Jan Vogt, Alexander Bönner, Michael Römmich, Malte Weiß,
Merih Türkoglu

Chapter 16
Energy Stock Price Forecast Based on Machine Learning and Sentiment Analysis – Which Approach Performs Best in Day Trading?

Abstract: We explore the application of machine learning (ML) methods to predict energy stock prices. We apply established ML methods, Gradient Boosted Regression Trees (GBRT), Multilayer Perceptron (MLP) and Long Short-Term Memory (LSTM) trained on the energy market stock data retrieved from Uniper from January 2019 to August 2020. Furthermore, we incorporate sentiment data linked to Twitter contributions in the aforementioned time period. We apply these algorithms to predict the next day's close price for three Energy companies (Uniper, Enel, EDF), and simulate a buy-and-hold trading strategy to measure the performance of our models. Our results indicate that MLP yields the most accurate predictions with the lowest mean absolute error. Applied to Uniper stock market data, our trading simulation significantly outperforms the buy-and-hold benchmark. Furthermore, the results show that the use of sentiment values improves trading performance significantly.

Keywords: energy market, stock price prediction, sentiment analysis, LSTM, GBRT, MLP

16.1 Introduction

In recent years, machine learning (ML) algorithms have been applied to various fields, such as physical simulations, weather forecasts, programming and language generation. In particular, stock market predictions represent a lucrative application that has gained significant attention (Islam, Hasan & Khan 2021; Liu, Dang & Yu 2020;

Jan Vogt, sipgate GmbH, Düsseldorf, e-mail: contact@j-vogt.com
Alexander Bönner, Hochschule Ruhr West, Mülheim a.d. Ruhr, e-mail: alexander.boenner@hs-ruhrwest.de
Michael Römmich, Hochschule Ruhr West, Mülheim a.d. Ruhr, e-mail: michael.roemmich@hs-ruhrwest.de
Malte Weiß, Hochschule Ruhr West, Bottrop, e-mail: malte.weiss@hs-ruhrwest.de
Merih Türkoglu, Hochschule Ruhr West, Bottrop, e-mail: merih.tuerkoglu@hs-ruhrwest.de

https://doi.org/10.1515/9783110790320-016

Wang & Guo 2020). While various market models have been developed using ML algorithms, there is little evidence of the use of ML algorithms in the energy stock market.

In this chapter, we explore three ML methods to predict stock prices in the energy market. We employ Gradient Boosted Regression Trees (GBRT), Multilayer Perceptron (MLP) and Long Short-Term Memory (LSTM) to create predictive models for three energy companies (i.e., Uniper, Enel and EDF). Our input vector consists of two data sets. The first data set constitutes stock prices (high, low, open and close) of Uniper in the period from January 2019 to August 2020. The second data set incorporates sentiment data extracted from Twitter publications of Uniper. We analyse the accuracy of stock predictions with and without sentiment data, and evaluate our model using a buy-and -hold strategy as a benchmark.

16.2 Related Work

16.2.1 Machine Learning and Sentiment Analysis Applied on Stock Price Predictions

Achkar et al. (2018) used an MLP and the LSTM to predict the daily closing prices of Alphabet, Facebook and Bitcoin shares. Their analysis indicates that the deviations measured as mean squared error (MSE) used to predict the closing price ranged from 3% to 16% for the MLP model and from 0.5% to 12% for the LSTM model, respectively. The authors find that neural networks can be used as an effective and promising tool for stock market predictions.

Torres et al. (2019) used an MLP and a decision tree to make predictions for the Apple stock. Historical price data (open, close, high, low and trading volume) over the last 250 trading days built the foundation of their analysis. The utilised methodology aimed to predict the closing price of a share on a particular day and, also, shows tolerable errors mean absolute error (MAE) results. The authors conclude that both methods are suitable for predicting stock prices and note that sentiment analysis could further enhance the accuracy of predictions.

Kolasani and Assaf (2020) investigated the effectiveness of using tweets to predict stock prices. Moreover, the authors looked at the effectiveness of neural networks in predicting stock price movements in contrast to traditional ML. The study elaborates on a study by Chakraborty et al. (2017) to predict price movements of the stock of the Apple and the Dow Jones index. Both studies examined several models, such as Logistic Regression, Support Vector Machine, Decision Tree, GBRT and Random Forests, to represent the manually tested sentiments in an ML model. Kolasani and Assaf (2020) found that, on average, the MLP performs better in terms of MAE and root mean squared error (RSME) in predicting the price difference of stocks than the GBRT. The authors note that the GBRT tends to be a more optimistic approach, while the MLP tends to be more pessimistic in predicting a price difference.

