Gated Recurrent Unit Networks and Discrete Wavelet Transforms Applied to Forecasting and Trading in the Stock Market

Victor Biazon[1], **Reinaldo A. C. Bianchi**[1]

[1]Centro Universitário da FEI
Av. Humberto de Alencar Castelo Branco, 3972-B - Assunção
São Bernardo do Campo - SP - Brasil, 09850-901

vbiazon@fei.edu.br, rbianchi@fei.edu.br

Abstract. *Trading in the stock market always comes with the challenge of deciding the best action to take on each time step. The problem is intensified by the theory that it is not possible to predict stock market time series as all information related to the stock price is already contained in it. In this work we propose a novel model called Discrete Wavelet Transform Gated Recurrent Unit Network (DWT-GRU). The model learns from the data to choose between buying, holding and selling, and when to execute them. The proposed model was compared to other recurrent neural networks, with and without wavelets preprocessing, and the buy and hold strategy. The results shown that the DWT-GRU outperformed all the set baselines in the analysed stocks of the Brazilian stock market.*

1. Introduction

The famous Dow's Theory is the foundation of the modern technical analysis and states that every possible aspect that could influence the price of a determined stock is already naturally taken in consideration by the market itself. Although news and climate disasters cannot be predicted, the market adapts rapidly assimilating the effects of such events on the current price [Brown et al. 1998].

The Efficient Market Hypothesis (EMH) says that stock market price time series are almost always unpredictable because every piece of relevant information that can influence the price is already taken in account including past values and volumes. This implies that the price responds immediately to new information and isn't bound to any trend or pattern. The EMH states that any type of prediction or forecasting is faded to have no better performance than random guessing and the stock price will always be the fair one, therefore unpredictable [Nobre and Neves 2019].

In this paper we propose that even if the market do oscillate in a pseudo-randon walk, there are some periods of predictability in some extent which can be exploited to acquire more profit than the Buy and Hold strategy (which consists on buying a certain stock with the confidence that its value will rise in the long run) in bull markets and having a positive performance in bear markets.

This paper is organized as follows: Section 2 presents the base theory behind the proposed model; Section 3 reviews some related works in the discussed areas; Section 4 proposes the methodology behind our research; Section 5 demonstrates the results and comparisons of the model and section 6 concludes our paper.

2. Theoretical Background

2.1. Recurrent Neural Networks

Recurrent Neural Networks (RNNs) [Weigend et al. 1991] are a different class of Artificial Neural Network (ANN) [Haykin et al. 2009] built specifically to analyse and learn from sequential data with, or without, multivariate dependencies. Generally RNNs can also process data of variable length, opposed to the method of processing sequential data on a Feed-forward Neural Network (FFNN) [Tang and Fishwick 1993] inputting current and past values as a fixed size array.

Several variations of RNNs were developed over the years, such as Long Short-Term Memory (LSTM) [Hochreiter and Schmidhuber 1997], Gated Recurrent Unit (GRU) [Chung et al. 2014] and Deep Recurrent Neural Networks (DRNN) [Pascanu et al. 2013]. All the models share the necessity to recall past inputs or results to predict the next state of the network and adapt itself [Goodfellow et al. 2016]. As cited by [Gopalswamy et al. 2017], there are several applications where the use of RNNs can be determinant to predict correctly the outcome of the analysed data, e.g. machine translation, speech recognition, language understanding, image captioning and time-series forecasting.

LSTM and GRU were specifically designed to overcome the difficulty created by the vanishing or exploding gradient, an inherent deep learning applications problem, and propagates its errors through internal states of special units controlled by each cell gates which are activated by the learning algorithm[Hochreiter and Schmidhuber 1997] [Chung et al. 2014].

2.2. Wavelets

Wavelets are small waves with flexible period and forms that can be stretched, phase shifted and dislocated in time to fit to the analysed signal [Nobre and Neves 2019]. The selection of the wavelet form to be used depends on the signal to be studied as said by [Galli et al. 1996].

