Predicting stock returns: ARMAX vs. machine learning

Darya Lapitskaya [1], Hakan Eratalay [24], Rajesh Sharma [34]

Abstract

In the modern world, online social and news media significantly impact society, economy, and financial markets. In this chapter, we compared the predictive performance of financial econometrics and machine learning and deep learning methods for the returns of the stocks of the SP100 index. The analysis is enriched by using COVID-19 related news sentiments data collected for a period of 10 months. We analyzed the performance of each model and found the best algorithm for such types of predictions. For the sample we analyzed, our results indicate that the autoregressive moving average model with exogenous variables (ARMAX) has a comparable predictive performance to the machine and deep learning models, only outperformed by the extreme gradient boosted trees (XGBoost) approach. This result holds both in the training and testing datasets.

Keywords: Sentiment Analysis, Machine Learning, ARMAX, Stock Returns Prediction, Deep Learning, COVID-19

1   Introduction

In the modern world, social and news media significantly impact society and the economy Bruhn et al. (2012), Hanusch and Tandoc (2019). They change the companies' business models and affect their performance and reputation as the opinions about a product or service now can be freely shared online. Hence, media also affect the stock markets and stock prices. Many researchers found dependencies between information and media and the company's performance at the stock market Steyn et al. (2020), Khan et al. (2020), Coyne et al. (2017).

The impact of information on the stock market and stock price volatilities has been investigated for many years Malkiel (2003). According to the earlier research and efficient market hypothesis, stock market prices are more affected by new information than by the present and past prices Arafat et al. (2013). Later on, it was proved that the public mood measured from posts on social media was correlated with the market prediction Arafat et al. (2013). Also, it was demonstrated that news and social media sentiments could predict future

---

[1] darya.lapitskaya@ut.ee, School of Economics and Business Administration, University of Tartu, Estonia

[2] hakan.eratalay@ut.ee, School of Economics and Business Administration, University of Tartu, Estonia

[3] rajesh.sharma@ut.ee, Institute of Computer Science, University of Tartu, Estonia

stock returns Leung et al. (2019). According to another research Mohan et al. (2019), there is a strong correlation between the stock price volatilities and the news articles. Hence, it can be concluded that online social media sentiments could be used to predict stock returns.

The prediction of stock returns is usually performed with different econometrical models, such as the autoregressive–moving-average with exogenous inputs (ARMAX). However, nowadays, machine learning (ML) Shah et al. (2018), deep learning Abe and Nakayama (2018), and graph neural networks Sharma and Sharma (2020) are actively used in financial econometrics and forecasts. The usage of ML forecasts brings significant financial benefits to investors, as in some cases, such methods have doubled the predictive performance of leading regression-based methodologies Gu et al. (2018). The interest in comparing ARMAX, ML and deep learning methods' predictive performance rests in the possible disadvantages of these methods. On the one hand, ARMAX models are restricted with the stationarity and invertability conditions on the coefficient estimators, which would trade off the predictive power with the stationary behaviour of the predicted variable. On the other hand, the ML and deep learning methods focus on better predictive performance considering many lags and/or functional forms of the variables but leaving the interpretability of the coefficients in the back plan. So, when the interest is in predicting the future behaviour of returns, it is expected that ML and deep learning methods would outperform the ARMAX method in terms of predictive performance.

This chapter gives an example of the mentioned comparison using highly volatile data due to the effect of COVID-19.In particular, this study is dedicated to analyzing the impact of COVID-19 related news on the Standard and Poor's 100 companies (SP100) stock returns by collecting news article for the period of 10 months. Specifically, we analyse stock returns predictions with sentiments scores by comparing two different prediction methods. The first one belongs to the traditional econometric domain (ARMAX modelling), and the second one is from the domain of machine learning (KNN, XGBoost and a deep learning neural network (LSTM)).

The comparison between different prediction techniques showed that the XGBoost had better predictive performance than the ARMAX model, while KNN and LSTM performed worse. Although the result is interesting that the ARMAX model performed well despite the restrictions on the parameters than KNN and LSTM, we acknowledge that the small sample size might be the culprit. The findings of this research contribute to the literature and allow understanding the impact of COVID-19 related news articles on the stock markets. Moreover, the application of both sentiment analysis and ML prediction techniques help create a more precise returns prediction model.

