of the order between time steps, and the Transformer achieves this by means of positional encoding. In our context, the Transformer architecture can be simplified: for example, for time series data, the embedding layers that are typically used for NLP tasks can be removed;

also, the decoder part of the Transformer can be removed to keep only the encoder as a prediction model. Another possibility is to use Multi-Head Attention, which essentially consists in combining multiple layers of attention in parallel to enhance the extraction of long-dependencies from the input sequence.

3 Stock Prediction with Sentiment Analysis

In this section, we report on a series of experiments that we performed during the Kaggle competition *Two Sigma: Using News to Predict Stock Movements*.[1] One of the key aspects of this competition is that it provided two different data sources: (1) market data, with typical indicators about stock prices, trading volumes and calculated returns; and (2) news data, with pre-calculated sentiment, novelty and volume counts for each news item. This made it possible to test several different models for stock prediction, both with and without sentiment analysis.

3.1 Data Description

The original dataset contained market and news data from the beginning of 2007 to the end of 2016. The market data included information on ∼3000 US-listed companies, containing over 4 million samples with features such as date, asset code, asset name, daily open and close prices, daily trading volume, and open-to-open and close-to-close returns. These returns were calculated both daily and for a 10-day period, and also both in raw and market-residualized form (i.e. by removing the movement of the market as a whole and leaving only the movement inherent to the asset).

On the other hand, the news data contained about 9 million samples (daily news items) regarding the same companies or related ones. For each news item, the sentiment was given as a probability distribution over 3 classes: positive, neutral and negative. Measures for the novelty and volume counts of each news item were also available, where novelty was calculated by comparing the asset-specific content of a news item against a cache of previous news items, and the volume was calculated by counting how many news items mentioned the asset within a certain time frame. However, for this work we used only the sentiment features of each news item.

As for the target variable, the competition used the open-to-open market-residualized return over a 10-day period into the future. In our experiments, and following common practice in the literature, we used only the directional movement (up or down) of this target variable. The problem then becomes a binary classification, which makes it especially easy to evaluate accuracy as the percentage of correct predictions, while also enabling the use of a standard loss function for training, i.e. binary cross-entropy.

3.2 Data Preprocessing

When joining the market data with the news data, it may be the case that there are several news items for a given asset on a given date. In this case, we group those news items together and compute their average sentiment.

[1] https://www.kaggle.com/c/two-sigma-financial-news

On the other hand, it may be the case that there are actually no news items for a given asset on a particular date. In this case, we have tried three different approaches to fill in the sentiment for the asset, namely:

– *propagation*, where we use the sentiment of the previous date as the current sentiment, hence propagating the sentiment across multiple days;
– *balancing*, where we use a uniform probability distribution over the three classes (positive, neutral and negative, with $\frac{1}{3}$ probability for each);
– *neutralizing*, where we assign zero probability to the positive and negative classes, and 1.0 to the neutral class.

After joining both data sources according to one of the strategies above, the dataset was split into 2007–2014 for training, 2015 for validation, and 2016 for testing.

3.3 Prediction Models

In our models, we used 2 categorical features and 15 numerical features as input. The 2 categorical features were asset code and asset name, because each asset name (i.e. company) may have several asset codes, but each asset code belongs to a single company. As for the numerical features, 12 of those features come from the market data, and the remaining ones correspond to the 3 sentiment classes from news data.

The way these categorical and numerical features have been combined is illustrated in Fig. 1. In essence, the categorical features go through an embedding layer before being concatenated together with the numerical features. This initial block provides the input for all subsequent models, except for the Feed-Forward Neural Network (FNN) which uses a single time-step rather than 10 time-steps as the other models.