Wavelets can be traced as an parallel to music tones and their assimilation and superposition with the right placement in time, with coefficients to describe their scale and intensity, can rebuild the transformed signal as it was music decomposed in tones. Generally, wavelets transforms can be classified in two main categories Discrete Wavelet Transforms (DWT) and Continuous Wavelets Transforms (CWT). We focus on the DWT which comprises the information in discrete coefficients obtained by processing the signal through an algorithm similar to the structure of a band-pass filter as illustrated in Fig.1, which image was presented by [Nobre and Neves 2019].

Wavelets were conceived to overcome the inherent difficulties in applying the Fourier transform to non-periodical signals. Although the Fourier transform is capable of expressing the frequency components of the analysed signal in series of sines and cosines, it lacks the capability of determining changes on the frequency spectrum from time to time [Valens 1999].

On the other hand, wavelets transforms considers both time and frequency on its analysis resulting in a better understanding of where and when in the signal each wavelets component manifests its effects. This specific characteristic of the wavelets transform

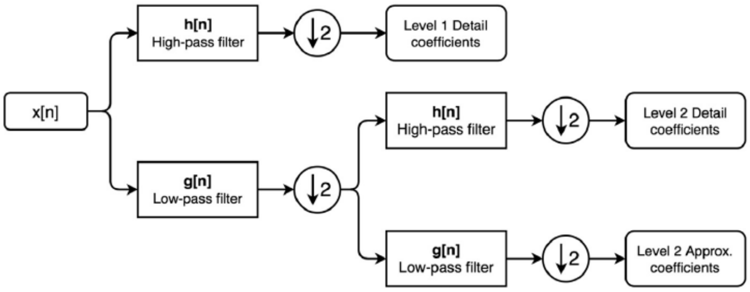

Figura 1. DWT Decomposition of a signal in two levels

makes it possible to analyse signals evolving in time simply by shaping and extracting new wavelets coefficients and positions then appending the values to the existing analysis.

The Empirical Wavelet Transform(EWT) is another way to analyse non-periodical signals and mainly consists on decomposing the signal in several steps using the DWT, the rebuilding only one wavelet detail layer at a time, resulting in several signals that if added together results in the original signal [Altan et al. 2019]. The result of the process is exemplified in Fig 2.

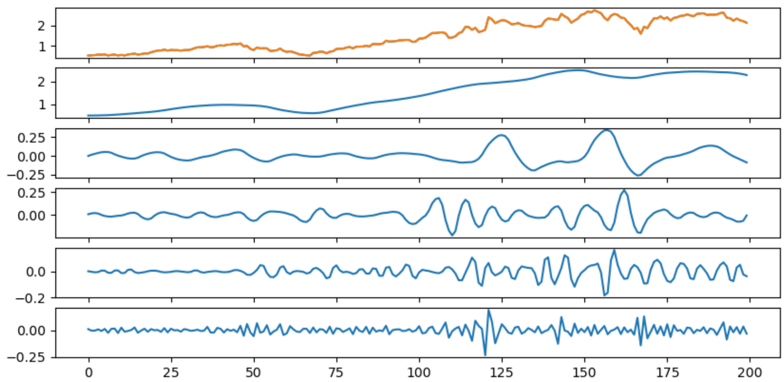

Figura 2. Empirical Wavelet Transform of the Stock PETR4 in 5 components. The original signals in orange, and the resulting Empirical wavelets in blue from the lowest to highest frequencies.

The algorithm developed by [Mallat 1999] decomposes the signal in several levels of high and low pass filters, each level being subsequently decomposed to get more detail about it. The wavelets transform is a powerful tool to analyse signals even from several dimensions of data. It also can be used as signal denoising as exemplified by [Sardy et al. 2001], which consists on not considering the higher frequency coefficients when reversing the wavelet transform to rebuild the signal. Figure 3 exemplifies the signal denoising effect by applying the DWT and taking out the highest frequencies on a given stock.

