engineering proceedings

Proceeding Paper

Foreign Exchange Forecasting Models: ARIMA and LSTM Comparison †

Fernando García [1], Francisco Guijarro [1], Javier Oliver [1,*] and Rima Tamošiūnienė [2]

1 Department of Economics and Social Sciencies, Universitat Politècnica de València, 46022 Valencia, Spain; fergarga@esp.upv.es (F.G.); fraguima@esp.upv.es (F.G.)
2 Department of Financial Engineering, Vilnius Gediminas Technical University, 10223 Vilnius, Lithuania; rima.tamosiuniene@vilniustech.lt
* Correspondence: jaolmun@ade.upv.es
† Presented at the 9th International Conference on Time Series and Forecasting, Gran Canaria, Spain, 12–14 July 2023.

Abstract: The prediction of currency prices is important for investors with foreign currency assets, both for speculation and for hedging the exchange rate risk. Classical time series models such as ARIMA models were relevant until the advent of neural networks. In particular, recurrent neural networks such as long short-term memory (LSTM) are show to be a good alternative model for the prediction of short-term stock prices. In this paper, we present a comparison between the ARIMA model and LSTM neural network. A hybrid model that combines the two models is also presented. In addition, the effectiveness of this model on Bitcoin's future contract is analysed.

Keywords: ARIMA; LSTM; foreign exchange prediction

check for updates

Citation: García, F.; Guijarro, F.; Oliver, J.; Tamošiūnienė, R. Foreign Exchange Forecasting Models: ARIMA and LSTM Comparison. *Eng. Proc.* **2023**, *39*, 81. https://doi.org/10.3390/engproc2023039081

Academic Editors: Ignacio Rojas, Hector Pomares, Luis Javier Herrera, Fernando Rojas and Olga Valenzuela

Published: 12 July 2023

1. Introduction

The foreign exchange market moves, on average, more than USD 6 million per day. It is a fundamental market for international transactions of services and goods. Hence, it is important to be able to have efficient models for predicting currency prices, as well as being able to determine their evolution. With this information, different economic agents and companies can establish their foreign currency risk levels and hedging strategies. On the other hand, it is also important for investors and speculators to know and understand the evolution of prices, and it is therefore necessary to improve predictions by reducing prediction errors as much as possible. There is a large number of currencies whose price is traded against the US dollar (USD). Those that are the most traded are considered majors. Others are considered exotic, as they are traded to a lesser extent, even though they are priced against the USD.

Since currencies can be considered time series, it is possible to apply different time series forecasting models such as the classical ARIMA model. There is a lot of applied literature on forecasting using these models for different currency pairs. For example, the author of [1] proposes a model for the prediction of the USD/EUR exchange rate taking into account the purchasing power parity theory. This theory is fundamentally based on the non-existence of arbitrage prices. It takes into account the price level differential between two countries. This model has also been applied to exotic currencies. For example, the author of [2] uses the Box–Jenkins methodology to apply an AR(1) model to predict the NGN to USD exchange rate for the period 1982–2011. On the other hand, ARIMA models are static once estimated [3]. It is necessary to create a dynamic model for long-term as well as short-term price forecasting. The main conclusion derived from this study is that the ARIMA model for short-term forecasting is more effective than for it is for long-term forecasting. Another example of exotic currency forecasting can be found in [4]. In this case, ARIMA is applied to the USD and PR currency pair with daily prices between April

2014 and May 2019. It highlights the importance of the stationarity of the series by taking first differences. The results obtained indicate a robust model with a difference between estimated prices and actual values of less than 1%. However, some authors have not found significant advantages of using the ARIMA model for the prediction of splits. For example, ref. [5] estimate a USD/EUR model for volatility prediction and not for price prediction. This may lead us to think that the ARIMA model may offer interesting results for the prediction of prices and their evolution, but not for determining their volatilities.

