International Research Journal of Engineering and Technology (IRJET) e-ISSN: 2395-0056
Volume: 10 Issue: 10 | Oct 2023 www.irjet.net p-ISSN: 2395-0072

Employing Advanced Neural Networks for Forecasting Time Series Data

Aayush Devgan[1]

[1]*Student, Computer Science Department, VIT University, Vellore, Tamil Nadu*

---***---

Abstract - *Intricate deep learning techniques are used in this paper to provide a method for identifying patterns in time series data. The method, which was tested on two different datasets, starts with data arrangement, moves on to neural network creation, and ends with a thorough visual examination. The findings show that when a dataset's time series exhibits consistent patterns, even if they differ somewhat, it is possible to train complex neural architectures with just one such series. These systems, however, produce results that are comparable to a straightforward technique that replicates the prior data entry in settings with less predictable datasets, such as stock market final prices. Every technique and experiment mentioned in this study is open to the public.*

Key Words: Deep learning, Machine learning, Time series, Forecasting

1. INTRODUCTION

This study offers a strategy based on the Long-Short Term Memory (LSTM) and Gated Recurrent Unit (GRU) deep learning architectures.

Recurrent neural networks (RNNs), a form of a deep learning network, are known to model dependencies over time [4]. They are pertinent to time series forecasting as a result. Thanks to its input, forget, and output gates, the LSTM chooses to retain content [3]. Reset and update gates make up the GRU [2]. Both networks have been widely applied to sequential issues since their introduction.

This paper focuses on investigating deep learning networks despite the fact that there are traditional approaches for time series forecasting, such as Autoregressive Integrated Moving Average (ARIMA). In fact, earlier research has demonstrated that LSTM beats ARIMA on financial data [5], and among several deep learning models, LSTM and GRU produce low predicting errors [1]. The procedure is then briefly discussed. You may read an extensive account of it in [6].

2. METHODOLOGY

As shown in Fig. 1, the procedure involves data preparation, model training, evaluation, and visual inspection. Data preparation includes normalization, defining how the train and test sets will be split up, and choosing a time series to train on. First, the data set's time series are each normalized between 0 and 1.

Then, as shown in Fig. 2, a time series of length Q samples is created, where w is the window size, f is the number of forecasting steps, and N is the number of training samples. The rest of the samples are put through testing.

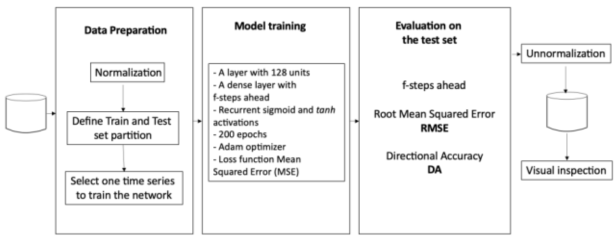

Fig. 1. A visual summary of the method

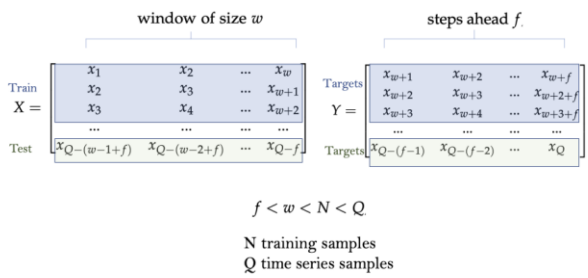

Fig. 2. Data partitioning. From [6]

An LSTM or GRU network is trained using a layer of 128 units, followed by a dense layer that provides an f-step ahead prediction. The networks are trained using the Adam optimizer for 200 iterations using the Mean Squared Error (MSE) Loss function.

Measurements for evaluation include Directional Accuracy (DA) and Root Mean Squared Error (RMSE) between actual and projected values on the test set. Every time series is then unnormalized and presented for visual analysis to help you better comprehend the outcomes.

3. EXPERIMENTS

In Fig. 3, the experimental configuration is shown. Two datasets with ten-time series each have been used to test the approach. The Activities dataset, which has ten

International Research Journal of Engineering and Technology (IRJET) e-ISSN: 2395-0056
Volume: 10 Issue: 10 | Oct 2023 www.irjet.net p-ISSN: 2395-0072

synthetic time series with five days of high activity and two days of low activity, is the initial dataset. This dataset might reflect, for instance, the number of calls made each week (see Fig. 4). The second dataset is the BANKEX dataset, which includes ten financial institutions' stock market closing prices (see Fig. 5). The impact of normalization between 0 and 1 is depicted in Fig. 6.