2.3. Stock Market Technical Analysis

Technical analysis in stock investment theory is an analysis methodology for forecasting the direction of prices through the research on past market data [Zhang et al. 2019]

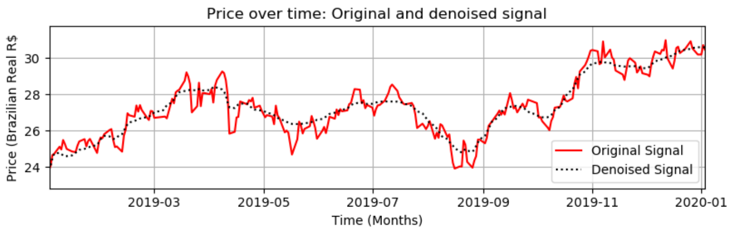

Figura 3. Example of DWT Denoising applied to the Stock PETR4 in 2019 year.

and has evolved from the desire to understand the behavior of the market to the development of several mathematical formulas and algorithms to describe and extract information of the analysed stocks.

Technical Indicators are said mathematical models to interpret the future direction and trends of stocks using the past historical data. These indicators are mostly used by short-term traders in identifying trading strategies and in turn taking timely investment decisions. There are hundreds of these indicators available in literature and they are mostly used in combinations to form a trading strategy. The choice of technical indicators is thus extremely important [Ullah et al. 2019].

In his article [Murphy 1999](cited in [Pimenta et al. 2018]) states that there are several types of technical indicators listed below:

- Trend followers: the indicators of this category identify the main movement direction of the asset prices at a certain period.
- Oscillators: the indicators of this category monitor the price variations of the asset in a certain range in order to identify possible reverse points.
- Band systems: band systems are constituted of three curves drawn around the prices. These curves are drawn from a particular distance of a moving average. The intermediate band is usually a simple moving average, and the intervals between the bands are determined by the price volatility.
- Divergence Identifier: these indicators are based on the principle that the whole trend goes through corrections. The divergences occur when comparing the behavior of the indicator in relation to the price movement of an asset.

3. Related Work

The following section presents some articles on which we have based our methodology and experiments, and are grouped by main theme, them being Recurrent Neural Networks, Wavelets and Stock Market Technical Analysis.

3.1. Recurrent Neural Networks

Since RNNs are capable of learning long and short-term patterns in the data, several papers in the last years tried to predict trends or prices in the stock market applying them. In his study [Alonso-Monsalve et al. 2020] created a hybrid algorithm called CLSTM that combines a Convolutional Neural Network (CNN) to preprocess the stock data consisting on the price and several technical indicators, then feeding the results in a LSTM network to predict the next tick price. The authors compared the results with a pure LSTM network, a Radial Basis Function Neural Network (RBFNN) and a CNN, concluding that the algorithm proposed was more accurate than the aforementioned models.

[Nelson et al. 2017] trained a LSTM model to predict the price movement of the BOVESPA (Brazil Stock Market) and accruing 55.9% of accuracy on predicting if the price is going up or down on the near future. His methodology used not only the historical price of the stock to train the LSTM, but several momentum technical indicators too.

In their paper [Huynh et al. 2017] developed a new Bidirectional Gated Recurrent Unit (BGRU), a model that relies on both online financial news and historical stock prices data to predict the stock movements in the future. Experimental results show that their model achieves accuracy of nearly 60% in S&P 500 index prediction, whereas the individual stock prediction is over 65%. Their results were compared to a simple Gated Recurrent Unit network (GRU) and a LSTM network, outperforming both of them.

[Shen et al. 2018] applied GRU networks allied with Support Vector Machine(SVM) creating a novel model called GRU-SVM. The research compared a simple GRU network, the GRU-SVM, Deep Neural Network (DNN) and the pure SVM. Every model was trained to predict if the stock trend was going up or down. The ones that applied GRU networks had on all cases higher accuracy in predicting the trend direction and surpassed 50% which validates its use in a dollar neutral approach.

This researches all set the precedent that the RNNs, in particular LSTM and GRU, are being successfully applied in the trend direction prediction but the best results are generally associations of RNNs allied to some other preprocessing method applied to the data.

3.2. Wavelets

In their paper [Nobre and Neves 2019] used the DWT combined with XGBoost (which is the abbreviation of Extreme Gradient Boosting, that uses a great number of weak predictors to make a strong one) and Principal Component Analysis (PCA), to analyse and predict the tendencies of the stock market to buy or sell effectively, based on the price of the stock and several technical indicators. Another researcher that used wavelets combined them with RNNs and Artificial Bee Colony (ABC) to predict the price of the stock and maximizing profits. The DWT was used in noise reduction in every feature fed to the XGBoost model [Hsieh et al. 2011].