The remainder of the chapter is organized in the following way: in the next section, the theoretical aspects of stock return predictions are presented as well as the analysis of

previous research done regarding the sentiments impact on the stock market. The third section is dedicated to the research method and its' methodology. In the fourth section, the results are presented together with a discussion followed by the study conclusions and future research perspectives in the last section.

2   Literature review

Stock prices forecasting is one of the critical fields in the financial econometrics Kordonis et al. (2016), hence the stock price and returns predictions are essential parts of the stock market analysis. Usually, the forecast is done with the historical stock price data, however media sentiments also affect the stock prices fluctuations Shah et al. (2018) and could be included in the prediction models. In general, sentiment scores are calculated based on news articles or social media/microblogging posts. The forecasting with microblogging data was studied by different researchers. For example, it was shown Kordonis et al. (2016) that there was a correlation between Twitter sentiments and stock prices. Also, another research Cazzoli et al. (2016) demonstrated that Twitter posts related to corporations could predict the financial market.

Moreover, stock-specific sentiments have a bigger impact on returns than market-specific sentiments Anusakumar et al. (2017). Furthermore, it was shown Wolf and Bergdorf (2019) that the sentiments derived from Twitter were useful in the individual stock returns predictions. In a recent study done during the COVID period, it is shown that tweets containing the term *stocks* have a substantial decline in log returns for US indices Goel et al. (2020). Also, there is a significant correlation between the changes in stock prices and the publication of news articles Mohan et al. (2019). Using news sentiments as a predictor variable leads to a directional accuracy of 70.59 per cent in short-term stock price movement trends prediction Shah et al. (2018).

Prediction computations can be done with econometrical and statistical approaches or with the adaptation of ML and deep learning algorithms. And nowadays, more and more researchers prefer using the ML algorithm for predictions or combine them with traditional methods. The usage of ML algorithms allows improving accuracy and overcome limitations of common econometrical models Rossi (2018). Moreover, the ML algorithms outperform the benchmark buy-and-hold strategy in the real-life simulations and the gradient boosting machine performs the best from the perspective of the statistical and economic evaluation criteria Nevasalmi (2020).

Deep Learning algorithms also show potential in stock returns predictions as they can analyze complex patterns and interactions in the dataset Vargas et al. (2017). Also, deep neural networks could outperform shallow neural networks, and some of them could even

outperform representative machine learning models Abe and Nakayama (2018). Typically, Convolutional Neural Networks (CNN) and Recurrent Neural Networks (RNN) can be used for stock market forecast Vargas et al. (2017). The above-mentioned deep and ML algorithms could be combined with sentiment analysis to receive better accuracy of returns. For example, the RNN models that used news articles text in the input performed better than ones that predicted future stock prices based only on historical stock prices Mohan et al. (2019). CNN outperform RNN on catching textual semantics, and RNN better catches the context information and performs better in modelling complex temporal characteristics Vargas et al. (2017).

3   Methodology

This chapter focuses on predictive modelling Shmueli (2011) and predictive powers of each chosen stock returns prediction method.

However, before constructing the models, it is crucial to understand the timing of the variables. The return vector is calculated using the log-difference formula, which is commonly used in finance:

$$r_t = log(P_t/P_{t-1}) = log(P_t) - log(P_{t-1}) \tag{1}$$

where P_t is the adjusted closing price of the stock market index. Therefore, a return is calculated as the change between the closing values of subsequent trading days. Consequently, any news published in the news sources played a role in deciding what would be the next closing price; hence, the return.

The model should also consider that returns are available only for the trading days, but the news is published daily, including the weekends and holidays. Ignoring the news data from the non-trading days would result in losing important information. The news that appears during the weekends and holidays affect the investors' behaviour, and this is reflected in the stock prices in the first subsequent trading day. The return that is calculated as the log-difference of closing prices before and after a weekend or holiday contains the news information during these days. That's why, in this study, the news data from weekends and holidays is merged to the first following trading day's news data. In fact, these effects are called weekend (a.k.a. Monday) and holiday effects in financial econometrics literature Basher and Sadorsky (2006), Marrett and Worthington (2009). Hence, a proper model should also consider the accumulated impact of the news from weekends and holidays.