Hutto and Gilbert (2014) implement a lightweight and rule-based model for general sentiment analysis called Valence Aware Dictionary for sEntiment Reasoning (VADER). It can be applied to social media style text and multiple domains such as news services or ratings. The model includes a generalisable, valence-based, human-organised lexicon for sentiment. VADER provides four values "positive," "negative," "neutral" and "composite" to identify sentiment. The attribute "composite" describes a composite value, namely the sum of the valence values of each word in the lexicon, which is calculated and then normalised to a value between –1 (negative) and +1 (positive) (Hutto & Gilbert, 2014).

Deepika and Nirupamabhat (2020) use different models to derive price direction prediction using stock data (open, close, adjusted close, high, low and trading volume) of Apple, Amazon, Infosys, Microsoft, Oracle and TCS. This was combined with technical analysis and sentiment analysis. An optimised Artificial Bee Colony (ABC)-LSTM proves to be the best model in terms of mean absolute percentage error (MAPE) compared to also optimised Least-squares Support Vector Machine (LSSVM), Gradient Boost, LSTM, ABC-LSSVM and ABC-Gradient Boost.

16.2.2 Machine Learning and Trading Strategy

In order to put ML approaches into trading practices, these forecasting approaches are trained and thus optimised with respect to metrics like the MSE, MAE or alike. However, from a trader's perspective, it is important that forecasting models perform well in monetary terms like profits or Sharpe Ratios. The following research articles deal with the performance of translating ML strategies into trading results.

Mittermayer (2004) forecasted intraday stock price trends for a short time immediately after the publication of press releases by his NewsCATS system, which used "Support Vector Machines" (SVM) for ML. Categorised press releases built the basis for trading strategies. The author showed that categorised press releases contain additional information in such a way that NewsCATS short-term trading strategy outperformed trading strategies by randomly buying or shorting stocks directly after the publication of press releases.

Fischer & Krauss (2018) used LSTM networks to predict out-of-sample directional stock movements for S&P 500 stocks only based on constituent S&P 500 stock data from 1992 to 2015. Their approach outperforms by daily returns and by Sharpe Ratio memory-free classification methods like random forests or logistic regression classifiers. Furthermore, the researchers successfully translated the output of the model (i.e., the selected stocks) into a rule-based short-term trading strategy to partially lighten up the black box character of the LSTM network.

Yuan et al. (2020) focussed on the Chinese stock market from 2010 until 2018 and incorporated different methods, including ML applied to long/short-term trading strategies. Their predictions on the stock price trend were based on three approaches: Support Vector Machine (SVM), Random Forests (RF) and an Artificial Neural Network (ANN). Long/short-trading strategies were evaluated by indicators like annualised returns, Sharpe Ratios or profit-loss ratios, etc. Overall, the RF approach performed best in their analysis.

Li et al. (2021) applied tensor theory to present fundamental stock information (high/low prices, volume, turnover, P/E-Ratio, P/B-Ratio, etc.) with media information (positive, negative or divergent media sentiment) to preserve the multifaceted and interrelated nature of both streams of information. An LSTM model is proposed to capture the relations between market information and stock movements. Based on the Sharpe Ratio reflecting the trade-off between risk and return, their approach outperformed trading algorithms in an investment simulation of the Chinese stock market.

16.3 Methodology and Data

16.3.1 Data Procurement and Preparation

Based on the findings from the literature review, we structure the methodology of our study as follows. In the first step, the data set is generated and forms the basis for training neural networks and decision trees to predict future stock price trends. Company profiles of Europe's three major energy groups, Enel, EDF and Uniper are identified. Data used to predict price movements consists of historical share price developments and user tweets retrieved from Twitter (which is used for sentiment analysis). The data was collected for the period from January 2019 to December 2020. Historical stock price movements were obtained from Yahoo Finance using the Python library, yfinance. For each of the three groups, information was collected over 512 trading days. Opening, closing, high and low prices were recorded for each trading day.

The Twitter REST API V2, which is reserved for non-commercial projects, was used to determine user contributions. The data collection was carried out in August 2021. The data collection included the number of tweets found on the public profiles of the three listed companies as can be seen in Table 16.1. Whenever company profiles were mentioned, a correlation between shares and posts on social media was established. For this purpose, all official profiles for each company were identified in advance.

The data was pre-processed for sentiment analysis. Twitter hashtags come in various forms, for example, as individual words such as nouns and adjectives, and as groups of words that are combined primarily with the camel or Pascal spelling. We normalised these hashtags by breaking them down into individual words. Finally, parts in tweets that would not contribute to sentiment analysis were removed before-

Table 16.1: Companies Twitter-Profile with the corresponding numbers of Tweets during the research period from January 2018 to December 2020. (Compiled by authors.)