Other studies apply different neural networks for price prediction. For example, the authors of [6] compare the ARIMA model to a Backpropagation neural network for the daily prediction of the NGN/CNY and NGN/USD currency pairs. In this work, it is concluded that the neural network improves the results with a lower prediction error. Other authors incorporate additional variables in the neural network. For example, the authors of [7] add different moving averages as inputs to the network for currency prediction. As a result, the neural network performs better for three different error measures. There is no consensus on the amount and type of inputs to incorporate in different neural networks. This leads to a diversity of results when comparing these models with those of other prediction models [8].

However, recurrent neural networks of the LSTM type generally perform well, even with a single price lag as input. For example, the authors of [9] compare this network with others such as Elman's, concluding on the superiority of LSTM for short-term predictions. It even improves against other econometric models such as the VaR model or support vector machine, as described in [10], where they predict the USD/INR currency exchange with 97.83% accuracy.

In [11], a review of the literature on foreign exchange modelling and forecasting is carried out, in which it is concluded that LSTM networks are one of the best solutions for short-term forecasting. However, there is the possibility of building hybrid models in which the combination of two or more methodologies is intended to improve the results over those of the best single model. Thus, hybrid models allow us to increase the accuracy of predictions by reducing the risk of the inadequate use of a single model [12]. The results obtained by combining a model with at least one neural network are promising [13].

Some authors have carried out important reviews of hybrid models. A growing interest in this type of model has been detected, highlighting the hybridisation with neural networks and ARIMA models. These hybrid models can be combined at the same time or sequentially, and may have benefits in terms of predictive power [14]. For example, the authors of [15] analyse the USD/ALL exchange rate with monthly data from 2000 to 2015. They compared an ARIMA model to a a hybrid ARIMA–ANN model sequentially. To carry this out, they initially estimated the ARIMA model using the residuals as inputs for the neural network. They used different performance indicators such as RMSE, MAE, and MAPE. In all of them, the improvement in the prediction of the hybrid model was evident. Therefore, the combination of linear and non-linear models is effective. Based on the same idea, the authors of [16] propose a hybrid multiplicative model for price forecasting in which the prediction of the non-linear components of the data series (obtained through the neural network) are multiplied by the predictions of the linear components obtained through the ARIMA model. This multiplicative model seems to work well except for some short-term forecasts.

This paper compares the ARIMA model with the LSTM recurrent neural network, as well as a hybrid ARIMA-LSTM model. For this purpose, these models are applied to the daily closing price prediction of the currencies AUD/USD, GBP/USD, JPY/USD, NZD/USD, and EUR/USD as well as the cryptocurrency Bitcoin (BTC/USD).

The following section sets out the methodology of the different models applied for the prediction of the selected currencies. Next, the main comparative results between the different models are presented. Finally, the main conclusions and limitations of this work are presented.

2. Methods

This section summarises the two main models used in forecasting the closing prices of the selected currencies. First, the classical ARIMA model, widely used in time series forecasting, is presented. Secondly, the long short-term memory (LSTM) neural network is described, which is a type of recurrent neural network that is very efficient in short-term forecasting. Finally, a hybrid model, ARIMA-LSTM, is established, in which the predictions obtained using the ARIMA model are added as inputs to the network.

2.1. ARIMA Model

The general autoregressive integrated moving average (ARIMA) model introduced by the authors of [17] includes auto-regressive as well as moving average parameters and explicitly includes differencing in the formulation of the model. Specifically, the three types of parameters in the model are the autoregressive parameter (p), the number of differencing passes (d), and moving average parameters (q). In the notation introduced by Box and Jenkins, the models are summarised as ARIMA (p,d,q).

With the ARIMA model, although a non-stationary process exhibiting homogeneity with respect to the class of series can occur in many ways, they could have a non-constant mean at time-varying second moments such as that of constant variance, or have both of these properties.

Models useful for representing such behaviours can be obtained by supposing a suitable difference in the process in order for the series to be stationary.

