Sentiment Analysis applied to IBOVESPA prediction

Yngwi Guimarães Vieira Souza[1] [0000-0002-2200-7153], Luís Tarrataca[1] [0000-0001-9359-5143], Douglas O. Cardoso[2] [0000-0002-1932-334X], and Laura Silva de Assis[1] [0000-0003-3081-9722]

[1] Celso Suckow da Fonseca Federal Center of Technological Education, Rio de Janeiro, RJ, Brazil
yngwi.guimaraes@aluno.cefet-rj.br, {luis.tarrataca, laura.assis}@cefet-rj.br
[2] Smart Cities Research Center, Polytechnic Institute of Tomar, Tomar, Portugal
douglas.cardoso@ipt.pt

Abstract. Social media is increasingly being used as a source of news, a trend which has resulted in large amounts of data. This work presents an evaluation strategy for assessing the impact that social media has on the Bovespa Index (IBovespa), the benchmark index of the Brazilian stock market. A total of 105000 tweets were collected from the twitter profile of "G1 Economia", one of the main Brazilian finance portals. This data was processed using sentiment analysis methods which were then incorporated into the development of an artificial neural network whose objective was to predict the IBovespa. A hyperparameter optimization study is also presented. The experimental results show that of the 1279 topologies studied, 82.4% exhibited better performance when using sentiment analysis in conjunction with historical data, against the baseline of using only the latter. Curiously, even though the average performance was higher, the absolute best result was obtained without the use of NLP techniques. In the context of the method developed and the data used, it appears that approaches using sentiment analysis alongside historical records may be more effective than using only one or the other.

Keywords: Artificial Neural Networks · Sentiment Analysis · Stock Market · VADER

1 Introduction

Artificial Neural Networks (ANN) have a wide range of applications in economy, business, industry and science. One such potential application is the use of ANN to predict market behavior [12]. Due to the complex nature of the stock market it is difficult to model an ANN capable of making accurate predictions. This difficulty may be due to the absence of features that adequately describe stock market evolution but can also be attributed to the non-linear essence of the market [31]. ANNs are trained on data sets correlating a set of input data to the corresponding outputs. Therefore, the quality of the prediction is dependent on that of the training set.

Typically, news and opinions that directly relate with the economic scenario of a corporation or country have an impact on the psychological attitude of investors towards a market. Market sentiment is therefore influenced by how investors process such news. Fortunately, the proliferation of economic online news channels in social networks has resulted in a large and rich dataset that is constantly being updated [20]. Such a dataset is ideally suited for training ANNs. This work assesses how market sentiment can be incorporated into ANNs through sentiment analysis (SA) using this data.

In order to do so, the following methodology is used: an ANN model is built using standard features that correlate strongly with stock markets, *e.g.,* price of gold, oil and exchange rates.This result is then expanded by using a SA method capable of mapping the sentiment of a text as either positive or negative. This method can then be fed with economic texts collected via online news channels to try to characterize market sentiment at a specific point in time. This metric can then be used with the original set of economic features to predict the IBovespa. The ANN model chosen is the Long Short-Term Memory (LSTM) that has a recurrent neural network architecture and can be applied to time series forecasting. The method chosen for natural language processing (NLP) was VADER (**V**alence **A**ware **D**ictionary and s**E**ntiment **R**easoner) [10], a tool capable of performing SA through well defined rules using a dictionary of words.

The main contributions of this work are the following: *(i)* adaptation of VADER to the Portuguese language; *(ii)* ANN development using different topologies for assessing their respective prediction accuracy; *(iii)* development of a SA strategy for economic news; and *(iv)* incorporation of the SA feature in the development of ANNs and performance evaluation. The remaining sections of this work are organized as follows: Section 2 presents the basic concepts of ANN and SA; Section 3 describes the related works; Sections 4 and 5 detail the VADER adjustment process as well as the dataset creation process; Section 6 details the results obtained; Section 7 presents the conclusions of this work.

2 Contextualization

This section provides a brief description of the concepts and tools employed in the development of this research. It aims at laying the ground for a proper understanding of the proposed methodology and results obtained. In this regard, Section 2.1 concerns Neural Networks (NN) while Section 2.2 covers SA.

2.1 Artificial Neural Networks

ANNs are biologically inspired computational systems that mimic the neural networks that can be found in animal brains [4]. They are the result of a cross-disciplinary approach involving neuroscience, mathematics and artificial intelligence [27]. ANNs represent an intelligent model that learns by example during a training stage. This knowledge can then be used for making future predictions or classifications [1]. The networks employed loosely mimic the neurons found in brains. These neurons can process the information provided and transmit their results to other neurons. Typically, the neurons and the respective interconnecting edges have weights that are adjusted during training. When a neuron is processing multiple inputs it may be possible that the output signal will only be produced if the aggregate signal crosses a threshold [18].