Company	Twitter-Profile	Numbers of tweets
UNIPER	@uniper_energy	37,179
ENEL	@Enelgroup, @EnelGroup	51,500
EDF	@edfenergy	53,840
		Total: 142,519

hand. These included URLs or evaluative usernames that were used whenever Twitter was mentioned. In addition, up to 43 different languages were identified with Google Language Detection. Following this, the multilingual posts were translated into English as the common target language.

16.3.2 Sentiment Analysis and Data Set Building

In this study, VADER analysis is used to extract sentiments from Twitter posts (Hutto & Gilbert, 2014). VADER was developed specifically for social media vocabulary and is therefore also suitable for Twitter posts. VADER is based on an English vocabulary. Consequently, it is required that messages are exclusively written in English. The process of sentiment analysis is as follows. Firstly, the so-called polarity of the sentiment of each tweet is determined sequentially for each company. Secondly, sentiment values were determined for the tweets. The mean values were calculated as overall sentiment values for a trading day. The sentiment of a day is thus defined by the average mean value of the sentiment for each day. The influence of these attributes could be investigated in later studies. When creating the data set, mean values returned from the sentiment analysis as well as the share price movements for each trading day were combined and formed the input vector for the ML models. The closing price of the following day was used as the target value in this study (t + 1).

16.3.3 Model Development

16.3.3.1 Hyperparameter Setting

Two versions of each of the models are trained, one using historical stock price movements to predict the next day's closing price and the other including sentiment data to investigate the impact. The split is 80:20, with 80% of the data used as training data and 20% of the data used as test data. The training data is used to train the model, while the test data is used for evaluation. For each model, multiple metrics are collected for evaluation. The metrics are primarily used to evaluate the accuracy and secondly to benchmark the two models. The metrics collected are the error values in the form of MSE, RMSE and MAE. In this study, as depicted in Table 16.2, three different levels of complexity are examined for each of the models (GBRT, MLP and LSTM).

Table 16.2: Complexity table. Displaying complexity of each model. (Compiled by authors.)

Complexity	GBRT	MLP	LSTM
Low	Estimator: 1.000 Max Depth: 4	Layers: 4-2-1	Layers: 4-2-1
Medium	Estimator: 1.000 Max Depth: 8	Layers: 8-4-1	Layers: 8-4-1
High	Estimator: 1.000 Max Depth: 16	Layers: 16-8-1	Layers: 16-8-1

In the course of the model developments, a series of measurements of different hyperparameters are carried out in order to compare them with each other and to find a suitable configuration. Different constellations of hyperparameters are tested and the resulting error values from the measurement series are compared with each other. Dropout (0%, 10%, 20%), batch (1, 2, 4) and the learning rate of (0.1, 0.01, 0.001) are considered. Stochastic Gradient Descent (SGD) is used as the optimiser to minimise the neural network errors. Furthermore, the MSE is chosen for the error function.

16.3.3.2 Hyperparameter Tuning with Uniper

Uniper is chosen to determine the hyperparameters. We hypothesise that the configuration of one model leads to similar results for other energy groups. The series of measurements are carried out sequentially, with one parameter being investigated without including another. GBRT, LSTM and MLP require a learning rate as a basis for optimising the model, so we start with this hyperparameter and then continue with batch size and dropout. The investigation is applied to all complexity levels “low,” “medium” and “high.” The MSE is used as a metric to evaluate the investigation.

The MLP results showed an increased error value at a learning rate of 0.1 across all complexity levels. The learning rate of 0.001 resulted in the lowest error values. No large deviations were found in the measurement series for the LSTM but led to the same values as the MLP at a learning rate of 0.001. The error values of the GBRT showed that a learning rate of 0.01 led to the lowest error values at all complexity levels, both for the data set with historical stock values and for the included sentiment with identical hyperparameter settings. This concludes the investigation for the GBRT, as only the learning rate was considered as a hyperparameter. After the learning rate leads to the lowest values in all complexity levels, the consideration of the lot sizes follows. A comparison between 1, 2 and 4 is carried out and applied to the learning rate of 0.001 determined previously.

In the series of experiments of the MLP, batch size 2 and 4 showed a lower error value in contrast to batch size 1. Among them, batch size 4 gave the lowest error values for MSE Stock and MSE Sentiment in all complexity levels. The batch size results for LSTM differ from the previously collected data. For the complexity level “simple” with the structure 4-2-1, lot size 2 leads to the lowest error values. For the complexity level “medium” and “complex,” the lowest error values were measured for lot size 4. For the following consideration of the LSTM of the dropout, different batch sizes are therefore used for the complexity levels.

After the learning rate and the batch size were determined, the last missing hyperparameter, dropout, was determined. Here, the hyperparameters that resulted in the lowest error values in the previously determined tests were used. The results of the MLP as well as the LSTM show that in all complexity levels, the use of dropouts leads to higher error values for both MSE-Stock and MSE-Sentiment.