Having a time series, X_t, where t represents the time index, the ARMA(p,q) model is expressed as follows:

$$X_t = \alpha_1 X_{t-1} + \alpha_2 X_{t-2} + \cdots + \alpha_p X_{t-p} - \theta_1 \varepsilon_{t-1} - \theta_2 \varepsilon_{t-2} - \cdots - \theta_q \varepsilon_{t-q} + \varepsilon_t \tag{1}$$

where α and θ are estimated coefficients and ε is the residual of forecasts. As can be seen in Equation (1), this model is constructed as a combination of the autoregressive process (AR) with past values of the variable and the moving average (MA) process with past predictions errors. On the other hand, the parameters p and q represent the number of lags selected, while the parameter d indicates the number of integrations (usually differencing or the application of logarithms) on the variable to make the series stationary.

2.2. LSTM Model

The LSTM model is a kind of recurrent neural network, that can capture the nonlinear and complex relationships between variables. This network was proposed by the authors of [18] and has been found to be effective at capturing long-term dependencies in sequential data and time series data.

Deep learning-based models, such as LSTM, have shown promising results in time series forecasting. This network can propagate activations to process different sequences including long distance dependencies [19]. To solve the vanishing problem, the recurrent unit is grouped into blocks with cells and three gates. These gates control the flow of information [20]. The LSTM architecture consists of a memory cell and the three gates (Figure 1). Each cell presents a different state (m_t) as information flows through each neuron. The different gates are activated depending on the previous state of the cell (m_t), the output from the previous neuron (h_{t-1}) and the new information input (x_t). Thus, the forget gate is activated to decide, on basis of the inputs, which part of the information to forget from the internal state of the cell. This gate is therefore used to remove or not remove a neuron. On the other hand, the input gate or relevant gate determines how much information from the past is incorporated into the neuron, i.e., how much information is memorised. Finally, the output gate calculates the output information of the cell taking into account the previous state of the cell and the new information.

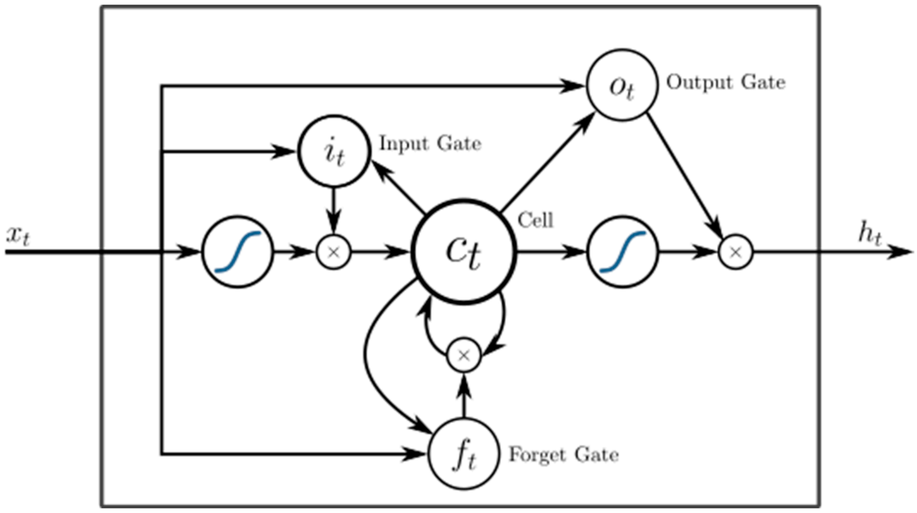

Figure 1. Internal cell structure of an LSTM. Source: "Creative Commons" by Eddie Antonio Santos licensed under BY CC-SA 4.0.

2.3. ARIMA-LSTM Model

As already indicated in the introduction, hybrid models generally show better predictions than models considered individually do. In this case, a two-stage hybrid ARIMA-LSTM model is considered. In the first stage, the ARIMA model is estimated using the Box–Jenkins methodology [17], taking into account the necessary transformations to obtain stationary series. Some authors use the residuals of the estimated model as the only input to the LSTM neural network in order to predict the non-linear patterns of the series. In this way, to obtain the final prediction, they use an additive or multiplicative model, combining in each case the prediction obtained via the ARIMA model (linear) with the non-linear patterns estimated using the LSTM network [21].