The LSTM is a Recurrent Neural Network (RNN) architecture that incorporates order dependence in its specification and is therefore used when predicting sequences. The ability of the model to remember values in arbitrary intervals is used to classify, process and predict time series with time intervals of unknown duration [8]. LSTMs were developed from the analysis of error flow in existing RNNs. This examination concluded that the backward propagation of errors decays exponentially [7]. An LSTM layer consists of a set of memory blocks connected in a recurrent manner. Each one contains one or more memory cells recurrently connected and three multiplicative units that provide continuous equivalents to the operations of recording (input gate), reading (output gate) and reset (forget gate) for the cells. The interactions between the neural network and the cells are performed exclusively through the cell gates [7]. LSTM networks are constituted by the following elements:

- **Forget Gate**: Removes information in the state of the cell that is no longer useful. Two inputs: x_t (input x in instant t) and h_{t-1} (output of the previous cell) are provided to the gate and multiplied by weight matrices. This is followed by the addition of a bias. The result is fed to an activation function that outputs a binary result. If the function evaluates to 0 then the information is forgotten, otherwise, the information is kept for future use;
- **Input Gate**: Adds useful information to the state of the cell. First, the information is provided to a sigmoid function responsible for filtering the values (in a similar manner to the forget gate) using the inputs h_{t-1} and x_t. In addition, a vector containing all possible values of h_{t-1} and x_t is transformed using the tanh (tangent) function. The outputs of the previous two stages are multiplied to obtain the useful information.
- **Output Gate**: Responsible for the task of extracting useful information from the current state cell. First, a vector is constructed using the tanh function in the cell. This is then fed to the sigmoid function that filters the values to be remembered by using the inputs h_{t-1} and x_t. The vector values and the regulated ones are multiplied and represent the output which can then be provided as input to another cell.

2.2 Sentiment Analysis

SA is used to refer to the automatic processing of opinions, sentiments and subjectivity in texts [16] and is an active research field in the area of NLP. The extraction of sentiments from a text can be performed in different levels, namely:

- **Document level analysis:** Assuming that each document expresses the opinions about a single entity this technique performs the evaluation of the overall sentiment in a specific document;

- **Sentence level analysis:** Determines the overall opinion polarity (positive or negative) expressed within each sentence;
- **Aspect level analysis:** Performs a deeper appreciation than the previous levels by contemplating polarity (negative or positive) and subjected to the perspective of an opinion.

SA can be performed through dictionary-based approaches and machine learning (ML) methods. The former is usually more practical and versatile since it does not require a training set whilst the latter tends to present better results. In this work, we opted to employ the dictionary-based tool VADER. Dictionary-based approaches explore a sentiment lexical that is constructed from a set of documents. The objective of this dictionary is to index the largest number of words that carry some opinion. In contrast, ML approaches require that text is properly classified with a sentiment so that these methods can learn to classify future instances as either postive, negative or neutral [35]. The main resources employed are: *(i)* words, *(ii)* bigrams; *(iii)* trigrams; *(iv)* grammatical class; and *(v)* polarity [30]. Some possible approaches that use ML techniques to perform classification are Support Vector Machine (SVM) and Naive Bayes (NB) classifiers [9]. More recently, deep learning methods have also been developed since they are able to outperform NB techniques [25] but require significantly larger datasets and computational resources [32].

3 Related Work

Stock market forecast is a highly active research field due to the potential financial gains to be had. Not surprisingly, the question of whether or not it is possible to accurately predict market trends using ML attracts significant attention [19]. The Efficient-market hypothesis (EMH) argues that stock prices reflect all available information at any given point in time. As a result of the EMH it is difficult to predict asset returns since any potential public information that would result in a profit would be quickly factored in the price of an asset. However, some works state that it is possible to predict market prices to some level [3].

In [24] a method is presented for performing SA using word dictionaries to process *tweets* applied to the financial market. The authors use a probabilistic analysis to assess how the words in a dictionary represent positive or negative sentiments. The work presents an investment simulation using the develop model that is able to reach 70% accuracy and that was able to achieve a 3.53% profit over the initial investment during 25 simulation days. A study of the relationship between *Apple* stock price and texts obtained from *Twitter* was presented in [6]. The results show a correlation coefficient of 0.7955 between the score of the *tweets* and stock market price variation. A method for SA obtained by combining SentiWordNet and the NB algorithm was described in [13]. The authors collected the texts of *Twitter* users which were then processed and used by SVM to predict the behavior of a stock market.

The work described in [3] applies SA to *Twitter* users by making use of the OpinionFinder software. This application performs NLP through a word dictionary and is capable of determining the polarity and intensity of each word within a sentence. The work made use of ANNs and reported an 87.6% accuracy when predicting the Dow Jones index. In [34] an approach was developed for integrating SA of social media alongside NB. The authors were able to achieve a decrease of 16% of average error when compared against the results obtained without using SA.