As a result, for the remainder of the study, the GBRT is used with a depth of 4 and the estimator of 1000 and a learning rate of 0.01. The MLP is used with the hyperparameters dropout 0%, LR 0.001, batch size 4 and the model setup of 16-8-1. The LSTM gave the same results as the MLP but with the difference that the model setup of 8-4-1 is used in this case.

16.4 Results

We use error functions as metrics like the L2-Error to mathematically validate the outcomes of each model for all companies. Then, using a variety of day trading methods based on the model, we test if the mathematical results correspond to actual world outcomes.

16.4.1 Error Scores and Analysis

The historical stock model and the sentiment model from 17.3 were individually applied to MSE, RMSE and MAE metrics, respectively. With the use of this analysis, we can determine whether the target value and prediction have any relationship to stock or sentiment values. The models have been trained on Uniper’s data and tested on three companies: Uniper, Enel and EDF. The input for the stock models included open, low, high and close values of the prior trading day. For the sentiment model, the same values along with the mean sentiment on the previous trading day were employed. The evaluation period was between 14 August 2020 and 29 December 2020 for each company with a total of 100 trading days. The following Table 16.3 displays the findings for each company and model assessed using the MSE. For more measurement results, please see the appendix.

Table 16.3: Analysis on each model with MSE. (Compiled by Authors.)

ML-Model	Model	Error	Uniper	Enel	EDF
GBRT	STOCK	MSE	0.1404 . . .	0.0195 . . .	0.0768 . . .
GBRT	SENTIMENT	MSE	0.1410 . . .	0.0228 . . .	0.0719 . . .
MLP	STOCK	MSE	0.1305 . . .	0.0151 . . .	0.0564 . . .
MLP	SENTIMENT	MSE	0.1287 . . .	0.0179 . . .	0.0570 . . .
LSTM	STOCK	MSE	0.1455 . . .	0.0468 . . .	0.0825 . . .
LSTM	SENTIMENT	MSE	0.1422 . . .	0.0405 . . .	0.0770 . . .

16.4.1.1 Uniper

The analysis shows that the MSE of Uniper GBRT using the stock model is 0.1404 US dollar (USD), MLP 0.1305 USD and LSTM 0.1455 USD, while the sentiment model GBRT is 0.1410 USD, MLP 0.1287 USD and LSTM 0.1422 USD (see Table 16.3 and Appendix). MLP and LSTM show a slightly better performance with sentiment, while GBRT performs worse. Interestingly, Uniper's analysis results have higher error values for MSE, RMSE and MAE compared to the other two companies. The probable reason for this is that Uniper's share value per stock is nearly twice of Enel and EDF. Figure 16.1 displays the 100-day evaluation period and the respective prediction of both MLP models. Stock and sentiment models have a small difference of 0.0018 USD on the analysis, but the figure exposes larger differences across all trading days.

16.4.1.2 Enel

As for Enel the GBRT value is 0.0195 USD, MLP 0.0151 USD and LSTM 0.0468 USD. While for the sentiment model, the GBRT is 0.0228 USD, MLP 0.0179 USD and LSTM 0.0405 USD.

In contrast to Uniper, only LSTM gained an improvement in performance by using sentiment values, while GBRT and MLP displayed a worsening in performance. Figure 16.2 shows the MLP predictions of Enel. Like Uniper, the projections are pessimistic, which means that most of them are below the target value. On some dates, the real values also have a small visual delay, for example between 1 November 2020 and 15 November 2020.