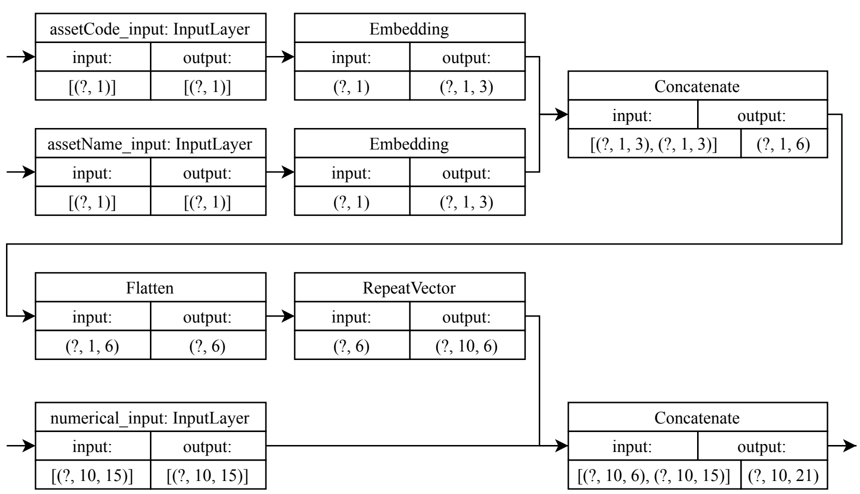

Fig. 1. Input block.

The first experiment, with a simple FNN (Fig. 2), already provides an interesting conclusion: when using this model, the sentiment features do not have any influence in the output predictions. In our view, this is due to the fact that the FNN does not incorporate any information on past behavior. Both with and without the sentiment features, this model scored a test accuracy of 53.3%.

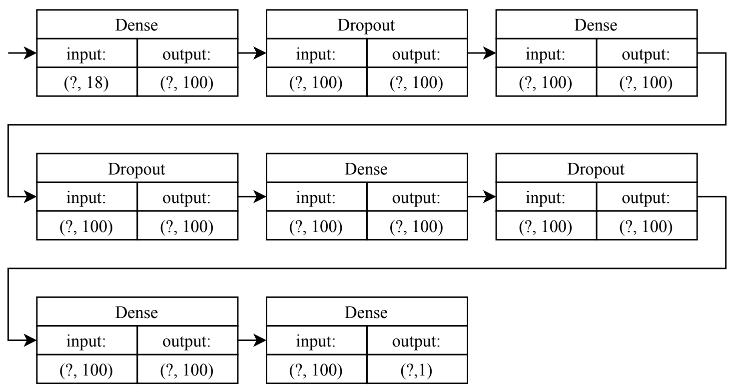

Fig. 2. FNN model.

Using a Convolutional Neural Network (CNN) (Fig. 3) with a 10 time-step window into the past improved the prediction accuracy. However, the test accuracy was higher without the sentiment features (55.1%) than with those features included (54.8%). We attribute this fact to overfitting, and dropout did not help in this case.

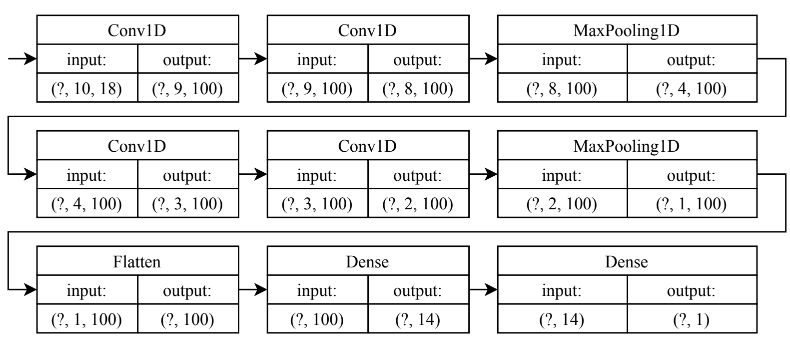

Fig. 3. CNN model.

As soon as we moved on to recurrent networks (Fig. 4), we began to observe more consistent results. A simple Recurrent Neural Network (RNN) provided the same accuracy (55.2%) with and without sentiment features, but a Long Short-Term Memory (LSTM) improved from 55.2% to 55.3% when the sentiment features were included.

We conclude that recurrence mechanisms are essential in order to leverage information from the past. However, a bi-directional LSTM did not improve the results, which is understandable, since the most recent history is probably more relevant for prediction than the distant past, so only the forward direction is useful.