In [21] a ML method using SVMs was presented for performing SA as well as stock market prediction. The results obtained were able to deliver a 9.83% accuracy increase when compared against models that did not employ SA during the training stages. The work described in [28] uses ANN and SVMs to improve the accuracy of stock market price prediction models by making use of SA of social media. The authors applied the developed techniques to predict the Dow Jones and obtained ambiguous results. Namely, the best performing method employed SVM alongside SA to obtain a prediction accuracy of 64.10%. However, this result does not improve significantly on the one the authors obtained when using an SVM without SA, respectively, 62.03%.

In [23] a NLP processing method using ANNs was utilized that was capable of classifying a sentence in three different ways: *i)* positive; *ii)* negative; or *iii)* neutral. The authors employed a dataset consisting of Brazilian online news sources, namely: G1, Estadão e Folha de São Paulo. The authors were able to correctly define the sentiment of a sentence with 86.5% accuracy. The work described in [33] aimed at predicting the S&P500 by using a solution based on different ML methods. The approach using a Convolutional Neural Network (CNN) presented the best result for SA whilst the approach using an ANN was able to achieve better stock market patterns.

A Multi Layer Perceptron alongside SA was used in [22] to predict the stock market of Ghana. The tests performed employed data collected from the Ghana stock exchange between 2010 and 2019 alongside texts obtained from Twitter, Google trends, forums and digital newspapers. The approach was able to obtain an average 73.9% accuracy. The work described in [17] attempted to understand the impact on stock markets caused by high magnitude events using SA. The work considered four companies from different countries. The authors collected 11.42 million *tweets* which were then processed using a lexical method and used for training a CNN. The results obtained showed that by employing SA there was a decrease in the average error of the predictions.

In [14] a method was developed based on ML and SA using lexical analysis of four different dictionaries to predict the Hong Kong stock exchange. In [29] the authors measured speech intensity through lexical analysis and employed it to better understand the relationship between social media and stock market. The work concludes that SA can be used to assess the opinion of a group of persons and has an impact on stock prices. A hybrid approach using CNNs and SA was presented in[11] with the objective of predicting the Xangai stock exchange.

The work detailed in [5] employs opinion mining to improve accuracy prediction of stock market models from several countries. The authors made use of a total of 6.48 million *tweets* of the *Xueqiu* social network, which is used by investors in the Chinese financial market. This data was then used for SA [36]. The authors applied a SVM to model and predict stock prices.

4 Sentiment Analysis Adaptation

VADER was chosen given its ability to perceive not only the polarity but also the intensity of sentiments. In addition, it also has higher classification accuracy when compared with other tools [10]. VADER was developed for the English language. Consequently, it was necessary to adapt it to Portuguese given that IBovespa is a Brazilian index and is therefore more susceptible to news in this language. This meant adapting VADER whilst maintaining its structural essence. Overall, the majority of changes consisted in translating English words of the VADER dictionary to Portuguese using *Google Translator* API. Some minor adaptations due to structural differences in grammar were also performed.

After concluding the translation process, it was necessary to assess the performance of the adapted version of VADER. In order to do so we employed 50 economy-related articles that were randomly chosen from a set consisting of 2000 articles obtained from G1 news website [26]. This test considered the text body, title and subtitle of each article. VADER was asked to classify these articles, and its predictions were compared to those of a human agent. The results matched in 78% of the articles. Table 1 brings some examples of these results. The full set of results are publicly available in the project repository [1].

Table 1: Prediction examples of the Portuguese-adapted version of VADER.

Article Title	VADER Score	Human Assessment
Sanitation improves, but half of Brazilians still have no sewage in the country	-0.7062	Negative
New labor law: termination no longer needs union approval	0.9389	Positive
S&P lowers credit score of Brazil	0.9816	Negative
Market breaks 20-week sequence of falling GDP forecast and projects lower inflation	-0.7783	Positive

5 Data Acquisition

When dealing with supervised ML methods it is important to choose a set of features that presents a strong correlation with the output variable [15]. As an initial approach we chose to employ a set of features that, typically, correlate

[1] https://gitlab.com/yngwi/SA_IBOVESPA_Prediction.

strongly with the IBovespa, namely: *(i)* gold; *(ii)* oil; *(iii)* real and dollar exchange rate; *(iv)* Dow Jones; *(v)* Nasdaq; and *(vi)* S&P 500. Each entry of our dataset describes the closing values of these items for a specific day. The data collected encompasses the time period from 4/17/2013 to 2/11/2014.

Figure 1 shows the correlation matrix (generated using the Pearson method [2]) between all the features. As it is possible see, most of the features chosen present high correlation with the IBovespa.