The initial time series of each dataset was utilized for data preparation, using a window of size w=60 days. Each series' final 251 samples were used for testing. LSTM, GRU networks, and a Baseline that only repeats the most recent observed value were used for forecasting. Each model underwent a one-step and twenty-step RMSE and DA evaluation.

The mean and standard deviation (SD) of RMSE and DA for the 10-time series on the test set of each dataset are summarized in Tables 2 through 5. It is ideal to have RMSE and DA values that are near zero. The results for the Activities dataset are shown in Tables 2 and 3. The blue highlights indicate the best outcomes.

GRU considerably outperforms LSTM and the Baseline for One-Step Forward regarding RMSE. However, on DA, both deep learning networks outperform the Baseline by a large margin. Considering RMSE and DA, LSTM is the undisputed winner for the twenty-step forecast. The networks demonstrate their ability to learn patterns that repeat, even with some change, on the Activities dataset.

The results for the BANKEX dataset are shown in Tables 4 and 5. Due to the nature of the stock market series, it's possible that the networks in this situation behave exactly like the Baseline. Last but not least, visual examination facilitates comprehension of the numerical data; see Figures 7 and 8.

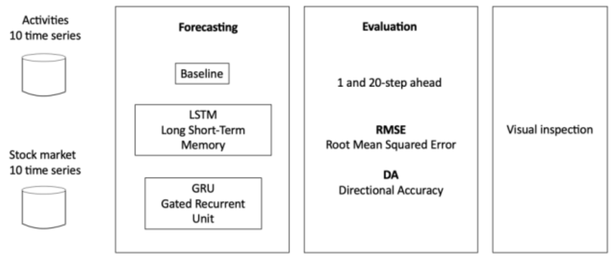

Fig. 3. Experimental setup

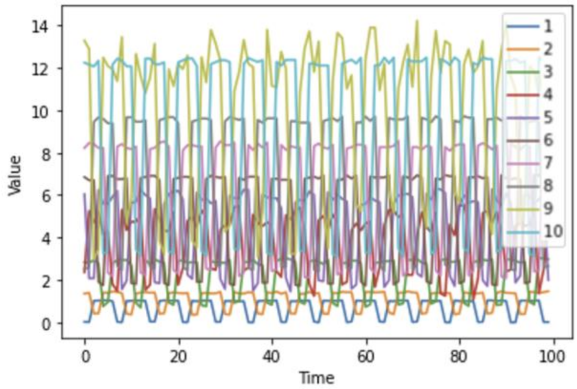

Fig. 4. Activities dataset (synthetic). The first 100 samples are plotted. The dataset contains 3 584 samples per series. Taken from [6]

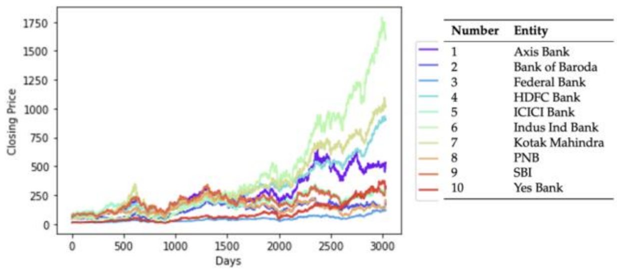

Fig. 5: BANKEX Dataset. Closing Price in Indian Rupee (INR). Daily samples retrieved between 12 July 2005 and 3 November 2017. All-time series with 3032 samples. Taken from [6]

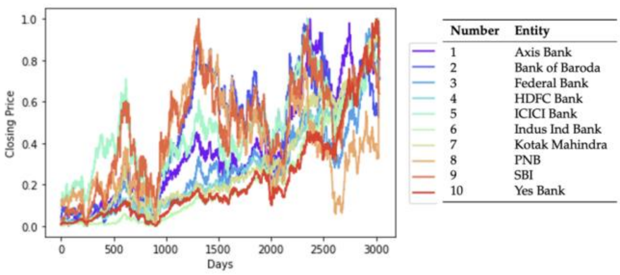

Fig. 6: Normalized BANKEX Dataset. Daily samples retrieved between 12 July 2005 and 3 November 2017 All-time series with 3 032 samples. Taken from [6]

4. CONCLUSIONS

The following are some key findings from this paper's demonstration of a technique for time series forecasting utilizing LSTM and GRU deep learning networks:

- It demonstrates that, provided the data is correctly prepared, LSTM and GRU networks may be trained for forecasting with a single time series in a dataset of series with patterns that repeat even with slight variation.