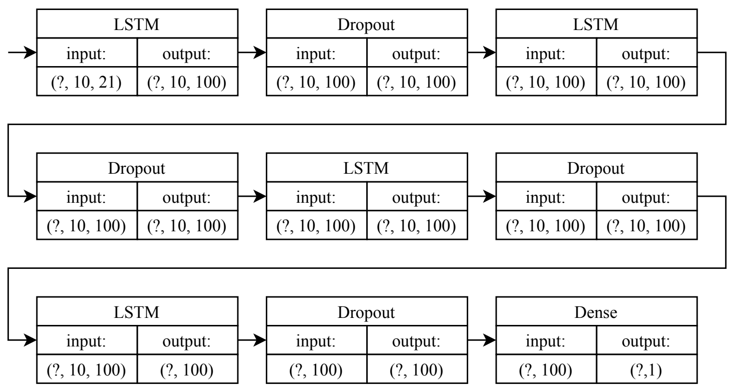

Fig. 4. LSTM model.

The accuracy improves even further when attention mechanisms are employed. An LSTM model with additive attention achieves 59.8% accuracy without sentiment features, and reaches 60.0% when those features are included. Moreover, a bidirectional LSTM with attention (Fig. 5) improves that result even further to 60.4%. Here, bidirectionality is definitely beneficial because it helps the attention mechanism learn which time-steps are the most relevant for prediction.

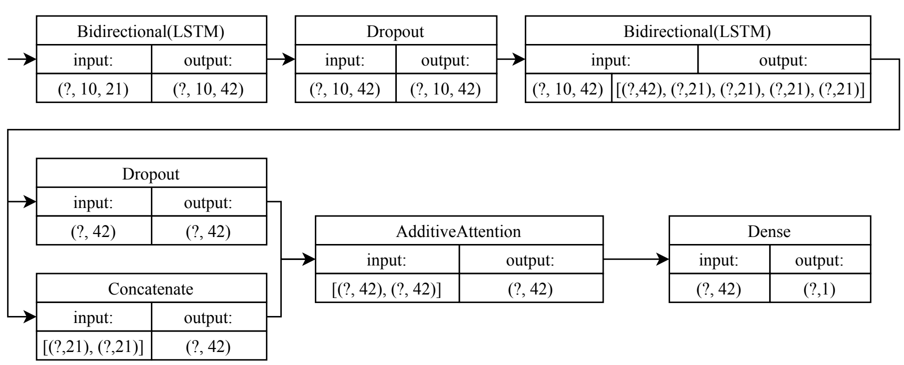

Fig. 5. Bidirectional LSTM with attention, the best performing model.

We have also experimented with a Transformer-like architecture, which relies on attention but not recurrence. Here the results were inferior, with 55.6% accuracy without sentiment features, and an increase to 55.7% when those features were included.

Table 3 provides an overview of the results, where it becomes apparent that the inclusion of sentiment features can account for an improvement of at most 0.6% in test accuracy, while the use of recurrence and attention mechanisms can provide an improvement of 5.6% over a baseline CNN that uses the same input data. Other metrics, such as area under the ROC curve (AUC) and Matthews correlation coefficient (MCC) have commensurate improvements with accuracy.

Table 3. Comparison of accuracy.

Model	Market only	Market+News propagated	Market+News balanced	Market+News neutralized
FNN	**0.533**	0.532	0.533	0.533
CNN	**0.551**	0.543	0.543	0.548
RNN	0.552	0.547	**0.552**	0.550
LSTM	0.552	0.551	**0.553**	0.551
Bi-LSTM	0.543	0.548	0.547	**0.549**
LSTM + Att	0.598	0.599	**0.600**	0.600
Bi-LSTM + Att	0.600	0.603	0.602	**0.604**
Transformer	0.556	0.552	**0.557**	0.550

All models have been implemented with TensorFlow and have been tested with different hyper-parameters (number of layers, number of filters, kernel sizes, dropout rates, etc.). Here we reported the results achieved with the best-performing version.

4 Conclusion

From the literature review that we presented in Section 2, we expected that sentiment analysis would improve stock prediction, but our experiments and results in Section 3 suggest that the improvement is rather modest compared to our initial expectation.

On the other hand, our findings point to a possible avenue for the continued improvement of prediction models and their accuracy. By making use of learning mechanisms such as recurrence and attention, and others that may appear along the way, it is possible to keep improving the results. It is also apparent that such mechanisms will work better in combination rather than in isolation; specifically, recurrence with attention works better than either of those mechanisms alone.

In addition, we found that propagating sentiment by keeping past sentiment when there are no news is not the best approach. In this respect, the market seems to forget past sentiment in a matter of days, to the point that it is no longer possible to clearly determine the current sentiment as being positive, negative, or neutral.