In the cryptocurrency exchange rates prediction field, [Altan et al. 2019] applied the Empirical Wavelets Transform (EWT) firstly proposed by [Gilles 2013] which consists on extracting different intrinsic modes of a time-series by building adaptive wavelets. The algorithm allied EWT to LSTM to predict the value of the criptocurrency, and Cuckoo Search (CS) to optimize the training of the LSTM. This specific work show that the wavelet decomposition in empirical wavelets resulted in a expressive reduction in the prediction losses.

In the work of [Stocchi and Marchesi 2018] a fast DWT is presented which consists on using compact supported Debeauchi wavelets to transform the signal in several coefficients and then using then inverse of the transform in an integrated group of estimating machines to predict cryptocurrency exchange prices.

The aforementioned works show that wavelets are a effective method to denoise and decompose data, and its application in financial data improved drastically the accuracy and reduced losses of the trained models.

3.3. Stock Market Technical Analysis

In their paper [Nobre and Neves 2019] generated 26 technical indicators with the Python library TA-lib, and added it to 5 raw price indicators (Open, High, Low, Close(OHLC) and volume) then normalized the data and fed it to a PCA to extract the most relevant ones.

[Alonso-Monsalve et al. 2020] applied a total of 18 indicators, all of them "momentum"indicators. They affirmed that, even though the results where good enough to validate the study, if more indicators were employed the results would be better. They justify that their study employed only "momentum"indicators because their CLSTM algorithm had a high computational cost and it would increase otherwise.

In the work of [Nelson et al. 2017] their LSTM input layer has a dimensionality of 180 features, that consists of the set of technical indicators (all of those generated by TA-Lib) plus the price raw data (OHLC and volume) and extracted an output using the tanh function connected to the network's output layer through 20 connections.

[Hsieh et al. 2011] executed the network training with 14 technical indicators (of momentum, bands and means) and raw price data preprocessed by Haar wavelets then fed them to a LSTM network to learn the patterns. His algorithm has its weights adjusted by the Artificial Bee Colony (ABC) optimization algorithm.

[Pimenta et al. 2018] employed 11 technical indicators subdivided in oscillators, trend followers, band systems and divergence identifiers. [Ullah et al. 2019] going in the same direction used Simple Moving Average (SMA), MACD, Stochastic Oscillator, Triple Exponential Average (TRIX) and Average Directional Index (ADX), then fed the values to a DNN to predict stock prices in the Pakistan Stock Exchange (PSE).

These papers show that the choice of the financial technical indicators can greatly influence the result of the study. As mentioned by some authors, more technical indicators usually result in better accuracy in the prediction models. We can extract from so many works applying this mathematical models that technical indicators can serve as a reliable source of information for models to learn complex financial trends.

4. Methodology

4.1. System Architecture

The proposed model consists on combining the denoising and decomposition capacity of the DWT and applying the pre-processed data to be trained by the recurrent neural network based primarily on the Gated Recurrent Unit Neural Network (GRUNN).This novel model called DWT-GRU network supposedly can out perform simple associations of only LSTM or GRU networks that don't previously pre-process the data with the DWT, as well as, the buy and hold strategy that is frequently used as base model for bench-marking the stock price prediction algorithms.

In our system, the DWT poses as a method to improve the RNNs learning efficiency. With the noise filtered by the DWT the model is not influenced by high frequency perturbations, allowing the learning of clearer tendencies in the data.

The proposed system architecture is composed of several processes consisting mostly on acquiring and preprocessing the data, training the neural network, computing results and evaluating the model. Figure 4 illustrates the actions executed by the system and the details of each individual process will be described in the next subsections.