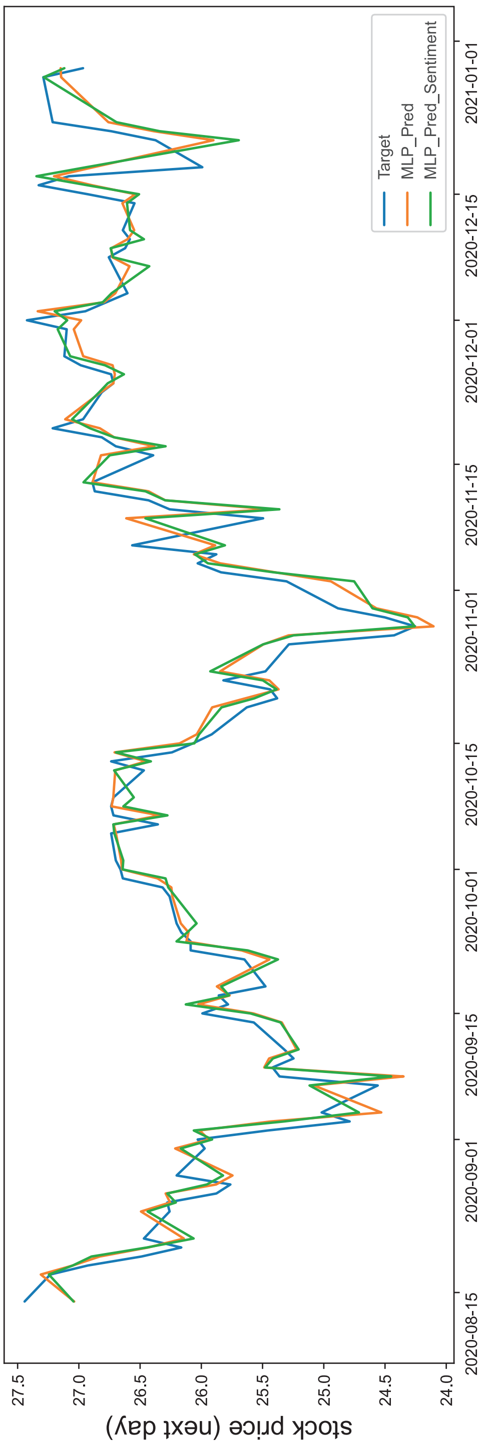

Figure 16.1: MLP validation on Uniper's 100 trading days. Predictions of stock and sentiment MLP. Target is the real value for the time period, while orange shows stock model prediction and green displays sentiment model. Sentiment and stock show strong similarities. (Compiled by authors.)

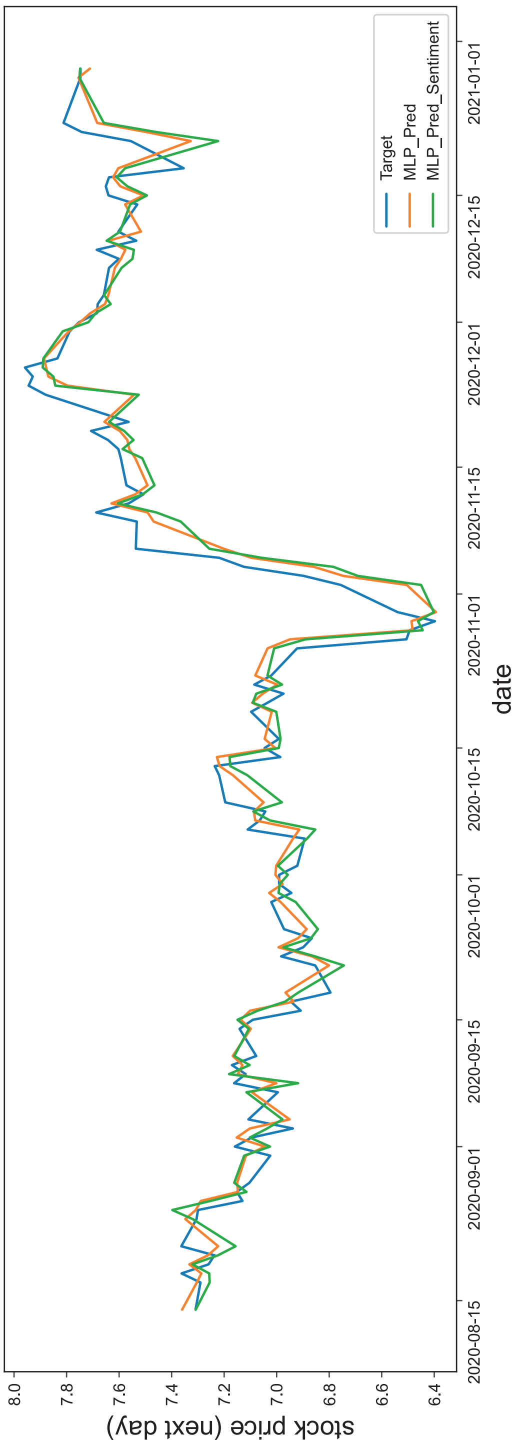

Figure 16.2: MLP Validation on Enel. (Compiled by authors.)