In this work, our prediction target was binary, in the form of a directional movement. In future work, we plan to apply similar models for regression tasks such as predicting price and volatility. In addition, the kind of evaluation that we provide here is by no means the end of the story; to derive actual benefits in the real-world, other components, such as a trading strategy to enter and exit the market, are necessary.

References

1. Fama, E.F.: The Behavior of Stock-Market Prices. The Journal of Business **38**(1), 34–105 (Jan 1965)
2. Elliott, R.N.: The Wave Principle (1938)
3. Fama, E.F.: Efficient Capital Markets: A Review of Theory and Empirical Work. The Journal of Finance **25**(2), 383–417 (May 1970)
4. LeBaron, B., Arthur, W.B., Palmer, R.: Time series properties of an artificial stock market. JEDC **23**(9-10), 1487–1516 (Sep 1999)
5. Emerson, S., Kennedy, R., O'Shea, L., O'Brien, J.: Trends and Applications of Machine Learning in Quantitative Finance. In: ICEFR 2019 (Jun 2019)
6. Vanstone, B., Finnie, G.: An empirical methodology for developing stock market trading systems using artificial neural networks. Expert Systems with Applications **36**(3), 6668–6680 (Apr 2009)
7. Hussain, A.J., Knowles, A., Lisboa, P.J.G., El-Deredy, W.: Financial time series prediction using polynomial pipelined neural networks. Expert Systems with Applications **35**(3), 1186–1199 (Oct 2008)
8. Box, G.E.P., Jenkins, G.: Time Series Analysis, Forecasting and Control. Wiley, 5 edn. (1990)
9. Fang, J., Jacobsen, B., Qin, Y.: Predictability of the simple technical trading rules: An out-of-sample test. Review of Financial Economics **23**(1), 30–45 (Jan 2014)
10. Cao, L.J., Tay, F.E.H.: Support vector machine with adaptive parameters in financial time series forecasting. IEEE Transactions on Neural Networks **14**(6), 1506–1518 (Nov 2003)
11. Krauss, C., Do, X.A., Huck, N.: Deep neural networks, gradient-boosted trees, random forests: Statistical arbitrage on the S&P 500. EJOR **259**(2), 689–702 (Jun 2017)
12. Malagrino, L.S., Roman, N.T., Monteiro, A.M.: Forecasting stock market index daily direction: A Bayesian Network approach. Expert Systems with Applications **105**, 11–22 (2018)
13. Chong, E., Han, C., Park, F.C.: Deep learning networks for stock market analysis and prediction: Methodology, data representations, and case studies. Expert Systems with Applications **83**, 187–205 (Oct 2017)
14. Kohzadi, N., Boyd, M.S., Kermanshahi, B., Kaastra, I.: A comparison of artificial neural network and time series models for forecasting commodity prices. Neurocomputing **10**(2), 169–181 (Mar 1996)
15. Zhang, G.P.: Time series forecasting using a hybrid ARIMA and neural network model. Neurocomputing **50**, 159–175 (Jan 2003)
16. Chatzis, S.P., Siakoulis, V., Petropoulos, A., Stavroulakis, E., Vlachogiannakis, N.: Forecasting stock market crisis events using deep and statistical machine learning techniques. Expert Systems with Applications **112**, 353–371 (Dec 2018)
17. Nazário, R.T.F., e Silva, J.L., Sobreiro, V.A., Kimura, H.: A literature review of technical analysis on stock markets. QREF **66**, 115–126 (Nov 2017)
18. Deng, Y., Bao, F., Kong, Y., Ren, Z., Dai, Q.: Deep Direct Reinforcement Learning for Financial Signal Representation and Trading. IEEE Transactions on Neural Networks and Learning Systems **28**(3), 653–664 (Mar 2017)
19. Si, J., Mukherjee, A., Liu, B., Li, Q., Li, H., Deng, X.: Exploiting Topic based Twitter Sentiment for Stock Prediction. In: ACL 2013. vol. 2, pp. 24–29 (Aug 2013)
20. Si, J., Mukherjee, A., Liu, B., Pan, S.J., Li, Q., Li, H.: Exploiting Social Relations and Sentiment for Stock Prediction. In: EMNLP. pp. 1139–1145 (Oct 2014)
21. Dickinson, B., Hu, W.: Sentiment Analysis of Investor Opinions on Twitter. Social Networking **4**(3), 62–71 (Jul 2015)