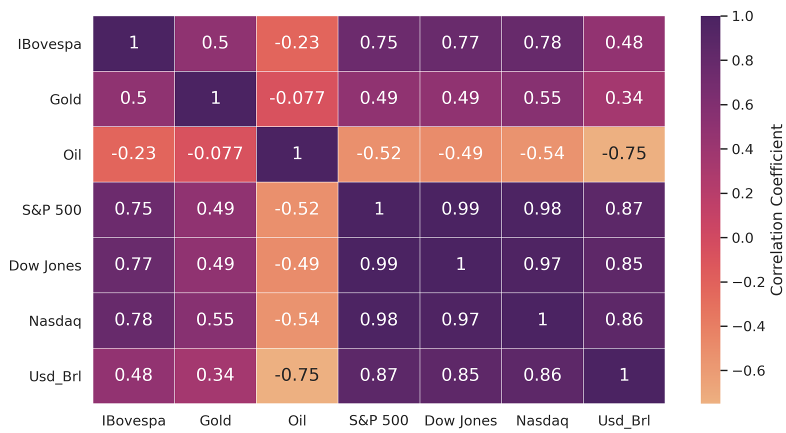

Fig. 1: Correlation matrix for the dataset.

We developed a web scraper for obtaining the news dataset. We opted to use news articles from Twitter given that many news channels have verified accounts in the platform. The web scraper works by performing an advanced query in Twitter and stores in *.csv* format the *tweets* returned. This approach allowed for the creation of a dataset consisting of 105000 *tweets* from the economy specific profile of G1, respectively, "G1 Economia" (in English, G1 Economy).

The adapted VADER method was then employed to perform SA to each *tweet* and classify it as either positive or negative. For each trading day (9:00 to 17:00), the original dataset of features was combined with the SA of the tweets that were in compatible time windows. This enabled the definition of a new feature to represent market sentimental state. The new feature values are obtained through a simple average of the sum of SA values of the day. The historical plot for the SA feature is presented in Figure 2.

It was also tested an alternative to expressing market sentimental state in a single feature. This second approach adds two features that describe the positive sentiment and the negative sentiment for a specific day. These features are composed, respectively, by: *(i)* the sum of the values that are larger than zero;

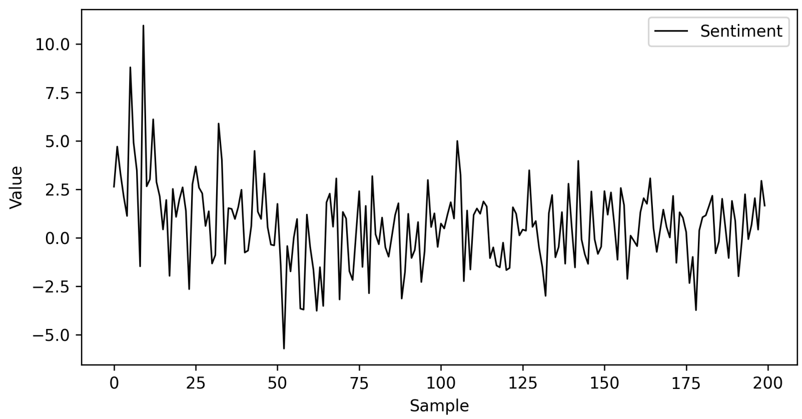

Fig. 2: Sentiment plot.

and *(ii)* the sum of the values that are smaller than zero. The historical plot for the positive and negative sentiment features are presented in Figure 3.

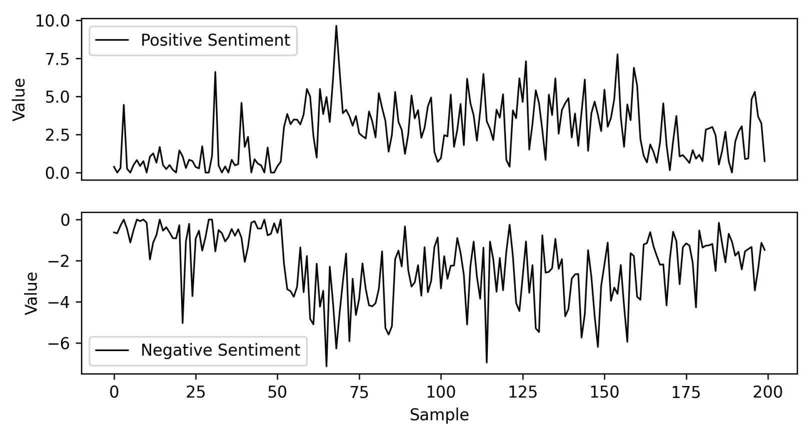

Fig. 3: Positive and negative sentiments plot.

6 Experimental Results

This section presents the results obtained for a diverse number of experiments using the set of hyperparameters described in Table 2. The values employed were the result of an exploratory analysis performed in order to gain a comprehension of the prediction ability of the model.

Table 2: Hyperparameters utilized and their respective values.