In this case, the use of the LSTM neural network as the prediction model is recommended. The first stage is identical to the process described above. That is, the ARIMA model is estimated. Then, in stage 2, two kinds of inputs are incorporated into the neural network. On the one hand, we add the lag closing prices of the time series. On the other hand, we add not the residuals, but the prediction of the prices achieved using the ARIMA model.

3. Results

This section compares forecasts for different currencies using ARIMA, LSTM and ARIMA-LSTM models. To make this comparison, different measures of prediction error are analysed, such as mean absolute error (MAE), mean absolute percentage error (MAPE) and root mean square error (RMSE).

3.1. Database

The following day closing prices have been selected for both majors and exotic currencies: EUR/USD, GBP/USD, JPY/USD, AUD/USD, and NZD/USD. The daily closing price of the Bitcoin cryptocurrency futures contract has also been selected to determine the behaviour of the models in this kind of asset. The database runs from 18 December 2017 to 27 January 2023.

3.2. Data Analysis and Processing

Currencies are time series with a regular presence of kurtosis and skewness. Table 1 shows the main descriptive statistics for EUR/USD as an example.

Table 1. Descriptive statistics for EUR/USD.

	Mean	Median	Sd	Max	Min	Skew.	Kurt.
EUR/USD	1.14	1.13	0.06	1.25	1.25	−0.54	0.02

On the other hand, for the use of ARIMA models, it is required that the series are stationary. The augmented Dickey–Fuller (ADF) test was used to calculate each currencies. An example of the ADF test on EUR/USD is shown in Table 2. As one can see, the p-value is greater than 0.05. Therefore, the null hypothesis of the stationarity of the series was rejected. In all cases, the stationarity hypothesis was rejected.

Table 2. Augmented Dickey–Fuller test for EUR/USD.

	Dickey–Fuller	p-Value
EUR/USD (original series)	−1.5054	0.07877
EUR/USD (log-diff series)	−11.83	0.01

These results imply the need for a transformation of the original series to make them stationary. For this purpose, the logarithm over a difference has been applied, obtaining stationary series (Table 2).

3.3. Model Estimation and Results

Once the stationary series were obtained, the different models described in Section 2 were estimated. The daily estimation and forecasting process for each currency was carried out by starting from the initial observations of the indicated database and adding each day after the forecast to estimate the new model. Different measures were used to compare the prediction error of the three models. Table 3 shows the main results.

Table 3. Measures of model prediction error.

Model	BTC	AUD/USD	GBP/USD	JPY/USD	NZD/USD	EUR/USD
ARIMA						
MAE	665.52	0.00616	0.00295	0.50610	0.00642	0.00394
MAPE	0.03212	0.00387	0.00337	0.00390	0.00377	0.00350
RMSE	1160.09	0.00831	0.00402	0.71309	0.00840	0.00513
LSTM						
MAE	28.81	0.00073	0.00010	0.18372	0.00059	0.00148
MAPE	0.00100	0.00047	0.00011	0.00136	0.00035	0.00137
RMSE	28.87	0.00075	0.00010	0.18429	0.00076	0.00148
ARIMA-LSTM						
MAE	23.57	0.00049	0.00018	0.17750	0.00078	0.00144
MAPE	0.00082	0.00032	0.00022	0.00131	0.00047	0.00133
RMSE	23.57	0.00049	0.00018	0.17750	0.00078	0.00144

Firstly, the advantage of the use of neural networks over the econometric ARIMA model is evident. In all the cases analysed, the LSTM neural network improves the prediction errors, reducing them by high percentages. For example, for EUR/USD, the percentage reduction in the error measures (MAE, MAPE, and RMSE) were, respectively, 62.4%, 60.9% and 71.2%. Similar results were obtained for other “Majors” currencies such as JPY/USD with percentages of 63.7%, 65.1% and 74.2%. However, these percentage reductions in the different measures of the prediction error increased for the “exotic” currencies analysed, including the BTC/USD cryptocurrencies. These percentage reductions reached levels around 90–91% for NZD/USD and between 95–97.5% for BTC/USD. This result may be related to the higher volatility in these exotic currencies and cryptocurrencies versus that of "Majors" currencies. For example, the volatility of EUR/USD is 0.0051 while for NZD/USD it increases to 0.0085.