Another issue to focus on is the asymmetric effect of the news on returns. It has been documented that the negative return shocks affect the returns and volatilities differently than the positive return shocks (See Maheu and McCurdy (2004), Puzanova and Eratalay

(2020),Ensor et al. (2020), Engle and Ng (1993)). The models in this paper incorporate this asymmetric effect of the news by considering positive and negative news sentiment scores separately.

3.1   ARMAX model

Taking into account the discussion above and adapting the approach of Puzanova and Eratalay (2020), the following ARMAX model is constructed:

$$
\begin{aligned}
r_t = {}& \mu + \mu_W D_t^W + \mu_{NO} D_t^{NO} + \mu_{PO} D_t^{PO} \qquad\qquad (2) \\
& + \sum_{i=1}^{p} \beta_i r_{t-i} + \varepsilon_t + \sum_{j=1}^{q} \theta_j \varepsilon_{t-j} \\
& + \delta_{1N} News_t^N + \delta_{1P} News_t^P \\
& + \delta_{2N} News_{t-1}^N + \delta_{2P} News_{t-1}^P \\
& + \gamma_{1N} Newscount_t^N + \gamma_{1P} Newscount_t^P \\
& + \gamma_{2N} Newscount_{t-1}^N + \gamma_{2P} Newscount_{t-1}^P
\end{aligned}
$$

where D_t^W is a dummy variable for the weekend and holiday effects; D_t^{NO} and D_t^{PO} are dummy variables for negative and positive outliers[4] respectively; $News_t^N$ and $News_t^P$ are negative and positive merged news sentiment scores, respectively; and $Newscount_t^N$ and $Newscount_t^P$ are the number of negative and positive news that occurred at day t. The ARMAX orders p and q are chosen by comparing the AIC values of the model estimates. The autoregressive parameters β_i and moving average parameters θ_j are restricted to satisfy the stationarity and invertability restrictions, respectively.

3.2   Sentiment analysis

To improve the prediction accuracy, we use news articles' sentiment-score as prediction variables. Hence, each article of the dataset was analyzed and assessed with a sentiment analysis algorithm. In general, the sentiment analysis is used to identify opinions expressed in the textual form, and it is based on a natural language processing algorithm where each word of the text has its sentiment-score (positive, neutral or negative) Luo et al. (2013). Sentiment analysis can be performed using supervised and non-supervised approaches. The non-supervised approach represents the classification done "based on a dictionary-based approach to convert the qualitative news articles into a quantitative measure" Li et al. (2018).

---

[4]Outliers were identified using Hampel filtering and Python calculations

The supervised approach of sentiment analysis uses historical trends and news patterns, and creates training data by automatic labelling of news and social media posts Yadav and Kumar (2019). In this research, we applied Valence Aware Dictionary and Sentiment Reasoner (VADER) method. This is a semi-supervised algorithm, which is a simple rule-based model for general sentiment analysis Hutto and Gilbert (2015) to the collected news articles. For each day, we calculated average news score. All the news articles with neutral sentiment scores (score equals to zero) were dropped from the calculation. Also, we have counted numbers of positive and negative news per day to use them as separate independent variables.

3.3   Machine learning and deep learning modelling

As an alternative to ARMAX model in this study, we used two ML algorithms and one deep learning neural network to predict stock returns: eXtreme Gradient Boosting (XGBoost), K-Nearest Neighbors (KNN), and Long Short-Term Memory (LSTM) regression models. These algorithms were chosen due to their high accuracy in the regression forecasting. For these algorithms, we used the same dataset, model, and dataset split ratio as in ARMAX model to create fair conditions for comparison. All the computations were performed in *Python*.

XGBoost: The first algorithm we applied to the chosen regression model was XGBoost ML algorithm designed for efficacy, computational speed and model performance that demonstrates good performance in solving regression and classification problems Malik et al. (2020). XGBoost is a tree boosting method that is considered a highly effective and widely used ML approach that can solve practical problems using a minimal amount of resources Chen and Guestrin (2016). While building the regression XGBoost uses a loss function to evaluate the prediction model. In particular, the XGBoost prediction model was constructed by using the *xgboost* library and the *xgboost.XGBRegressor* function.