Hyperparameter	Description	Value
B_s	Batch size	1024
shuffle	Shuffle input data	No
N_t	Proportion of dataset used for training	80%
N_v	Proportion of dataset used for tests	20%
$L_{\max}$	Maximum number of neurons	256
P_s	Past time series steps used to produce input pattern	1
SP_s	Past time SA steps used to produce input pattern	1
T_s	Number of predictions to be made from an input pattern	1
L_s	Maximum number of hidden layers	5
E_t	Training error	$[0, 1]$
E_v	Validation error	$[0, 1]$
E_s	Number of epochs for training	500

The next sections presented results for the following experiments: *(i)* a topology exploratory analysis; and *(ii)* an assessment of the number of main components employed. These results can be found, respectively, in Section 6.1 and Section 6.2. The results concerning the use of SA are described in Section 6.3.

6.1 Exploratory Analysis of Topologies

These experiments were performed in order to determine which NN topology produced the best results. Accordingly, we opted to test these by varying the maximum number of layers and neurons. In the subsequent text we opted to represent each topology as a sequence of numbers $\{a_1, a_2, \cdots, a_m\}, \forall\ i \in [1, l_s]$ describing the number of neurons in each layer. Tables 3 and 4 present, respectively, the set of values for each parameter tested and the top 10 results.

6.2 Principal Components Analysis

This section aims to assess how the number of principal components impacts prediction accuracy. Accordingly, the 30 topologies that exhibited the smallest E_v's were then tested against different numbers of principal components. Namely, for each topology we trained an ANN with different numbers of principal components, $C_s = \{1, 2, 3, 4, 5, 6\}$. Table 5 presents the best results obtained for the topologies considered. The best performing topology was $(256, 128, 128, 32, 16)$ that used 6 principal components and was able to achieve $E_v = 2.5831\%$.

Table 3: Parameters used in exploratory of the topology analysis.

Parameter	Description	Value
l_s	Number of layers	$\{1, 2, 3, 4, 5\}$
$N_{\max}$	Maximum number of neurons	256
N_v	Step increase in the number of neurons	2^p
p	Power index	$\{1, 2, 3, 4, 5, 6, 7, 8\}$
C_s	Number of principal components	$\{1, 2, 3, 4, 5, 6\}$
E_t	Training Error	$[0, 1]$
E_v	Validation Error	$[0, 1]$
E_s	Number of training epochs	500

Table 4: 10 best results of topology analysis.

Index	Topology	E_v	E_t
1	128, 64, 32, 4, 4	2.986%	2.258%
2	128, 128	3.027%	2.105%
3	256, 64, 16	3.027%	2.298%
4	256, 64, 32, 32, 32	3.059%	2.343%
5	128, 64, 16, 16, 16	3.063%	2.085%
6	128, 128, 128, 32, 8	3.108%	2.007%
7	128, 128, 128, 64, 4	3.109%	2.256%
8	256, 32	3.208%	1.998%
9	256, 256, 128	3.255%	3.221%
10	256, 64, 32, 8, 4	3.272%	2.621%

6.3 Sentiment Analysis Approach

This section presents the experiments that result from adding the SA strategy described in Section 5. The first experiment consisted in considering the original dataset consisting of 2000 articles from G1. Namely, the texts from G1 for each day of the dataset were analyzed using SA and incorporated as a new feature into the original dataset. Table 6 shows the training results for the best topologies for three different strategies, namely: *(i)* without using SA; *(ii)* with a single overall SA feature; and *(iii)* when positive and negative SA features are employed.

The results presented in Table 6 did not lead to conclusive observations, although, this was to be expected given the small size of the dataset employed (2000 articles). As such, we opted to perform a similar analysis but this time using the 105000 *tweets* from “G1 Economy”. Table 7 shows the training results for the best topologies for three different strategies, namely: *(i)* without using SA; *(ii)* with a single overall SA feature; and *(iii)* when positive and negative SA features are employed.

In addition, besides the *tweets* from “G1 Economy”, an additional 25000 *tweets* from the “O Globo Economia” portal were collected. Table 8 shows the corresponding training results for the best topologies for three different strate-

Table 5: 10 best results of C_s analysis.

Index	Topology	E_v	E_t	C_s
1	32, 32, 32, 32, 8	3.9335%	3.9186%	2
2	32, 32, 32, 32, 16	4.0057%	3.9292%	2
3	256, 256, 4, 4, 2	3.2855%	3.2091%	3
4	64, 32, 4, 4, 4	3.4233%	3.4531%	3
5	256, 256, 256, 64, 32	2.8519%	2.856%	4
6	32, 32	3.1184%	3.0928%	4
7	256, 256, 128, 32	2.7873%	2.6994%	5
8	64, 32, 16, 16, 2	2.9027%	2.7293%	5
9	**256, 128, 128, 32, 16**	**2.5831%**	**2.4897%**	**6**
10	256, 256, 64, 2, 2	2.6071%	2.4916%	6

Table 6: Training results when using SA from G1 texts.