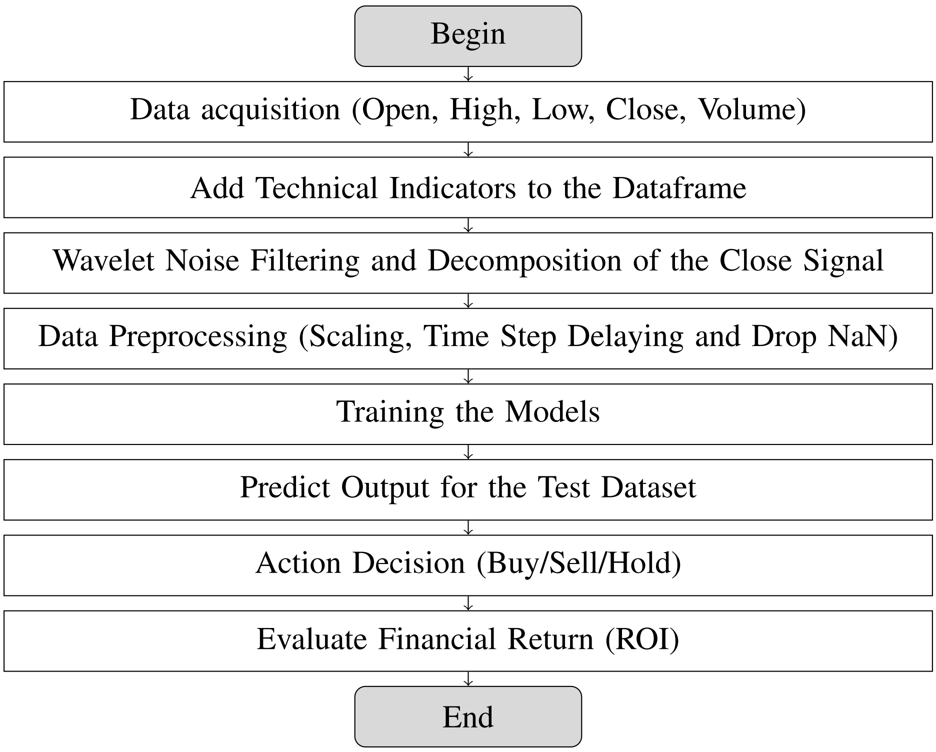

Figura 4. System process fluxogram

4.2. Data acquisition

The data used in this study was acquired from the Yahoo Finance databate for each stock individually in daily raw information of the stock price (Low, High, Open, Close, Volume). The database exports a .csv file (comma separated values) and the data is imported in the system as a Python's Pandas dataframe.

4.3. Data Preprocessing

The data preprocessing is subdivided in several steps described below:

- Adding Technical indicators to the stock data-frame using the Python library TA-lib, totalling in this step 74 variables;
- Applying the DWT of the Close price signal and filtering it rebuilding the signal without the 2% higher frequencies, then applying the first and second derivative operation to the rebuilt filtered signal, and finally adding the results to the data-frame, accumulating 77 variables;
- Decomposing the Close signal into 6 empirical wavelets and adding the output to the data-frame, resulting in 83 variables;
- Scaling the data-frame applying the StandardScaler function of the Python Sklearn package;

- Applying for each dimension of the data frame (the 83 data columns) the time delay shift, which first copies the columns a determined number of times, then shift it to make each observation be composed of n number of observation like x(t-n)...x(t-1),x(t). The number of days used in this approach were 5 days in total;
- The desired output to train the models is selected here, taking in this case the filtrated signal of the Close signal on two days after the time steps delay, and the aims of the training will be then to predict two days in the future based on the known data available;
- The final step is to split the data in three parts, following the 80/20% distribution between training data and validation data, then the validation data is split with a distribution of 75/25% between validation data and test data.

4.4. Proposed Model

The GRUNN was chosen based on its on par capacity with LSTM networks but being less demanding in terms of computational effort, thus being trained faster and converging easier into satisfactory results. Fig. 5 illustrates the layers used in building the neural network model responsible for the prediction of the system. The first layer consists on a GRU layer with its dimension based on multiplying the input shape by 2, and counting on a 0.4 dropout and 0.3 L2 regularizer for over-fitting prevention and a single fully-connected layer with a single neuron to predict the stock price as the network output.