On the other hand, the hybrid ARIMA-LSTM model seems to have improved the results obtained using the univariate LSTM model, although these improvements were relatively small. However, this model fails for the GBP/USD and NZD/USD currencies. The reason for the model's failure is unclear since while the GBP/USD currency is considered to be in the "Majors" category, the NZD/USD currency belongs to the "exotics" category. While GBP/USD has a volatility, measured by the deviation, of 0.004, NZD/USD has a higher volatility of 0.008. However, EUR/USD has a deviation of 0.0051 between the other two currencies. Therefore, it is also not a justification for the failure of the hybrid model. Further analysis in other time periods and an extension to other currencies would be desirable to determine the percentage of currencies where the hybrid ARIMA-LSTM model, as described in this paper, outperforms the univariate LSTM model.

4. Conclusions

In this paper, we have compared between the classical time series model (ARIMA) and the recurrent neural network LSTM. For this purpose, we have modelled and predicted the daily closing prices of different currencies, some of them considered majors and others exotic, as well as the cryptocurrency Bitcoin. The neural network was initially applied as a univariate model in which the input corresponded to a single lag of the closing price. The results suggest that this neural network is very efficient for short-term predictions, i.e., in this case, for the next period. Next, a hybrid ARIMA-LSTM model built in two phases was proposed. The first required the corresponding forecasts to be made using the ARIIMA model. In the second phase, these forecasts served as inputs to the network together with a price lag. This approach differs from that of other authors who propose that the input of the neural network should only be the residuals of the ARIMA model, given that these include the non-linear patterns. Finally, either additively or multiplicatively, the non-linear prediction of the network was combined with the linear prediction of the ARIMA model price. The results obtained with the hybrid model suggest a slight improvement with respect to the univariate LSTM neural network and, of course, with respect to the ARIMA model.

One of the limitations of this work was the lack of determination of whether or not this hybrid model also improves the prediction results when the sample data are not daily as in this case, i.e., when there are data with different timeframes (5 min, 15 min, etc.). On the other hand, the hybrid model proposed was not compared with other hybrid models in which the input is the residuals of the ARIMA model prediction, so the advantage of this one over the others is unknown. These analyses and comparisons are left for future work.

Author Contributions: Conceptualisation, F.G. (Fernado García) and F.G. (Francisco Guijarro); literature review, J.O. and R.T.; writing—original draft, J.O., F.G. (Fernando García) and R.T.; writing—review, F.G. (Francisco Guijarro). All authors have read and agreed to the published version of the manuscript.

Funding: This research received no external funding.

Institutional Review Board Statement: Not applicable.

Informed Consent Statement: Not applicable.

Data Availability Statement: Publicly available datasets were analysed in this study. This data can be found here https://www.visualchart.es.

Conflicts of Interest: The authors declare no conflict of interest.