Topology	C_s	Mean Average Error %		
		Without SA	Overall SA	+/- SA
256, 128, 128, 32, 16	6	**2.14%**	**2.62%**	**2.32%**
256, 256, 128, 32	5	2.42%	2.69%	2.49%
256, 256, 256, 64, 32	4	2.57%	2.98%	2.64%

Table 7: Training results using SA from “G1 Economy” *tweets*.

Topology	C_s	Mean Average Error %		
		Without SA	Overall SA	+/- SA
256, 128, 128, 32, 16	6	2.48%	**2.41%**	2.50%
256, 256, 64, 2, 2	6	2.56%	2.49%	**2.38%**
64, 32, 16, 16	6	2.94%	**2.35%**	2.81%
256, 256, 256, 64, 32	6	2.45%	2.24%	**2.22%**
64, 32, 16, 16, 2	6	2.87%	2.75%	**2.44%**
256, 256, 4, 4, 4	6	2.49%	**2.40%**	15.75%

gies, namely: *(i)* without using SA; *(ii)* with a single overall SA feature; and *(iii)* when positive and negative SA features are employed.

The overall conclusions of the experiments performed can be found in Table 9 which shows the top three results obtained for each approach. From the data presented it is possible to observe that the approach without SA exhibits the best results in absolute values.

However, for most topologies, the approach that incorporates the positive and negative features outperforms the approach without SA. This behavior can be seen in Table 10 that describes how each approach fared against the others, namely: the first two columns indicate the approaches being compared, with the

Table 8: Training results using SA from "O Globo Economia" *tweets*.

Topology	C_s	Mean Average Error %		
		Without SA	Overall SA	+/- SA
256, 128, 128, 32, 16	6	2.49%	2.71%	**2.21%**
256, 256, 64, 2, 2	6	2.26%	2.27%	**2.04%**
64, 32, 16, 16	6	3.04%	2.69%	**2.47%**
256, 256, 256, 64, 32	6	2.16%	2.21%	**2.03%**
64, 32, 16, 16, 2	6	2.85%	2.83%	**2.64%**
256, 256, 4, 4, 4	6	2.39%	**2.25%**	16.96%

Table 9: Top three results for each approach.

Topology	Approach	E_t	E_v
256, 256, 256	Without SA	1.851 %	1.948 %
256, 256, 256, 128, 128	Without SA	1.791 %	2.004 %
256, 256, 256, 256, 64	Without SA	1.848 %	2.005 %
256, 256, 256, 256, 2	Overall SA	1.904 %	2.014 %
256, 128, 128, 128, 64	Overall SA	1.816 %	2.023 %
256, 128, 128, 64	Overall SA	1.734 %	2.023 %
256, 256, 128, 128	+/- SA	1.624 %	1.952 %
256, 256, 256, 256, 32	+/- SA	1.780 %	2.002 %
256, 256, 256, 128, 64	+/- SA	1.762 %	2.015 %

third column describing the percentage improvement. The last three columns characterize the minimum, maximum and average difference between the approaches being compared. For instance, the first line of Table 10 compares the approach without SA against the one that employs the overall SA strategy. In this specific case, $33,46\%$ of the topologies analyzed without SA showcased better performance than the ones with the overall SA feature. This means that the method employing the overall SA technique exhibited better performance in $100\% - 33,46\% = 66,54\%$ of the topologies considered. The same line states that: *(i)* the best E_v value of the approach without SA was just 0.08% better than the best value obtained from using the overall SA strategy; *(ii)* the minimum E_v value was 1.12% worse; and *(ii)* on average, not using SA resulted in a penalty of 0.86%

Overall, the main conclusion that can be derived from Table 10 is that the approaches using SA outperformed the baseline for the majority of the topologies considered. The best result achieved was when the +/- SA approach was compared against the baseline which resulted in $82,40\%$ of the topologies having better performance.

The previous set of experiment employed $SP_s = 0$ to convert the time series. The next set of experiments (Table 11) describe the results obtained for $SP_s = [0, 1, 2, 3]$. In addition, E_s was altered to 20000 epochs. Overall, the

Table 10: Comparison results between the different approaches.

Approach 1	Approach 2	Topologies Improved	Winning approach	Differences (E_v)		
				Max.	Min.	Avg.
Without SA	Overall SA	33.46%	Overall SA	0.08%	-1.12%	-0.86%
Without SA	+/- SA	17.59%	+/- SA	0.07%	-2.22%	-1.73%
Overall SA	Without SA	66.30%	Overall SA	1.12%	-0.08%	0.86%
Overall SA	+/- SA	24.94%	+/- SA	0.01%	-1.13%	-0.88%
+/- SA	Without SA	82.40%	+/- SA	2.22%	-0.07%	1.73%
+/- SA	Overall SA	74.98%	+/- SA	1.13%	-0.01%	0.88%

results obtained indicate that the approaches without SA showcased the best absolute values.