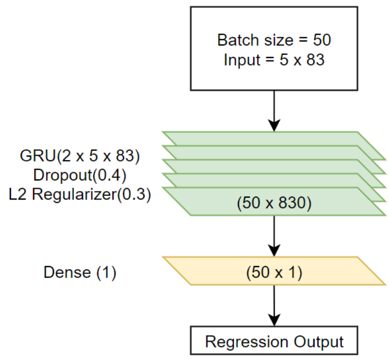

Figura 5. Proposed Neural Network Model

4.5. Action Decision

After training and validating, the next step consists on using the trained model to predict the output of the test dataset. The predicted signal is then treated in a specific function to decide whether it will indicate buy, sell or hold actions.

The decision is based on detecting zero-crossings on the first derivative of the predicted signal and evaluating if its going up or down. E.g if the value crosses the zero line and the next value is positive, there is a tendency of the price to go up, therefore the action to be taken is ”BUY”. If the derivative of the output crosses the zero line and the next value is negative, then the action to be taken is ”SELL”. Every other case that doesn’t fall in these two situations is by elimination ”HOLD”.

4.6. The BOVESPA Dataset

The chosen data for the study is based on the Brazilian Stock Market (BOVESPA) and is mainly focused on the ETF BOVA11 and the blue chips ITUB4, ABEV3, PETR4 and VALE3, that are well known for the high market value and count on a long history of consistent and reliable growth in the long run. The data sets were taken from Yahoo Finance a well known and established stock data source. All the stocks financial data dates back to year 2000, except for BOVA11 which dates back to 2008's January.

5. Experiments

The experiments were conducted testing the trained models in the 2001-2016 historical data on the 2019 historical data. The DWT-GRU model was compared to several baselines including the DWT-LSTM, and pure GRU and LSTM networks. The ROI of each model, as well as the buy-and-hold strategy, will also be discussed. The model was trained to predict two days ahead as this increases the capability of the model to take actions on the best moments possible, without losing accuracy.

Figure 6 illustrates the behavior of the algorithm in the 2019 historical data of stock PETR4, and shows a much greater profit using the DWT-GRU algorithm than just buying and holding the stock for the period.

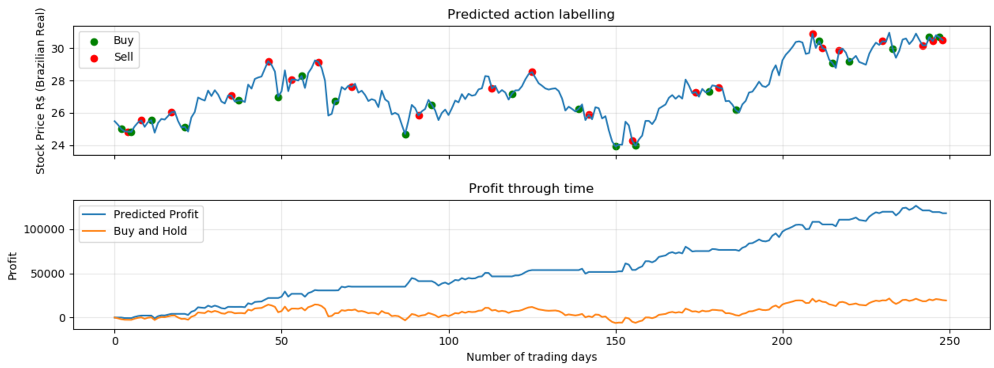

Figura 6. Algorithm actions and profit of PETR4 stock

5.1. Analysis

The evaluation of the Return Over Investment (ROI) is based on starting with a limited quantity of "money"available for the algorithm to use in its actions. Every time the algorithm finds a "Buy"action it evaluates how many shares it can buy with the available balance and emulates a transaction subtracting from the balance the value of how many shares, multiples of 100, it "bought"then transfers that value to another variable called "invested value". If the algorithm finds a "Sell"action, the algorithm checks how many shares does it have on custody and the current price of the stock to assert how much is it worth and terminates its position "selling"the shares by transferring the current value to the variable "balance". When the period of evaluation defined in the testing data-frame ends, the algorithm calculates how much profit the strategy returned following Equation (1):

$$ROI = \frac{(B_F + I_V) - B_0}{B_0} x100\% \tag{1}$$

where: B_0 = Initial Balance
B_F = Final Balance
I_V = Invested Value

The novel DWT-GRU model had quite a good performance compared to the baseline algorithms and the buy and hold strategy as illustrated in Fig. 7. As we can observe the total profit return of the proposed model was higher in every stock outperforming every baseline models we compared to. Brokerage was not taken in account in the results as there are several investment platforms that do not charge for operations nowadays.