22. Mittal, A., Goel, A.: Stock Prediction Using Twitter Sentiment Analysis. Standford University, CS229 (2012)
23. Fischer, T., Krauss, C.: Deep learning with long short-term memory networks for financial market predictions. EJOR **270**(2), 654–669 (Oct 2018)
24. Lee, T.K., Cho, J.H., Kwon, D.S., Sohn, S.Y.: Global stock market investment strategies based on financial network indicators using machine learning techniques. Expert Systems with Applications **117**, 228–242 (Mar 2019)
25. Guresen, E., Kayakutlu, G., Daim, T.U.: Using artificial neural network models in stock market index prediction. Expert Systems with Applications **38**(8), 10389–10397 (2011)
26. Wang, H., Lu, S., Zhao, J.: Aggregating multiple types of complex data in stock market prediction: A model-independent framework. Knowledge-Based Systems **164**, 193–204 (2019)
27. Kim, H.Y., Won, C.H.: Forecasting the volatility of stock price index: A hybrid model integrating LSTM with multiple GARCH-type models. Expert Systems with Applications **103**, 25–37 (Aug 2018)
28. Long, W., Lu, Z., Cui, L.: Deep learning-based feature engineering for stock price movement prediction. Knowledge-Based Systems **164**, 163–173 (Jan 2019)
29. Selvin, S., Vinayakumar, R., Gopalakrishnan, E.A., Menon, V.K., Soman, K.P.: Stock price prediction using LSTM, RNN and CNN-sliding window model. In: ICACCI (Sep 2017)
30. Schumaker, R.P., Chen, H.: Textual Analysis of Stock Market Prediction Using Financial News Articles. In: AMCIS 2006. vol. 3, pp. 1422–1430 (Dec 2006)
31. Ding, X., Zhang, Y., Liu, T., Duan, J.: Using Structured Events to Predict Stock Price Movement: An Empirical Investigation. EMNLP 2014 pp. 1415–1425 (Oct 2014)
32. Vargas, M.R., De Lima, B.S.L.P., Evsukoff, A.G.: Deep learning for stock market prediction from financial news articles. In: CIVEMSA 2017. pp. 60–65 (Jun 2017)
33. Nguyen, T.H., Shirai, K.: Topic Modeling based Sentiment Analysis on Social Media for Stock Market Prediction. In: ACL-IJCNLP 2015. vol. 1, pp. 1354–1364 (Jul 2015)
34. Weng, B., Lu, L., Wang, X., Megahed, F.M., Martinez, W.: Predicting short-term stock prices using ensemble methods and online data sources. Expert Systems with Applications **112**, 258–273 (Dec 2018)
35. Nasseri, A.A., Tucker, A., de Cesare, S.: Quantifying StockTwits semantic terms' trading behavior in financial markets: An effective application of decision tree algorithms. Expert Systems with Applications **42**(23), 9192–9210 (Dec 2015)
36. Deng, S., Mitsubuchi, T., Shioda, K., Shimada, T., Sakurai, A.: Combining Technical Analysis with Sentiment Analysis for Stock Price Prediction. In: DASC2011. pp. 800–807 (Dec 2011)
37. Di Persio, L., Honchar, O.: Artificial neural networks architectures for stock price prediction: Comparisons and Applications. Int. J. Circuits, Syst. Signal Process. **10**, 403–413 (2016)
38. Hiransha, M., Gopalakrishnan, E.A., Menon, V.K., Soman, K.P.: NSE Stock Market Prediction Using Deep-Learning Models. Procedia Computer Science **132**, 1351–1362 (2018)
39. Li, X., Xie, H., Chen, L., Wang, J., Deng, X.: News impact on stock price return via sentiment analysis. Knowledge-Based Systems **69**, 14–23 (Oct 2014)
40. Li, J., Bu, H., Wu, J.: Sentiment-aware stock market prediction: A deep learning method. In: ICSSSM'17 (June 2017)
41. Kim, S., Kang, M.: Financial series prediction using Attention LSTM. arXiv:1902.10877 (Feb 2019)
42. Akita, R., Yoshihara, A., Matsubara, T., Uehara, K.: Deep learning for stock prediction using numerical and textual information. ICIS'16 **1**, 1–6 (Jun 2016)
43. Zhang, J., Cui, S., Xu, Y., Li, Q., Li, T.: A novel data-driven stock price trend prediction system. Expert Systems with Applications **97**, 60–69 (May 2017)