Table 11: Top three results for each SP_s value, ordered by SP_s and E_v.

Topology	E_v	E_t	SP_s
256, 256, 256, 128, 128,	1.784%	1.124%	0
256, 256, 256, 128,	1.796%	1.316%	0
256, 256, 256, 256, 32,	1.801%	1.162%	0
256, 256, 128, 128, 128,	1.831%	1.291%	1
256, 256, 256, 64,	1.900%	1.211%	1
256, 256, 256, 256, 128,	1.920%	1.363%	1
256, 256, 256, 128, 128,	1.784%	1.124%	2
256, 256, 256, 128,	1.796%	1.316%	2
128, 128, 128, 128,	1.801%	1.162%	2
256, 256, 256, 256, 64,	2.042%	1.264%	3
256, 256, 128, 128, 128,	2.088%	1.287%	3
256, 128, 128, 128, 64,	2.137%	1.635%	3
256, 256, 128, 128, 128,	1.463%	1.079%	Without SA
256, 128, 128, 128,	1.570%	1.135%	Without SA
256, 256, 128, 128, 64,	1.578%	1.225%	Without SA

7 Conclusion

This work assessed the impact of features resulting from sentiment analysis of news text excerpts in Portuguese on Brazilian stock market forecasting. Such features were produced by a version of the VADER tool whose dictionary was translated from English to Portuguese, enabling its use for a language which differs from the one it originally targeted. Among all considered ANN-based predictors, over 82% had a superior performance when such features integrated the input compared to their counterparts which considered only conventional economic indicators. In spite of such fact, it is interesting to notice that the best performance overall ($E_v = 1.948\%$) was obtained by one whose input did not include such features, while the second best did ($E_v = 1.952\%$).

As a future work, it could be interesting to limit the dictionary to economy-related terms: this could avoid misinterpretations resulting from taking into account words which could be considered neutral in this specific context. Another possible continuation would be the employment or even the development of a tool similar to VADER but with Portuguese as its original target language, since such specialized version would probably produce better results than the adapted version used.

Acknowledgements This work has been financially supported by national grant from the FCT (Fundação para a Ciência e a Tecnologia), under the project UIDB/05567/2020.

References

1. Barreto, J.M.: Introdução às redes neurais artificiais. V Escola Regional de Informática. Sociedade Brasileira de Computaçao, Regional Sul, Santa Maria, Florianópolis, Maringá pp. 5–10 (2002) 2
2. Benesty, J., Chen, J., Huang, Y., Cohen, I.: Pearson correlation coefficient. In: Noise reduction in speech processing, pp. 1–4. Springer (2009) 7
3. Bollen, J., Mao, H., Zeng, X.: Twitter mood predicts the stock market. Journal of computational science **2**(1), 1–8 (2011) 4
4. Conway, C.M.: How does the brain learn environmental structure? ten core principles for understanding the neurocognitive mechanisms of statistical learning. Neuroscience & Biobehavioral Reviews **112**, 279–299 (2020) 2
5. Derakhshan, A., Beigy, H.: Sentiment analysis on stock social media for stock price movement prediction. Engineering applications of artificial intelligence **85**, 569–578 (2019) 5
6. Dey, S.K.: Stock market prediction using twitter mood. International Journal of Scientific & Engineering Research **5**(5), 44–47 (2014) 4
7. Graves, A., Schmidhuber, J.: Framewise phoneme classification with bidirectional lstm networks. In: Proceedings. 2005 IEEE International Joint Conference on Neural Networks, 2005. vol. 4, pp. 2047–2052 vol. 4 (2005) 3
8. Graves, A., Liwicki, M., Fernández, S., Bertolami, R., Bunke, H., Schmidhuber, J.: A novel connectionist system for unconstrained handwriting recognition. IEEE Transactions on Pattern Analysis and Machine Intelligence **31**(5), 855–868 (2009) 3
9. Hasan, A., Moin, S., Karim, A., Shamshirband, S.: Machine learning-based sentiment analysis for twitter accounts. Mathematical and Computational Applications **23**(1), 11 (2018) 4
10. Hutto, C., Gilbert, E.: Vader: A parsimonious rule-based model for sentiment analysis of social media text. In: Proceedings of the International AAAI Conference on Web and Social Media. vol. 8 (2014) 2, 6
11. Jing, N., Wu, Z., Wang, H.: A hybrid model integrating deep learning with investor sentiment analysis for stock price prediction. Expert Systems with Applications **178**, 115019 (2021) 5