References

1. Ghalayini, L. Modeling and Forecasting the US Dollar/Euro Exchange Rate. *Int. J. Econ. Financ.* **2014**, *6*, 194–207. [CrossRef]
2. Nwankwo, S.C. Autorregressive Ingrated Moving Average (ARIMA) Model for Exchange Rate (Naira to Dollar). *Acad. J. Interdiscip. Stud.* **2014**, *3*, 429–433. [CrossRef]
3. Ufuk Yildiran, C.; Fettahoglu, A. Forecasting USDTRY rate by ARIMA method. *Cogent Econ. Financ.* **2017**, *5*, 1335968. [CrossRef]
4. Asadullah, M. Forecast Foreing Exchange Rate: The Case Study of PKR/USD. *Mediterr. J. Soc. Sci.* **2020**, *11*, 129–137. [CrossRef]
5. Dunis, C.; Huang, X. Forecasting and trading currency volatility: An application of recurrent neural regression model and model combination. *Liverp. Bus. Sch. Woring Pap.* **2002**, *21*, 317–354. [CrossRef]
6. Mbaga, Y.V.; Olubusoye, O.E. Foreign Exchange Prediction: A Comparative Analysis of Foreign Exchange Neural Network (FOREXNN) and ARIMA Models. 2014. Available online: https://www.researchgate.net/publication/280040546 (accessed on 5 April 2023).
7. Kamruzzaman, J.; Sarker, R.A. Comparing ANN Based Models with ARIMA for Prediction of Forex Rates. *ASOR Bull.* **2003**, *22*, 2–11.
8. Huang, W.; Lai, K.K.; Wang, S. Forecasting Foreign Exchange Rates With Artificial Neural Networks: A Review. *Int. J. Inf. Technol. Decis. Mak.* **2004**, *3*, 145–165. [CrossRef]
9. Escudero, P.; Alcocer, W.; Paredes, J. Recuerrent Neural Networks and ARIMA Models for Euro/Dollar Exchange Rate Forecasting. *Appl. Sci.* **2021**, *11*, 5658. [CrossRef]
10. Kaushik, M.; Giri, A.K. Forecasting Foreign Exchange Rate: A Multivariate Comparative Analysis between Traditional Econometric, Contemporary Machine Learning & Deep Learning Techniques. *arXiv* **2002**, arXiv:2002.10247. [CrossRef]
11. Islam, M.S.; Hossain, E.; Rahman, A.; Shahadat, M.; Andersson, K. A Review on Recent Advancements in Forex Currency Prediction. *Algorithms* **2020**, *13*, 186. [CrossRef]
12. Hajirahimi, Z.; Khashei, M. Hybrid structures in time series modeling and forecasting: A review. *Eng. Appl. Artif. Intell.* **2019**, *86*, 83–106. [CrossRef]
13. Khashei, M.; Bijari, M. A new class of hybrid models for time series forecasting. *Expert Syst. Appl.* **2012**, *39*, 4344–4357. [CrossRef]
14. Zougagh, N.; Charkaoui, A.; Echchatbi, A. Artificial intelligence hybrid models for improving forecasting accuracy. *Procedia Comput. Sci.* **2021**, *184*, 817–822. [CrossRef]
15. Mucaj, R.; Sinaj, V. Exchange Rate Forecasting using ARIMA, NAR and ARIMA-ANN Hybrid Model. *J. Multidiscip. Eng. Sci. Technol.* **2017**, *4*, 8581–8586.
16. Wang, L.; Zou, H.; Li, L.; Chaudhry, S. An ARIMA-ANN Hybrid Model for Time Series Forecasting. *Syst. Res. Behav. Sci.* **2013**, *30*, 244–259. [CrossRef]
17. Box, G.E.P.; Jenkins, G.M. *Time Series Analysis: Forecasting and Control*; Holden-Day: San Francisco, CA, USA, 1970.
18. Hochreiter, S.; Schmidhuber, J. Long Short-Term Memory. *Neural Comput.* **1997**, *9*, 1735–1780. [CrossRef] [PubMed]
19. Kelleher, J. *Deep Learning*; The MIT Press: Cambridge, MA, USA, 2019; ISBN 9780262537551.
20. Altché, F.; de La Fortelle, A. An LSTM Network for Highway Trajectory Prediction. In Proceedings of the IEEE 20th International Conference on Intelligent Transportation Systems (ITSC), Yokohama, Japan, 16–19 October 2017; pp. 353–359. [CrossRef]
21. Zhang, G.P. Time series forecasting using a hybrid arima and neural network model. *Neurocomputing* **2003**, *50*, 159–175. [CrossRef]