Especially for BOVA11 we compared the results of an independent investment fund called Alaska Black and the proposed algorithm. The proposed system also surpassed the ROI of this fund[1] in the analysed period. We can also note that the DWT-GRU was consistently better than the DWT-LSTM, and we consider that this result came from the better training the DWT-GRU attained based on its simpler architecture, thus being easier to train in such a complex environment.

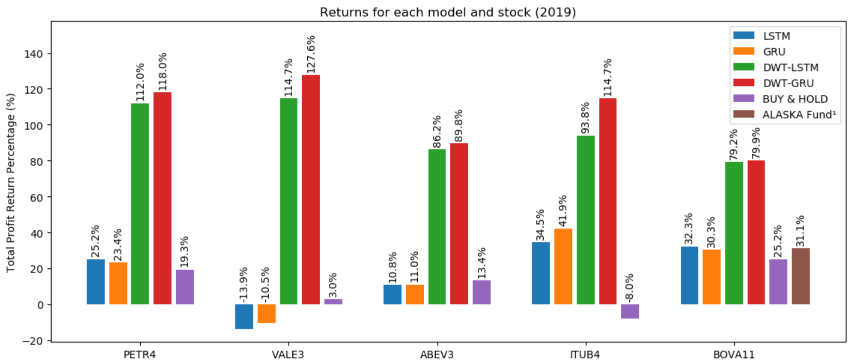

Figura 7. Profit return over investments for each stock

The following Table 1 groups the model metrics for each stock, crossing the expected label/action the algorithm should do in each trading day and the predicted or actual action taken. As we can see the accuracy was satisfactory for a forecasting model using so many input variables and with this kind of complexity.

Tabela 1. DWT-GRU Model Metrics

Stock	Accuracy	Precision	Recall	F1
PETR4	0.852	0.875	0.852	0.861
VALE3	0.788	0.733	0.788	0.806
ABEV3	0.740	0.787	0.740	0.759
ITUB4	0.756	0.801	0.756	0.773
BOVA11	0.700	0.790	0.700	0.736

[1]https://www.alaska-asset.com.br/pdf/Fundos/BLACK_INSTITUCIONAL.pdf [Accessed: 21/04/2020]

6. Conclusion and future work

In this work we proposed a comparison between RNN models allied, and not, with discrete wavelet transforms to preprocess the stock data helping the models better understand the trends in each analysed stocks. The developed system was compared to pure GRU and LSTM models, to the buy and hold strategy and particularly in the case of the ETF BOVA11 to a independent investment fund called Alaska Black.

We can observe that in general the proposed model had a better Return Over Investments than the compared algorithms and the buy-and-hold strategy. The implementation of the DWT-GRU shows promising results over improving the capacity to predict and take actions in the stock market using the DWT. The statistical metrics of the model were satisfactory either. The main contribution of this paper is the appplication of the wavelet preprocessing which drastically improved result in both LSTM and GRU networks. This validates the application of the model in other experiments with stocks, commodities and cryptocurrencies, pursuing the continuity of this research.

For future work we propose to compare the results to other baselines as simpler forecasting methods like ARIMA or complex models like XGBoost. Even thought the results of the proposed model are already better than the compared baselines we intend to improve the training and the parameter tuning with several techniques e.g. evolutionary algorithms and genetic algorithms.

Referências

Alonso-Monsalve, S., Suárez-Cetrulo, A. L., Cervantes, A., and Quintana, D. (2020). Convolution on neural networks for high-frequency trend prediction of cryptocurrency exchange rates using technical indicators. *Expert Systems with Applications*, page 113250.

Altan, A., Karasu, S., and Bekiros, S. (2019). Digital currency forecasting with chaotic meta-heuristic bio-inspired signal processing techniques. *Chaos, Solitons & Fractals*, 126:325–336.