KNN: The second ML algorithm we used for returns prediction was KNN. KNN algorithm is a widespread ML algorithm for regression analysis. Its' choice is justified by its simplicity and easy adaptation process, hence it is commonly used for time series analysis and forecast Ban et al. (2013). In the KNN regression algorithm, the dependent variable of a time series forecast is described as a sequence of interval scaled values. Then, based on the pattern, the KNN algorithm identifies similar past patterns and combines their future values to form predictions Ban et al. (2013). The KNN model was created with *sklearn.neighbors* library and *KNeighborsRegressor* function.

LSTM: Finally, we created the LSTM regression model. LSTM is a type of recurrent neural network (one of the general classes of neural networks) deep learning-based algorithm which is commonly used in times series forecasting Elsworth and Güttel (2020), Sherstinsky

(2020). The LSTM-based regression algorithm is a multi-step univariate forecast algorithm that demonstrates a good accuracy in processing the dependency among the dependent variables Siami Namini et al. (2018). The LSTM-based regression was estimated with *Keras* library.

The main issue with ML and deep learning algorithms is that they usually require big volumes of data to properly learn and provide accurate results. In this study, we were using a relatively small dataset, so we were also testing whether the chosen algorithms could outperform ARMAX-based prediction with the small volumes of information to process.

3.4   Dataset

The dataset used for ARMAX and ML modelling consists of SP100 historical data, the negative and positive sentiment scores and the number of news.

The news data was collected using the web scraping method. In total, we have collected over 6000 news articles related to the COVID-19 pandemic that was later cleaned, pre-processed (duplicates were deleted, and the dataset was sorted by date) and were used to conduct sentiment analysis. This dataset covers the period from 27.01.2020 until 10.12.2020. Even though the spread of COVID-19 started in December 2019, there were not enough news articles data for that time to make a sufficient analysis. Hence the data collection started from 27.01.2020 to have the sufficient number of news regarding COVID-19 published on a daily basis. The dataset gives 223 observations for the returns and 11 exogenous variables. The return observations contain outliers identified by Hampel filtering that are shown in Figure 1a. To have a visual presentation of the returns distribution, a quantile-to-quantile (QQ) plot is presented in Figure 1b.

The descriptive statistics can be found in Table 2. The skewness and in particular the kurtosis values suggest that the returns are not normally distributed. This is confirmed by the p-value of the Jarque-Bera test of normality for the returns. What is partially responsible for the high kurtosis is the existence of outliers. In the QQ plot presented in Figure 1b, it can be seen that there are some positive and negative outliers at the top right and bottom left of the figure, respectively. These outliers are identified by the Hampel filter as shown in Figure 1a, where time series plots of the true returns and the returns cleaned from these outliers are presented. It should be noted that the identified outliers in the return series are not removed or smoothed out. Instead, as shown in Section 3.1, we add dummy variables to the model to control for the positive and negative outliers.

Table 1: Descriptive statistics of the SP100 returns for the data period.

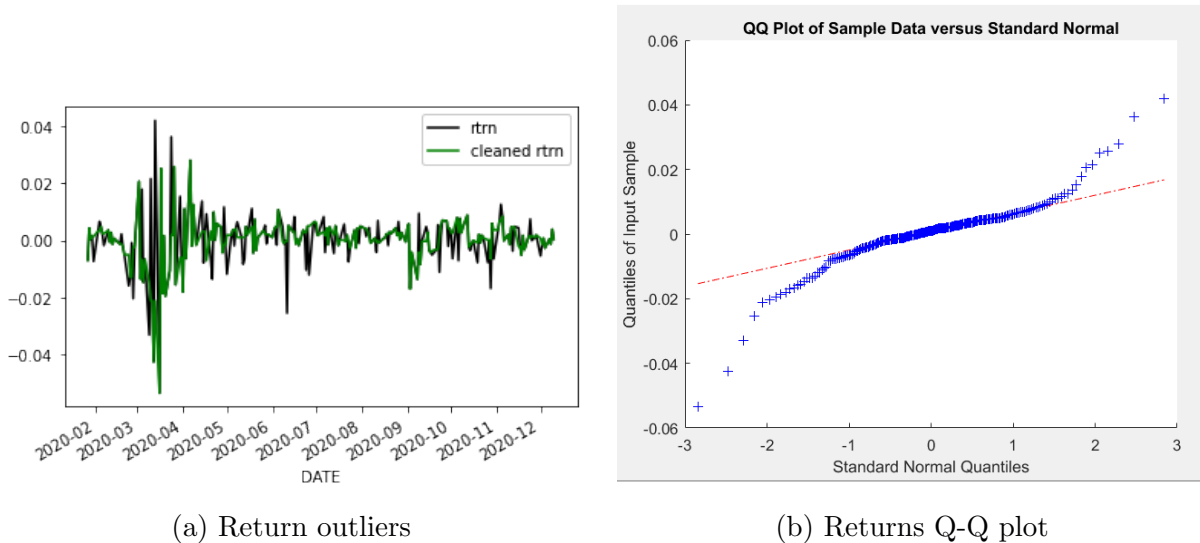| Mean	Median	Variance	Skewness	Kurtosis	JB test pval.
0.0002	0.0010	0.0001	-0.6850	9.9869	0.0010