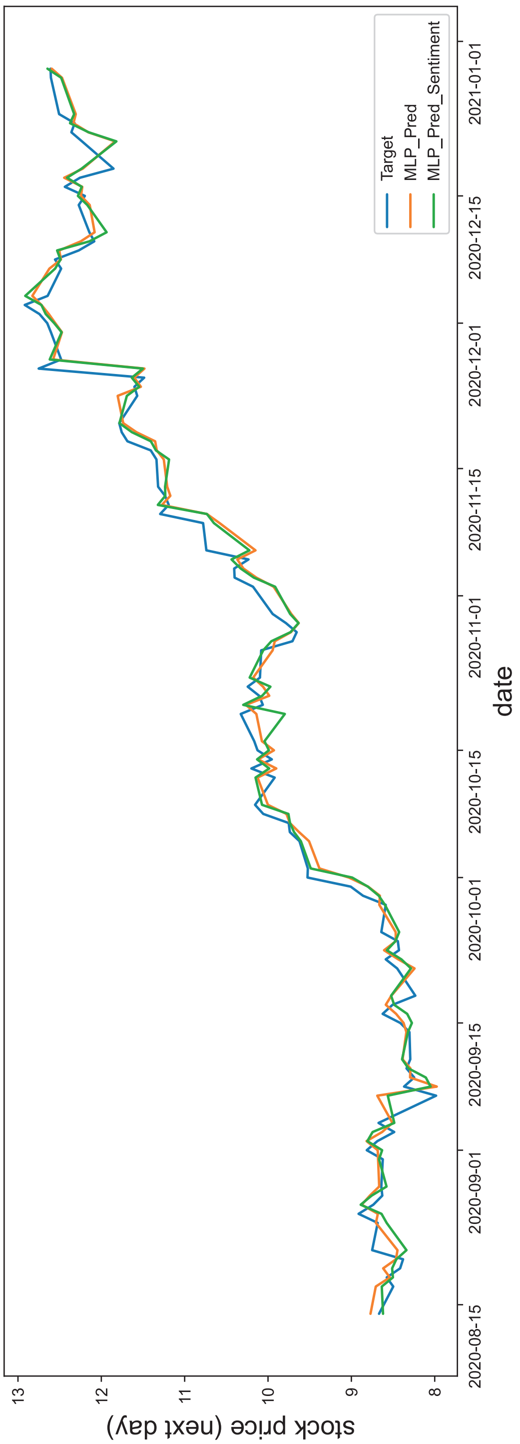

Figure 16.3: MLP Validation on EDF. (Compiled by authors.)

16.4.1.3 EDF

For EDF, the GBRT stock model produced values of 0.0768 USD, 0.0564 USD for MLP and 0.0825 USD for LSTM, while the sentiment model result in values of 0.0719 USD for GBRT, 0.0570 USD for MLP and 0.0770 USD for LSTM, respectively. The MLP's prediction for the EDF is shown in Figure 16.3. Compared to Uniper and Enel, the sentiment and stock prices seem to be more in line here.

Given that the sentiment has somewhat improved performance on the MSE, except for MLP, the results for the MSE on EDF are completely different from those for Uniper and Enel. MSE performed somewhat worse with sentiment for the MLP. Overall, all models for EDF have slightly improved sentiment performance.

In our study of the three distinct companies, MLP shows the best performance and the lowest error rates across all three metrics. Interestingly, the LSTM produced the highest error for the three companies.

16.4.2 Applied Trading Strategy

16.4.2.1 Derivation and Implementation of a Straightforward Trading Strategy

In the following, a trading strategy is derived for the stocks of Uniper, Enel and EDF. The basis for the trading strategy are the forecast values of the model with the most robust performance, the MLP with layers 16-8-1, a learning rate of 0.001, a batch size 4 and a dropout 0. Both variants, that is, the forecast values with sentiment analysis and without sentiment analysis, are examined for the test data period.

The aim is to achieve the best possible return compared to a buy-and-hold strategy (benchmark) for the next 100 trading days. The investable amount of money at day 1 is 100 EUR. Following a signal, it is possible to buy the shares of the respective company at the open price at the beginning of a trading day and automatically sell them at the close price at the end of the trading day. On the other hand, it is also possible, after receiving a respective signal, to sell shares at the value of the investable amount of money at the open price at the beginning of a trading day and automatically buy them back at the close price at the end of the trading day. The resulting amount is assumed to be investable (long or short strategy accordingly) again on the next day. Borrowing fees and transaction costs are not further considered in the analysis. The investment in fractional shares is possible.

In the next step, buy and sell signals are determined. For this purpose, the forecast value for the day, measured at the end of the previous day, is compared to the close price of the previous day. If the forecast value is greater than the close price, shares at the value of the investable amount are purchased at the open price on the following day. In the opposite case, shorting shares at the value of the investable amount at the open price is undertaken.[1]

16.4.2.2 Results and Interpretation

After the 100 trading days of the test data period, the total returns for all three companies are as shown in Table 16.4.

Table 16.4: Total returns of the trading strategy compared with benchmark. (Compiled by authors.)

	Return with sentiment	Return without sentiment	Return benchmark
Uniper	30.28%	15.03%	−6.66%
Enel	10.91%	5.09%	5.56%
EDF	5.12%	−18.06%	46.74%

On the one hand, it can be seen that the model with input of the sentiment analysis leads to significantly better results for all three companies than without input of the sentiment analysis. Moreover, the model with input of the sentiment analysis leads to positive returns for all three companies. According to an applied t-test, Uniper shows a significant outperformance compared to the buy-and-hold strategy (benchmark). Enel´s result is positive compared to the buy-and-hold strategy as well, but not on a significant level. In contrast, the returns forecasted for EDF are outweighed by the buy-and-hold strategy. The same pattern is valid for the respective Sharpe ratios.