- It displays how the approach performed on two datasets. The method is suitable for time series with repeating patterns, such as those found in weekly activities, but it is not suitable for stock market data, potentially because certain information is not stored in the closing price alone or because of the nature of the problem.
- The ability to foresee twenty steps ahead as well as one step ahead is versatile.
- Visual inspection aids in understanding the numerical results in addition to the numerical evaluation offered by RMSE and DA.

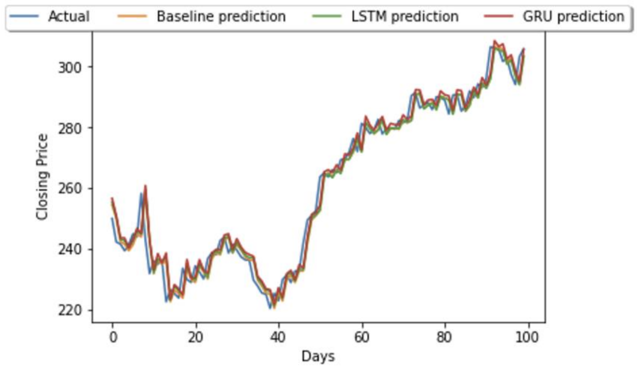

Fig. 7: An example of a 1-step ahead forecast. Actual and predicted closing price over the first 100 days of the test set Yes Bank. Closing Price in Indian Rupee (INR). Taken from [6]

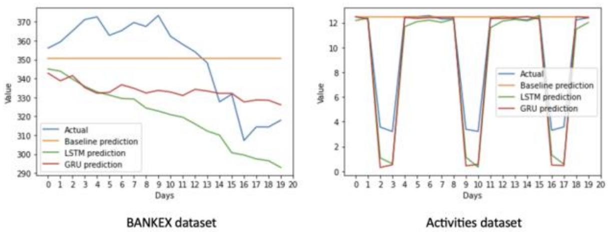

Fig. 8: Examples of 20-step ahead forecast. Taken from [6]

Table 2: One-step ahead forecast on Activities dataset.

	RMSE			DA		
	LSTM	GRU	Baseline	LSTM	GRU	Baseline
Mean	0.2949	0.1268	0.3730	0.6360	0.6236	0.4212
SD	0.0941	0.0425	0.0534	0.0455	0.0377	0.0403

Table 3: Twenty-step ahead forecast on Activities dataset.

	RMSE			DA		
	LSTM	GRU	Baseline	LSTM	GRU	Baseline
Mean	0.1267	0.2048	0.4551	0.6419	0.6261	0.4805
SD	0.0435	0.0683	0.0678	0.0331	0.0255	0.0413

Table 4: One-step ahead forecast on BANKEX dataset.

	RMSE			DA		
	LSTM	GRU	Baseline	LSTM	GRU	Baseline
Mean	0.0163	0.0163	0.0161	0.4884	0.4860	0.4880
SD	0.0052	0.0056	0.0056	0.0398	0.0385	0.0432

Table 5: Twenty-step ahead forecast on BANKEX dataset.

	RMSE			DA		
	LSTM	GRU	Baseline	LSTM	GRU	Baseline
Mean	0.0543	0.0501	0.0427	0.5004	0.5004	0.4969
SD	0.0093	0.0064	0.0113	0.0071	0.0087	0.0076

5. CODE AVAILABILITY

The developed code is made available on GitHub https://github.com/aayushdevgan/time-series-forecasting

REFERENCES

[1] Balaji, A.J., Ram, D.H., Nair, B.B.: Applicability of deep learning models for stock price forecasting an empirical study on bankex data. Procedia computer science 143, 947–953 (2018)

[2] Cho, K., Van Merri¨enboer, B., Bahdanau, D., Bengio, Y.: On the properties of neural machine translation: Encoder-decoder approaches. arXiv preprint arXiv:1409.1259 (2014)

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

[4] Rumelhart, D.E., Hinton, G.E., Williams, R.J.: Learning representations by back-propagating errors. nature 323(6088), 533–536 (1986)

[5] Siami-Namini, S., Tavakoli, N., Namin, A.S.: A comparison of arima and lstm in forecasting time series. In: 2018 17th IEEE international conference on machine learning and applications (ICMLA). pp. 1394–1401. IEEE (2018)

[6] Velarde, G., Bra˜nez, P., Bueno, A., Heredia, R., Lopez-Ledezma, M.: An open source and reproducible implementation of lstm and gru networks for time series forecasting. Engineering Proceedings 18(1), 30 (2022), https://doi.org/10.3390/engproc2022018030