(a) Return outliers

(b) Returns Q-Q plot

Figure 1: Return outliers and QQ-plot of the returns.

4   Results and discussions

We chose the mean absolute error (MAE) and root mean squared error (RMSE) as comparison indicators of predictive performances of the models. The dataset was split in training and testing dataset to perform the prediction modelling. The training dataset includes the period from 27.01.20 to 31.08.20 (68 per cent of the whole dataset) and the test dataset covers the period from 01.09.2020 until 10.12.20 (32 per cent of the whole dataset). This corresponds to almost 70:30 per cent split, which is a splitting ratio usually used in ML.

The ARMAX model orders of the equation 2 were chosen as p=1 and q=1, by identifying the ARMAX specification that gave the lowest AIC value. In Figure 2, the histogram of the residuals, in modulus, of the ARMAX model estimation is plotted. Most of the residuals are concentrated towards zero, while a few of the residuals lie in the tail of the histogram. A Ljung-Box test on the residuals up to 15 lags showed that the autocorrelation in the returns is successfully captured by the ARMAX model.

Despite the fact that the ARMAX model requires stationarity and invertibility restrictions on the model parameters (see Section 3.1), it is interesting to see that it learned on the training set well. Figure 3a shows that the ARMAX model was able to predict closely most of the highest and lowest returns, probably because of the dummy variables for the outliers in the model. However, its predictions were smoother than the true returns and couldn't catch the variation in it so well. In contrast, the ARMAX model produced better results for the test dataset. As Figure 3b suggests, it wasn't able to predict the high and low points as much in the test dataset, but the spread of the predictions was slightly higher, which resulted in better predictive performance. The prediction MAE results were *0.00511* for the
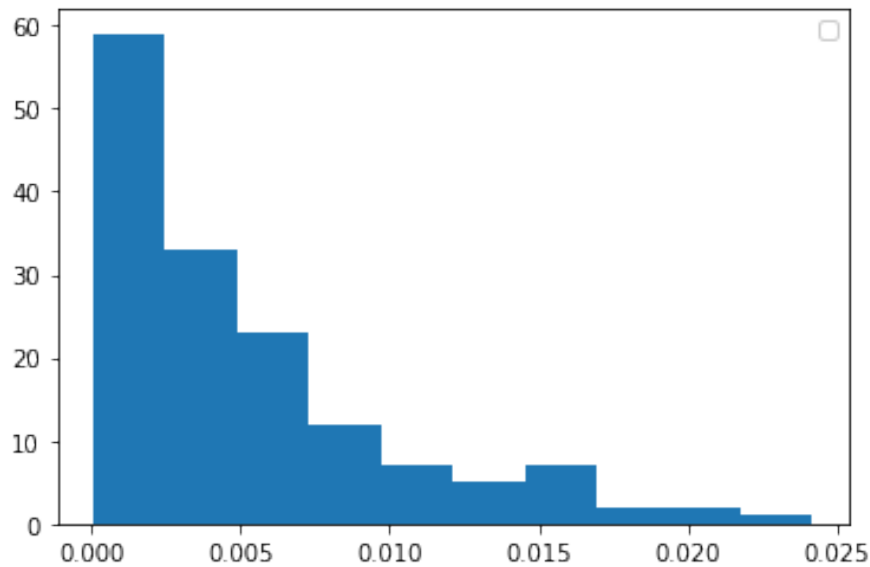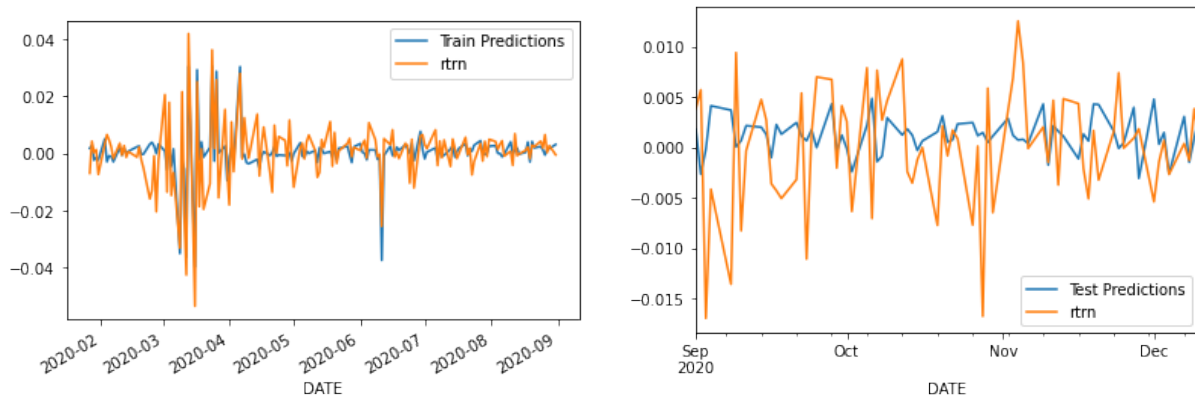
Figure 2: Returns residuals