A possible explanation for the different performance of the trading strategy compared to the buy-and-hold strategy could be that the share prices of Uniper and Enel tended to move sideways (Uniper slightly down, Enel slightly up) during the test period. On the other hand, the share price of EDF has developed strongly and positively with a significantly higher volatility in share price than Uniper and Enel. This could be interpreted in such a way that the trading strategy works better in less volatile markets without a clear trend but seems less suitable in nervous and strongly trending markets. However, it should be noted that EDF also generated a positive return in the chosen period.

1 The strategy was further refined, including the inclusion of volatility or the previous day’s forecast quality. However, there were no significant improvements in the results.

16.5 Limitations and Future Work

The present analysis is subject to various restrictions. For example, it is limited to three levels of complexity in the construction of the models. Furthermore, the subsequent calibration of the models, including hyperparameter setting, focuses on only one company. Both restrictions are due to the available computer capacities. A detailed investigation of the hyperparameters could lead to different and possibly better results. Other ML methods, such as CNN, could also be applied in principle.

The sentiment analysis was based on Twitter tweets about the respective companies in all languages used. Since free products did not have sufficient translation quality, Google Translate was used. Thus, the sentiment analysis in this study could only be carried out for three companies and especially tweet leader RWE could not be further investigated due to these restrictions.

Furthermore, additional factors such as likes, number of positive and negative comments or the evaluation of the content behind the links used could provide further insights. Likewise, alternative news portals can contribute further useful information on price changes. On top of this, the influence of different time intervals on the sentiment values could be looked at, which include several days or a lag, for example.

In addition, a follow-up study could use tools that are more specialised in sentiment analysis in the financial sector, such as FinBERT, instead of VADER. Furthermore, the study was limited to a period of 2 years due to capacity bottlenecks in the processing of sentiment inputs. A longer period could lead to lower error values.

Especially interesting is the pattern where the error values for the MLP model of Enel and EDF are slightly lower for the stock models whereas the trading strategy performs better for the sentiment models. The signals of the sentiment models seem to be more significant. Why this is the case should be investigated in the next step.

The analysis presented could be made even more sophisticated, especially in the trading section, if the data set of share prices, which is limited to open, close, high and low, was extended to include, for example, shorter time intervals such as minute-by-minute price movements.

For the purposes of finding out whether the hyperparameters can also be applied to other price movements in the energy sector, we transfer the hyperparameters used in the study to the 48 largest energy companies. We use the same time period for this purpose and measure the error value with the MSE. The results based on historical stock price movements are promising. The hyperparameters for ML models yielded the lowest error values for 22 GBRT, 6 MLP and 4 LSTM companies. Looking at the first quartiles according to the hyperparameters obtained for Uniper, of the 48 possibilities, 44 could be counted for GBRT (91.66%), 45 for MLP (93.75%) and 41 for LSTM (85.41%). These results encourage us to perform further studies in this area in the future.

16.6 Conclusion

Exploring the application of various ML techniques showed that sentiment analysis adds value to day-trading strategies to predict stock prices in energy markets. Generally, our results indicate that MLP yields the most accurate predictions with lowest mean absolute error. Applied to Uniper stock, our trading simulation significantly outperforms a 100-day buy-and-hold strategy as a benchmark.

Appendix

Tables 16.5–16.7 show the results of the analysis on each company.

Table 16.5: Results of the different metrics used for each model trained on Uniper. (Compiled by authors.)

ML-Model	Model	MSE	RMSE	MAE
GBRT	STOCK	0.1404 . . .	0.3747 . . .	0.2892 . . .
GBRT	SENTIMENT	0.1410 . . .	0.3756 . . .	0.2897 . . .
MLP	STOCK	0.1305 . . .	0.3612 . . .	0.2751 . . .
MLP	SENTIMENT	0.1287 . . .	0.3586 . . .	0.2569 . . .
LSTM	STOCK	0.1455 . . .	0.3814 . . .	0.2979 . . .
LSTM	SENTIMENT	0.1422 . . .	0.3771 . . .	0.2851 . . .

Table 16.6: Results of the different metrics used for each model trained on Enel. (Compiled by authors.)

