



Comparative Study of Predicting Stock Index Using Deep Learning Models

Harshal Patil^{1,3}, Bharath Kumar Bolla^{2(✉)}, E. Sabeesh⁴,
and Dinesh Reddy Bhumireddy⁵

¹ Upgrad Education Pvt. Ltd., Mumbai, India

² Salesforce, Hyderabad, India

bolla111@gmail.com

³ Liverpool John Moores University, London, UK

⁴ IIT Jodhpur, Jodhpur, India

⁵ Cardinal Health, Bengaluru, India

Abstract. Time series forecasting has seen many methods attempted over the past few decades, including traditional technical analysis, algorithmic statistical models, and more recent machine learning and artificial intelligence approaches. Recently, neural networks have been incorporated into the forecasting scenario, such as the LSTM and conventional RNN approaches, which utilize short-term and long-term dependencies. This study evaluates traditional forecasting methods, such as ARIMA, SARIMA, and SARIMAX, and newer neural network approaches, such as DF-RNN, DSSM, and Deep AR, built using RNNs. The standard NIFTY-50 dataset from Kaggle is used to assess these models using metrics such as MSE, RMSE, MAPE, POCID, and Theil's U. Results show that Deep AR outperformed all other conventional deep learning and traditional approaches, with the lowest MAPE of 0.01 and RMSE of 189. Additionally, the performance of Deep AR and GRU did not degrade when the amount of training data was reduced, suggesting that these models may not require a large amount of data to achieve consistent and reliable performance. The study demonstrates that incorporating deep learning approaches in a forecasting scenario significantly outperforms conventional approaches and can handle complex datasets, with potential applications in various domains, such as weather predictions and other time series applications in a real-world scenario.

Keywords: ARIMA · SARIMA · SARIMAX · RNN · CNN · LSTM · GRU · DeepAR · DSSM · DF-RNN · Deep Renewal · POCID · Thiels'U

Time series forecasting has been implemented traditionally using standard methods such as ARIMA, SARIMA, and SARIMAX [1]. A significant drawback of these methods has been their inability to handle multivariate datasets where exogenous variables significantly affect the forecasting predictions [2]. Furthermore, their accuracies of predictions have not been satisfying enough in many complex real-world scenarios [2]. The advent of Deep Learning has helped bridge this gap. Neural networks and their ability to achieve universal approximation is a well-established theory, as seen in scenarios such as regression and classification [3]. In the last two decades, models based on recurrent neural

networks (RNNs) and LSTMs (Long-short-term memory) have been widely used in the forecasting scenario with promising results, and their ability to process sequential data has been exploited to solve complex time series scenarios [4]. However, in the last decade, newer architectures such as Deep Factor RNN(DF-RNN) [5], DSSM [6], Deep AR [7], and Deep Renewal [8] have been shown to outperform classical RNN and LSTM-based deep learning models in various scenarios. Very little experimentation has been done using these approaches on the Stock Market data. Hence, a comparative study of these models on a widely established dataset such as the NIFTY 50 index would help establish the superiority of deep learning models over traditional approaches and evaluate the effectiveness of recent deep learning models. The research objectives are as follows.

- To evaluate the superiority of neural networks over traditional approaches in forecasting augmentation.
- To evaluate the performance of the models, varying levels of train data (50% and 25%), keeping the test data the same to assess the effect of lesser data on model's performance.
- To evaluate if the better models are performing consistently on all metrics (MSE, RMSE, POCID, Theil's U, MAPE).

1 Literature Review

Various research has been done in the recent past to increase the efficiency of time series forecasting by incorporating deep learning methodologies. While traditional approaches have been used to solve time series problems in a univariate scenario, deep learning approaches have been used to approximate multivariate datasets with significantly higher efficiency.

1.1 Traditional Approaches

Time series forecasting has been an important research field since humans started to predict values associated with a time component. According to De Gooijer and Hyndman [9], the earliest statistical models for time series analysis, namely the Auto Regressive (AR) and Moving Average (MA) models, were developed in the 1940s.