12. JingTao, Y., Chew Lim, T.: Guidelines for financial prediction with artificial neural networks (2009) 1
13. Li, J., Meesad, P.: Combining sentiment analysis with socialization bias in social networks for stock market trend prediction. International Journal of Computational Intelligence and Applications **15**(01), 1650003 (2016) 4
14. Li, X., Wu, P., Wang, W.: Incorporating stock prices and news sentiments for stock market prediction: A case of hong kong. Information Processing & Management **57**(5), 102212 (2020) 5
15. Lippmann, R.: An introduction to computing with neural nets. IEEE Assp magazine **4**(2), 4–22 (1987) 6
16. Liu, B.: Sentiment analysis and opinion mining. Synthesis lectures on human language technologies **5**(1), 1–167 (2012) 3
17. Maqsood, H., Mehmood, I., Maqsood, M., Yasir, M., Afzal, S., Aadil, F., Selim, M.M., Muhammad, K.: A local and global event sentiment based efficient stock exchange forecasting using deep learning. International Journal of Information Management **50**, 432–451 (2020) 5
18. McCulloch, W.S., Pitts, W.: A logical calculus of the ideas immanent in nervous activity. The bulletin of mathematical biophysics **5**(4), 115–133 (1943) 3
19. Nabipour, M., Nayyeri, P., Jabani, H., Mosavi, A., Salwana, E., et al.: Deep learning for stock market prediction. Entropy **22**(8), 840 (2020) 4
20. Nguyen, T.H., Shirai, K.: Topic modeling based sentiment analysis on social media for stock market prediction. In: Proceedings of the 53rd Annual Meeting of the Association for Computational Linguistics and the 7th International Joint Conference on Natural Language Processing (Volume 1: Long Papers). pp. 1354–1364 (2015) 2
21. Nguyen, T.H., Shirai, K., Velcin, J.: Sentiment analysis on social media for stock movement prediction. Expert Systems with Applications **42**(24), 9603–9611 (2015) 5
22. Nti, I.K., Adekoya, A.F., Weyori, B.A.: Predicting stock market price movement using sentiment analysis: Evidence from ghana. Appl. Comput. Syst. **25**(1), 33–42 (2020) 5
23. de Oliveira Carosia, A.E., Coelho, G.P., da Silva, A.E.A.: Investment strategies applied to the brazilian stock market: A methodology based on sentiment analysis with deep learning. Expert Systems with Applications **184**, 115470 (2021) 5
24. Pagolu, V.S., Reddy, K.N., Panda, G., Majhi, B.: Sentiment analysis of twitter data for predicting stock market movements. In: 2016 international conference on signal processing, communication, power and embedded system (SCOPES). pp. 1345–1350. IEEE (2016) 4
25. Paredes-Valverde, M.A., Colomo-Palacios, R., Salas-Zárate, M.d.P., Valencia-García, R.: Sentiment analysis in spanish for improvement of products and services: A deep learning approach. Scientific Programming **2017** (2017) 4
26. e Participações S.A, G.C.: G1 - o portal de notícias da globo (2000), `g1.globo.com` 6
27. Pehlevan, C., Chklovskii, D.B.: Neuroscience-inspired online unsupervised learning algorithms: Artificial neural networks. IEEE Signal Processing Magazine **36**(6), 88–96 (2019) 2
28. Porshnev, A., Redkin, I., Shevchenko, A.: Machine learning in prediction of stock market indicators based on historical data and data from twitter sentiment analysis. In: 2013 IEEE 13th International Conference on Data Mining Workshops. pp. 440–444. IEEE (2013) 5

29. Reed, M.: A study of social network effects on the stock market. Journal of Behavioral Finance **17**(4), 342–351 (2016) 5
30. Sindhu, C., ChandraKala, S.: A survey on opinion mining and sentiment polarity classification (2013) 4
31. Soni, S.: Applications of anns in stock market prediction: a survey. International Journal of Computer Science & Engineering Technology **2**(3), 71–83 (2011) 1
32. Tang, D., Qin, B., Liu, T.: Document modeling with gated recurrent neural network for sentiment classification. In: Proceedings of the 2015 conference on empirical methods in natural language processing. pp. 1422–1432 (2015) 4
33. Vargas, M.R., De Lima, B.S., Evsukoff, A.G.: Deep learning for stock market prediction from financial news articles. In: 2017 IEEE international conference on computational intelligence and virtual environments for measurement systems and applications (CIVEMSA). pp. 60–65. IEEE (2017) 5
34. Zhang, G., Xu, L., Xue, Y.: Model and forecast stock market behavior integrating investor sentiment analysis and transaction data. Cluster Computing **20**(1), 789–803 (2017) 5
35. Zhang, Q., Wang, B., Wu, L., Huang, X.: Fdu at trec 2007: Opinion retrieval of blog track. In: TREC. pp. 500–274. Citeseer (2007) 4
36. Zhang, X., Shi, J., Wang, D., Fang, B.: Exploiting investors social network for stock prediction in china's market. Journal of computational science **28**, 294–303 (2018) 6