ML-Model	Model	MSE	RMSE	MAE
GBRT	STOCK	0.0195 . . .	0.1398 . . .	0.1100 . . .
GBRT	SENTIMENT	0.0228 . . .	0.1509 . . .	0.1152 . . .
MLP	STOCK	0.0151 . . .	0.1229 . . .	0.0931 . . .
MLP	SENTIMENT	0.0179 . . .	0.1337 . . .	0.1017 . . .
LSTM	STOCK	0.0468 . . .	0.2164 . . .	0.1743 . . .
LSTM	SENTIMENT	0.0405 . . .	0.2013 . . .	0.1615 . . .

Table 16.7: Results of the different metrics used for each model trained on EDF. (Compiled by authors.)

ML-Model	Model	MSE	RMSE	MAE
GBRT	STOCK	0.0768 . . .	0.2772 . . .	0.1979 . . .
GBRT	SENTIMENT	0.0719 . . .	0.2682 . . .	0.1971 . . .
MLP	STOCK	0.0564 . . .	0.2374 . . .	0.1725 . . .
MLP	SENTIMENT	0.0570 . . .	0.2387 . . .	0.1719 . . .
LSTM	STOCK	0.0825 . . .	0.2872 . . .	0.2148 . . .
LSTM	SENTIMENT	0.0770 . . .	0.2775 . . .	0.2108 . . .

Bibliography

Achkar, R., Elias-Sleiman, F., Ezzeddine, H., & Haidar, N. (2018). Comparison of BPA-MLP and LSTM-RNN for stocks prediction. *International Symposium on Computational and Business Intelligence (ISCBI)*, 48–51. https://doi.org/10.1109/ISCBI.2018.00019.

Chakraborty, P., Pria, U. S., Rony, M. R. A. H., & Majumdar, M. A. (2017). Predicting stock movement using sentiment analysis of Twitter feed. *International Conference on Informatics, Electronics and Vision*, 1–6. https://doi.org/10.1109/ICIEV.2017.8338584.

Deepika, N., & Nirupamabhat, M. (2020). An optimized machine learning model for stock trend anticipation. *IngéNierie des Systèmes D Information*, *25*(6), 783–792. https://doi.org/10.18280/isi.250608.

Fischer T., & Krauss C. (2018). Deep learning with long-short term memory networks for financial market predictions. *European Journal of Operational Research*, *270*(2), 654–669. https://doi.org/10.1016/j.ejor.2017.11.054.

Islam, S. B., Hasan, M. M., & Khan, M. M. (2021). Prediction of stock market using recurrent neural network. *2021 IEEE 12th Annual Information Technology, Electronics and Mobile Communication Conference (IEMCON)*, 479–483. https://doi.org/10.1109/IEMCON53756.2021.9623206.

Hutto, C. J., & Gillbert, E. (2014). VADER: A parsimonious rule-based model for sentiment analysis of social media text. *Proceedings of the International AAAI Conference on Web and Social Media*, *8*(1), 216–225. https://doi.org/10.1609/icwsm.v8i1.14550.

Kolasani, S. V., & Assaf, R. (2020). Predicting stock movement using sentiment analysis of Twitter feed with neural networks. J*ournal of Data Analysis and Information Processing*, *8*(4), 309–319. https://doi.org/10.4236/jdaip.2020.84018.

Li, Q., Tan, J., Wang, J., & Chen, H. (2021). A multimodal event-driven LSTM model for stock prediction using online news. *IEEE Transactions on Knowledge and Data Engineering*, *33*(10), 3323–3337. https://doi.org/10.1109/TKDE.2020.2968894.

Liu, Z., Dang, Z., & Yu, J. (2020). Stock price prediction model based on RBF-SVM algorithm. *2020 International Conference on Computer Engineering and Intelligent Control (ICCEIC)*, 124–127. https://doi.org/10.1109/ICCEIC51584.2020.00032.

Mittermayer, M.-A. (2004). Forecasting Intraday stock price trends with text mining techniques. *37th Annual Hawaii International Conference on System Sciences*. 10. https://doi.org/10.1109/HICSS.2004.1265201.

Torres P. E. P., Hernández-Álvarez, M., Hernández, E. A. T., & Yoo, S. G. (2019). Stock market data prediction using machine learning techniques. In: Á. Rocha, C. Ferrás, & M. Paredes (Eds.), Information Technology and Systems. *Proceedings of ICITS 2019* (pp. 539–547). Springer. https://doi.org/10.1007/978-3-030-11890-7_52.

Wang, Y., & Guo, Y. (2020). Forecasting method of stock market volatility in time series data based on mixed model of ARIMA and XGBoost. *China Communications*, *17*(3), 205–221. https://doi.org/10.23919/JCC.2020.03.017.

Yuan, X., Yuan, J., Jiang, T., & Ain, Q. U. (2020). integrated long-term stock selection models based on feature selection and machine learning algorithms for China Stock Market. *IEEE Access*, *8*. https://doi.org/10.1109/ACCESS.2020.2969293.