training dataset and *0.00487* for the test dataset. Meanwhile, not all the ML algorithms outperformed the ARMAX model in predictive performance. (see Table 2).
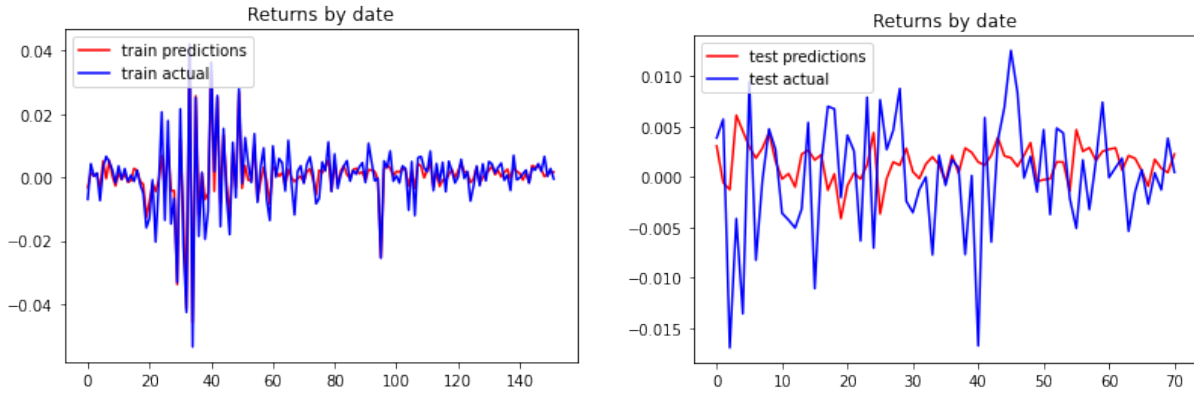

(a) ARMAX training dataset returns predictions   (b) ARMAX test dataset returns predictions

Figure 3: ARMAX Prediction results.

The results in Table 2 show that the XGBoost algorithm gave the best prediction result on the given dataset. Figures 4a and 4b show the XGBoost predictions of the returns and the true returns in the training set and testing set, respectively. the XGBoost algorithm was able to learn very well from the training set, although the predictive performance in the training set is not so much above the one for the ARMAX model. The deep learning (LSTM) model could not properly learn due to the small sample size, and it was outperformed by ARMAX predictions both in the training and testing sets. The results were worse for KNN: the MAE and RMSE results for the KNN were almost double the ones for the ARMAX

Table 2: Results for stock returns prediction models.