44. Paiva, F.D., Cardoso, R.T.N., Hanaoka, G.P., Duarte, W.M.: Decision-making for financial trading: A fusion approach of machine learning and portfolio selection. Expert Systems with Applications **115**, 635–655 (Jan 2019)
45. Ding, X., Zhang, Y., Liu, T., Duan, J.: Deep Learning for Event-Driven Stock Prediction. IJCAI 2015 pp. 2327–2333 (Jul 2015)
46. Zhou, F., Zhou, H.m., Yang, Z., Yang, L.: EMD2FNN: A strategy combining empirical mode decomposition and factorization machine based neural network for stock market trend prediction. Expert Systems with Applications **115**, 136–151 (Jan 2019)
47. Bollen, J., Mao, H., Zeng, X.: Twitter mood predicts the stock market. Journal of Computational Science **2**(1), 1–8 (Mar 2011)
48. Moews, B., Herrmann, J.M., Ibikunle, G.: Lagged correlation-based deep learning for directional trend change prediction in financial time series. Expert Systems with Applications **120**, 197–206 (Apr 2019)
49. Baek, Y., Kim, H.Y.: ModAugNet: A new forecasting framework for stock market index value with an overfitting prevention LSTM module and a prediction LSTM module. Expert Systems with Applications **113**, 457–480 (Dec 2018)
50. Prosky, J., Song, X., Tan, A., Zhao, M.: Sentiment Predictability for Stocks. arXiv:1712.05785 (Dec 2017)
51. Hollis, T., Viscardi, A., Yi, S.E.: A Comparison of LSTMs and Attention Mechanisms for Forecasting Financial Time Series. arXiv:1812.07699 (Dec 2018)
52. Day, M.Y., Lee, C.C.: Deep learning for financial sentiment analysis on finance news providers. In: ASONAM. pp. 1127–1134 (Aug 2016)
53. Almahdi, S., Yang, S.Y.: An adaptive portfolio trading system: A risk-return portfolio optimization using recurrent reinforcement learning with expected maximum drawdown. Expert Systems with Applications **87**, 267–279 (Nov 2017)
54. Jeong, G., Kim, H.Y.: Improving financial trading decisions using deep Q-learning: Predicting the number of shares, action strategies, and transfer learning. Expert Systems with Applications **117**, 125–138 (Mar 2019)
55. Mäntylä, M., Graziotin, D., Kuutila, M.: The Evolution of Sentiment Analysis - A Review of Research Topics, Venues, and Top Cited Papers. Computer Science Review **27**, 16–32 (Feb 2018)
56. Devitt, A., Ahmad, K.: Sentiment Polarity Identification in Financial News: A Cohesion-based Approach. In: ACL 2007. pp. 984–991 (Jun 2007)
57. Feuerriegel, S., Prendinger, H.: News-based trading strategies. Decision Support Systems **90**, 65–74 (Oct 2016)
58. Sehgal, V., Song, C.: SOPS: Stock Prediction Using Web Sentiment. In: ICDM 2007. pp. 21–26 (Oct 2007)
59. Cambria, E., Schuller, B., Xia, Y., Havasi, C.: New Avenues in Opinion Mining and Sentiment Analysis. IEEE Intelligent Systems **28**(2), 15–21 (Mar 2013)
60. Wuthrich, B., Cho, V., Leung, S., Permunetilleke, D., Sankaran, K., Zhang, J.: Daily stock market forecast from textual web data. In: SMC'98. vol. 3, pp. 2720–2725 (Oct 1998)
61. Xiong, R., Nichols, E.P., Shen, Y.: Deep Learning Stock Volatility with Google Domestic Trends. arXiv:1512.04916 (Dec 2015)
62. Agarwal, S., Kumar, S., Goel, U.: Stock market response to information diffusion through internet sources: A literature review. IJIM **45**, 118–131 (Apr 2019)
63. Arévalo, R., García, J., Guijarro, F., Peris, A.: A dynamic trading rule based on filtered flag pattern recognition for stock market price forecasting. Expert Systems with Applications **81**, 177–192 (Sep 2017)
64. Hollis, T.: Deep Learning Algorithms Applied to Blockchain-Based Financial Time Series. Tech. rep., University of Manchester (2018)