Brown, S. J., Goetzmann, W. N., and Kumar, A. (1998). The dow theory: William peter hamilton's track record reconsidered. *The Journal of finance*, 53(4):1311–1333.

Chung, J., Gulcehre, C., Cho, K., and Bengio, Y. (2014). Empirical evaluation of gated recurrent neural networks on sequence modeling. *arXiv preprint arXiv:1412.3555*.

Galli, A., Heydt, G., and Ribeiro, P. (1996). Exploring the power of wavelet analysis. *IEEE Computer Applications in Power*, 9(4):37–41.

Gilles, J. (2013). Empirical wavelet transform. *IEEE transactions on signal processing*, 61(16):3999–4010.

Goodfellow, I., Bengio, Y., and Courville, A. (2016). *Deep learning*. MIT press.

Gopalswamy, S., Tighe, P. J., and Rashidi, P. (2017). Deep recurrent neural networks for predicting intraoperative and postoperative outcomes and trends. In *2017 IEEE EMBS International Conference on Biomedical & Health Informatics (BHI)*, pages 361–364. IEEE.

Haykin, S. S. et al. (2009). Neural networks and learning machines/simon haykin.

Hochreiter, S. and Schmidhuber, J. (1997). Long short-term memory. *Neural computation*, 9(8):1735–1780.

Hsieh, T.-J., Hsiao, H.-F., and Yeh, W.-C. (2011). Forecasting stock markets using wavelet transforms and recurrent neural networks: An integrated system based on artificial bee colony algorithm. *Applied soft computing*, 11(2):2510–2525.

Huynh, H. D., Dang, L. M., and Duong, D. (2017). A new model for stock price movements prediction using deep neural network. In *Proceedings of the Eighth International Symposium on Information and Communication Technology*, pages 57–62.

Mallat, S. (1999). *A wavelet tour of signal processing*. Elsevier.

Murphy, J. J. (1999). *Technical analysis of the financial markets: A comprehensive guide to trading methods and applications*. Penguin.

Nelson, D. M., Pereira, A. C., and de Oliveira, R. A. (2017). Stock market's price movement prediction with lstm neural networks. In *2017 International joint conference on neural networks (IJCNN)*, pages 1419–1426. IEEE.

Nobre, J. and Neves, R. F. (2019). Combining principal component analysis, discrete wavelet transform and xgboost to trade in the financial markets. *Expert Systems with Applications*, 125:181–194.

Pascanu, R., Gulcehre, C., Cho, K., and Bengio, Y. (2013). How to construct deep recurrent neural networks. *arXiv preprint arXiv:1312.6026*.

Pimenta, A., Nametala, C. A., Guimarães, F. G., and Carrano, E. G. (2018). An automated investing method for stock market based on multiobjective genetic programming. *Computational Economics*, 52(1):125–144.

Sardy, S., Tseng, P., and Bruce, A. (2001). Robust wavelet denoising. *IEEE Transactions on Signal Processing*, 49(6):1146–1152.

Shen, G., Tan, Q., Zhang, H., Zeng, P., and Xu, J. (2018). Deep learning with gated recurrent unit networks for financial sequence predictions. *Procedia computer science*, 131:895–903.

Stocchi, M. and Marchesi, M. (2018). Fast wavelet transform assisted predictors of streaming time series. *Digital Signal Processing*, 77:5–12.

Tang, Z. and Fishwick, P. A. (1993). Feedforward neural nets as models for time series forecasting. *ORSA journal on computing*, 5(4):374–385.

Ullah, S., Javed, N., Hanif, A., and Abdullah, A. (2019). Stock price forecast using recurrent neural network. *Urdu News Headline, Text Classification by Using Different Machine Learning Algorithms*, page 47.

Valens, C. (1999). A really friendly guide to wavelets. *ed. Clemens Valens*.

Weigend, A. S., Rumelhart, D. E., and Huberman, B. A. (1991). Back-propagation, weight-elimination and time series prediction. In *Connectionist models*, pages 105–116. Elsevier.

Zhang, K., Zhong, G., Dong, J., Wang, S., and Wang, Y. (2019). Stock market prediction based on generative adversarial network. *Procedia computer science*, 147:400–406.