Method	MAE Training	MAE Test	RMSE Training	RMSE Test
ARMAX	0.00511	0.00487	0.00718	0.00640
XGBoost	**0.00316**	**0.00479**	**0.00433**	**0.00639**
KNN	0.00544	0.00534	0.00893	0.00682
LSTM	0.01049	0.00742	0.01347	0.00854

(a) XGBoost training dataset returns predictions (b) XGBoost test dataset returns predictions

Figure 4: XGBoost Prediction results.

predictions. The problem could again be connected to the small dataset size.

To summarize, we can see that the XGBoost algorithm outperformed ARMAX for the training dataset, but gave a similar performance with ARMAX in the testing dataset. The other ML approaches couldn't perform that well. One could conclude here that the XGBoost was the most suitable algorithm for this specific sample, followed by the ARMAX model. It is important to point out here that the findings in this analysis only apply to this particular data, which is not a very large sample, volatile and with some outliers.

5   Conclusions and future research

The impact of information on stock markets was investigated by many researchers. The previous studies suggest that there is a correlation between news and media sentiments and stock returns. This chapter contributes to the literature in several dimensions. On the one hand, the effect of the news sentiments on the returns of the SP100 index was analyzed considering the possibility of the asymmetric effect of negative and positive news. On the other hand, the analysis was conducted using the period when the markets were very volatile and very sensitive to the news about COVID-19. Lastly, the analysis compared the predictive performance of ARMAX and the ML algorithms. The results of the analysis demonstrate that the sentiments score inclusion and usage of the ML algorithm significantly

increase the accuracy of the prediction. We found that the XGBoost prediction model showed the best results and had the highest predictive power. In terms of comparing the predictive performances of the mentioned models, this is not an exhaustive study. Therefore the findings only relate to the specific data and period under consideration. Future research could extend the comparison of the predictive performances by increasing the sample size of the returns and considering many different data characteristics related to the distribution of the returns. Moreover, an exhaustive simulation study on the comparison of these methods using data with many different statistical properties is planned by the authors in future research. There could be many factors that could be considered for this comparison, some of which are the volatility of the data at hand, the number of outliers in the training and testing sets, the autocorrelation structure, structural breaks, misspecification of the distribution.

Bibliography

Abe, M. and Nakayama, H. (2018). Deep learning for forecasting stock returns in the cross-section.

Anusakumar, S., Ali, R., and Hooy, C. (2017). The effect of investor sentiment on stock returns: Insight from emerging asian markets. *Asian Academy of Management Journal of Accounting and Finance*, 13:159–178.

Arafat, J., Habib, M. A., and Hossain, R. (2013). Analyzing public emotion and predicting stock market using social media. pages 265–275.

Ban, T., Zhang, R., Pang, S., Sarrafzadeh, A., and Inoue, D. (2013). Referential knn regression for financial time series forecasting. pages 601–608.

Basher, S. A. and Sadorsky, P. (2006). Day-of-the-week effects in emerging stock markets. *Applied Economics Letters*, 13(10):621–628.

Bruhn, M., Schoenmueller, V., and Schäfer, D. (2012). Are social media replacing traditional media in terms of brand equity creation? *Management Research Review*, 35:770–790.

Cazzoli, L., Sharma, R., Treccani, M., and Lillo, F. (2016). A large scale study to understand the relation between twitter and financial market. In *2016 third European network intelligence conference (ENIC)*, pages 98–105. IEEE.

Chen, T. and Guestrin, C. (2016). Xgboost: A scalable tree boosting system. pages 785–794.

Coyne, S., Madiraju, P., and Coelho, J. (2017). Forecasting stock prices using social media analysis. *2017 IEEE 15th Intl Conf on Dependable, Autonomic and Secure Computing, 15th Intl Conf on Pervasive Intelligence and Computing, 3rd Intl Conf on Big Data Intelligence and Computing and Cyber Science and Technology Congress(DASC/PiCom/DataCom/CyberSciTech)*, pages 1031–1038.

Elsworth, S. and Güttel, S. (2020). Time series forecasting using lstm networks: A symbolic approach.

Engle, R. F. and Ng, V. K. (1993). Measuring and testing the impact of news on volatility. *The journal of finance*, 48(5):1749–1778.