65. Hoseinzade, E., Haratizadeh, S.: CNNpred: CNN-based stock market prediction using a diverse set of variables. Expert Systems with Applications **129**, 273–285 (Sep 2019)
66. Huang, C.J., Yang, D.X., Chuang, Y.T.: Application of wrapper approach and composite classifier to the stock trend prediction. Expert Systems with Applications **34**(4), 2870–2878 (May 2008)
67. Huang, W., Nakamori, Y., Wang, S.Y.: Forecasting stock market movement direction with support vector machine. Computers & Operations Research **32**(10), 2513–2522 (Oct 2005)
68. Nassirtoussi, A.K., Aghabozorgi, S., Wah, T.Y., Ngo, D.C.L.: Text mining of news-headlines for FOREX market prediction: A Multi-layer Dimension Reduction Algorithm with semantics and sentiment. Expert Systems with Applications **42**(1), 306–324 (Jan 2015)
69. Li, B., Chan, K.C.C., Ou, C., Ruifeng, S.: Discovering public sentiment in social media for predicting stock movement of publicly listed companies. Information Systems **69**, 81–92 (Sep 2017)
70. Kar, S., Maharjan, S., Solorio, T.: RiTUAL-UH at SemEval-2017 Task 5: Sentiment Analysis on Financial Data Using Neural Networks. In: SemEval-2017. pp. 877–882 (Aug 2017)
71. Weng, B., Ahmed, M.A., Megahed, F.M.: Stock market one-day ahead movement prediction using disparate data sources. Expert Systems with Applications **79**, 153–163 (2017)
72. Cabanski, T., Romberg, J., Conrad, S.: HHU at SemEval-2017 Task 5: Fine-Grained Sentiment Analysis on Financial Data using Machine Learning Methods. In: SemEval-2017. pp. 832–836 (Aug 2017)
73. Schumaker, R.P., Chen, H.: A quantitative stock prediction system based on financial news. Information Processing and Management **45**(5), 571–583 (May 2009)
74. Kara, Y., Boyacioglu, M.A., Baykan, Ö.K.: Predicting direction of stock price index movement using artificial neural networks and support vector machines: The sample of the Istanbul Stock Exchange. Expert Systems with Applications **38**(5), 5311–5319 (May 2011)
75. Kim, K.: Financial time series forecasting using support vector machines. Neurocomputing **55**(1-2), 307–319 (Sep 2003)
76. Sen, J., Chaudhuri, T.D.: Decomposition of Time Series Data of Stock Markets and its Implications for Prediction – An Application for the Indian Auto Sector. In: ABRMP ’2016 (Jan 2016)
77. Zhiqiang, G., Huaiqing, W., Quan, L.: Financial time series forecasting using LPP and SVM optimized by PSO. Soft Computing **17**(5), 805–818 (May 2013)
78. Arévalo, A., Nino, J., Hernandez, G., Sandoval, J.: High-Frequency Trading Strategy Based on Deep Neural Networks. In: Springer (ed.) LNCS’16. vol. 9773, pp. 424–436 (08 2016)
79. Dixon, M., Klabjan, D., Bang, J.H.: Classification-based financial markets prediction using deep neural networks. Algorithmic Finance **6**(3-4), 67–77 (Jan 2017)
80. Ghosal, D., Bhatnagar, S., Akhtar, M.S., Ekbal, A., Bhattacharyya, P.: IITP at SemEval-2017 Task 5: An Ensemble of Deep Learning and Feature Based Models for Financial Sentiment Analysis. In: SemEval-2017. pp. 899–903 (Aug 2017)
81. Pivovarova, L., Escoter, L., Klami, A., Yangarber, R.: HCS at SemEval-2017 Task 5: Sentiment Detection in Business News Using Convolutional Neural Networks. In: SemEval-2017. pp. 842–846 (Aug 2017)
82. Mansar, Y., Gatti, L., Ferradans, S., Guerini, M., Staiano, J.: Fortia-FBK at SemEval-2017 Task 5: Bullish or Bearish? Inferring Sentiment towards Brands from Financial News Headlines. arXiv:1704.00939 (Apr 2017)
83. Moore, A., Rayson, P.: Lancaster A at SemEval-2017 Task 5: Evaluation metrics matter: predicting sentiment from financial news headlines. arXiv:1705.00571 (May 2017)
84. Nelson, D., Pereira, A., de Oliveira, R.: Stock market’s price movement prediction with LSTM neural networks. IJCNN’17 (May 2017)