Ensor, K. B., Han, Y., Ostdiek, B., and Turnbull, S. M. (2020). Dynamic jump intensities and news arrival in oil futures markets. *Journal of Asset Management*, 21(4):292–325.

Goel, R., Ford, L. J., Obrizan, M., and Sharma, R. (2020). Covid-19 and the stock market: evidence from twitter. *arXiv preprint arXiv:2011.08717.*

Gu, S., Kelly, B., and Xiu, D. (2018). Empirical asset pricing via machine learning. *International Political Economy: Globalization eJournal.*

Hanusch, F. and Tandoc, E. C. (2019). Comments, analytics, and social media: The impact of audience feedback on journalists' market orientation. *Journalism*, 20:695 – 713.

Hutto, C. and Gilbert, E. (2015). Vader: A parsimonious rule-based model for sentiment analysis of social media text.

Khan, W., Ghazanfar, M. a., Azam, M. A., Karami, A., Alyoubi, K., and Alfakeeh, A. (2020). Stock market prediction using machine learning classifiers and social media, news. *Journal of Ambient Intelligence and Humanized Computing.*

Kordonis, J., Symeonidis, S., and Arampatzis, A. (2016). Stock price forecasting via sentiment analysis on twitter.

Leung, W. S., Wong, W. K., and Wong, G. (2019). Social-media sentiment, limited attention, and stock returns. *Economics of Networks eJournal.*

Li, M., Yang, C., Zhang, J., Puthal, D., Luo, Y., and Li, J. (2018). Stock market analysis using social networks. ACSW '18, New York, NY, USA. Association for Computing Machinery.

Luo, T., Chen, S., Xu, G., and Zhou, J. (2013). *Sentiment Analysis*, pages 53–68.

Maheu, J. M. and McCurdy, T. H. (2004). News arrival, jump dynamics, and volatility components for individual stock returns. *The Journal of Finance*, 59(2):755–793.

Malik, S., Harode, R., and Kunwar, A. (2020). Xgboost: A deep dive into boosting (introduction documentation).

Malkiel, B. (2003). The efficient market hypothesis and its critics. *Journal of Economic Perspectives*, 17:59–82.

Marrett, G. J. and Worthington, A. C. (2009). An empirical note on the holiday effect in the australian stock market, 1996–2006. *Applied Economics Letters*, 16(17):1769–1772.

Mohan, S., Mullapudi, S., Sammeta, S., Vijayvergia, P., and Anastasiu, D. (2019). Stock price prediction using news sentiment analysis. pages 205–208.

Nevasalmi, L. (2020). Forecasting multinomial stock returns using machine learning methods. *The Journal of Finance and Data Science*, 6:86–106.

Puzanova, Y. and Eratalay, M. H. (2020). Effect of real estate news sentiment on the stock returns of swedbank and seb bank. *Available at SSRN 3571166*.

Rossi, A. G. (2018). Predicting stock market returns with machine learning.

Shah, D., Isah, H., and Zulkernine, F. (2018). Predicting the effects of news sentiments on the stock market.

Sharma, S. and Sharma, R. (2020). Forecasting transactional amount in bitcoin network using temporal gnn approach. *IEEE/ACM International Conference on Advances in Social Networks Analysis and Mining (ASONAM)*.

Sherstinsky, A. (2020). Fundamentals of recurrent neural network (rnn) and long short-term memory (lstm) network. *Physica D: Nonlinear Phenomena*, 404:132306.

Shmueli, G. (2011). To explain or to predict? *Statistical Science*, 25.

Siami Namini, S., Tavakoli, N., and Siami Namin, A. (2018). A comparison of arima and lstm in forecasting time series. pages 1394–1401.

Steyn, D. H. W., Greyling, T., Rossouw, S., and Mwamba, J. M. (2020). Sentiment, emotions and stock market predictability in developed and emerging markets.

Vargas, M., Lima, B., and Evsukoff, A. (2017). Deep learning for stock market prediction from financial news articles. pages 60–65.

Wolf, F. and Bergdorf, O. (2019). Twitter sentiment and stock returns.

Yadav, R. and Kumar, A. (2019). News-based supervised sentiment analysis for prediction of futures buying behaviour. *IIMB Management Review*, 31.