85. Bertoluzzo, F., Corazza, M.: Testing different Reinforcement Learning configurations for financial trading: Introduction and applications. Procedia Economics and Finance **3**, 68–77 (Jan 2012)
86. Pendharkar, P.C., Cusatis, P.: Trading financial indices with reinforcement learning agents. Expert Systems with Applications **103**, 1–13 (Aug 2018)
87. Rojas-Barahona, L.M.: Deep learning for sentiment analysis. Language and Linguistics Compass **10**(12), 701–719 (Dec 2016)
88. Loia, V., Senatore, S.: A fuzzy-oriented sentic analysis to capture the human emotion in Web-based content. Knowledge-Based Systems **58**, 75–85 (Mar 2014)
89. Barbaglia, L., Consoli, S., Manzan, S.: Monitoring the business cycle with fine-grained, aspect-based sentiment extraction from news. In: Mining Data for Financial Applications. LNAI, vol. 11985, pp. 101–106. Springer (2020)
90. Ahern, K.R., Sosyura, D.: Who Writes the News? Corporate Press Releases during Merger Negotiations. The Journal of Finance **69**(1), 241–291 (Feb 2014)
91. Vega, C.: Stock price reaction to public and private information. Journal of Financial Economics **82**(1), 103–133 (Oct 2006)
92. Niederhoffer, V.: The Analysis of World Events and Stock Prices. The Journal of Business **44**(2), 193–219 (Apr 1971)
93. Hu, Y., Liu, K., Zhang, X., Su, L., Ngai, E.W.T., Liu, M.: Application of evolutionary computation for rule discovery in stock algorithmic trading: A literature review. Applied Soft Computing **36**, 534–551 (Nov 2015)
94. Tetlock, P.C.: Giving Content to Investor Sentiment: The Role of Media in the Stock Market. The Journal of finance **62**(3), 1139–1168 (Jun 2007)
95. Kim, Y.: Convolutional Neural Networks for Sentence Classification. arXiv:1408.5882 (Oct 2014)
96. Fawaz, H.I., Forestier, G., Weber, J., Idoumghar, L., Muller, P.: Deep learning for time series classification: a review. Data Mining and Knowledge Discovery **33**(4), 917–963 (Jul 2019)
97. Sutskever, I., Martens, J., Hinton, G.E.: Generating Text with Recurrent Neural Networks. In: ICML'11. pp. 1017–1024 (2011)
98. Bengio, Y., Simard, P., Frasconi, P.: Learning Long-Term Dependencies with Gradient Descent is Difficult. IEEE Transactions on Neural Networks **5**(2), 157–166 (Mar 1994)
99. Hochreiter, S., Schmidhuber, J.: Long Short Term Memory. Neural Computation **9**(8), 1735–1780 (Nov 1997)
100. Cho, K., van Merrienboer, B., Gulcehre, C., Bougares, F., Schwenk, H., Bengio, Y.: Learning Phrase Representations using RNN Encoder-Decoder for Statistical Machine Translation. arXiv:1406.1078 (Jun 2014)
101. Hochreiter, S.: The Vanishing Gradient Problem During Learning Recurrent Neural Nets and Problem Solutions. IJUFKS '98 **6**(2), 107–116 (1998)
102. Lin, Z., Feng, M., dos Santos, C.N., Yu, M., Xiang, B., Zhou, B., Bengio, Y.: A Structured Self-attentive Sentence Embedding. ICLR'17 (Mar 2017)
103. Bahdanau, D., Cho, K., Bengio, Y.: Neural Machine Translation by Jointly Learning to Align and Translate. arXiv:1409.0473 (2014)
104. Sutskever, I., Vinyals, O., Le, Q.V.: Sequence to Sequence Learning with Neural Networks. In: NIPS'14. pp. 3104–3112 (2014)
105. Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A.N., Kaiser, L., Polosukhin, I.: Attention is All You Need. In: NIPS'17. pp. 6000–6010 (Dec 2017)