These models aimed to describe time series autocorrelation and were limited to linear forecasting problems. As researchers delved deeper into the subject, they factored in parametric influences. In the early 1970s, Box et al. [10] developed the Box-Jenkins method, a three-step iterative process for determining time series, which became a popular approach for time-series modeling. With the advent of computers and increasing processing power, Autoregressive integrated Moving Average (ARIMA) models were used empirically for univariate and multivariate time series forecasting [11]. In the 1980s and 1990s, researchers started incorporating seasonality in time series modeling. Various methods, including X-11, X-12-ARIMA, etc., used decomposition to obtain seasonality and apply it in time-series forecasting [12].

In the past few decades, many methods have been implemented to forecast various domains. These methods include simple traditional technical analysis (also known as

“charting”) of price charts [13], algorithmic statistical models [14], and more recent Machine Learning and Artificial Intelligent approaches [15]. Computational time series forecasting has applications in various fields, from weather and sales forecasting to finance-related forecasting (budget analysis, stock market price forecasting). It is an indispensable tool for all fields that rely on time factors. Methods including Autoregression, Box Jenkins, and Holt-Winters were used to yield generally acceptable results.

1.2 Deep Learning Approaches

Recently, novel techniques and models have emerged utilizing deep learning methodologies. For instance, the Long- and Short-Term Time-Series Network (LSTNet) incorporates Convolutional Neural Network (CNN) and Recurrent Neural Network (RNN) to capture both short-term and long-term dependencies in time-series data [16]. Another approach proposes using the Gaussian Copula process (GP-Copula) in conjunction with RNN [17]. The Neural Basis Expansion Analysis for Interpretable Time Series forecasting (NBEATS) achieved state-of-the-art performance in the recent M4 time-series prediction competition [18]. It has been observed that deep learning methods possess an edge over traditional techniques with regard to overfitting, as evidenced in previous research by [19].

Many studies have shown that classical deep learning and machine learning models outperform ARIMA models in time-series forecasting. Various complex models, including Multi-Layer Perceptron, CNN, and LSTM, have been implemented and analyzed for time-series forecasting. These models can handle multiple input features, leading to higher accuracy than conventional methods. Feature extraction is a critical step in improving the performance of predictive models, even when simple features are used. Some studies have used modified deep networks to extract frequency-related features from time-series data using EMD and Complete Ensemble Empirical Mode Decomposition with Adaptive Noise (CEEDMAN). The extracted features were then fed to LSTM to predict one-step-ahead forecasting [20, 21]. Other studies have used image data features by decomposing raw time-series data into IMFs using IF and providing CNN to learn features automatically [22]. Data augmentation approaches such as adding external text-based sentiment data to the model-generated features, were also used [23].

Additionally, auto-regressive models have been proposed, such as DeepAR, which uses high-dimensional related time-series features to train Autoregressive Recurrent Neural Networks and has demonstrated superior performance compared to other competitive models [7]. Another study proposed a Multi-Step Time-Series Forecaster that uses various related time-series features to forecast demand on Amazon.com [24]. Furthermore, several state-of-the-art methods have been developed and proven to be highly promising in generalized competitions like M4 [25]. Finally, it has been shown that an ensemble of models consistently performs better than any single model [26].

2 Research Methodology

A novel python-based library, GluonTS, has been introduced to provide models, tools, and components for time-series forecasting [27]. Techniques such as DF-RNN, DSSM, Deep AR, and LSTNet have been implemented using relevant libraries from the GluonTS framework. The models used as baselines are ARIMA, SARIMA, SARIMAX, and Facebook’s Prophet. Facebook’s Prophet API is used for implementing the Prophet model.

2.1 Dataset Exploration

The dataset used in this research consists of the NIFTY 50 index consisting of the closing value of the stock indices from Jan 2011 till Feb 2022 (Fig. 1) [28]. A rolling window of length ten has been taken with context and prediction length of five each. Mean, and Standard deviation is calculated for ten window period to test for the stationarity of the time series. As seen in Fig. 2, where there is a variation in the mean and Standard deviation across the ten-window time frame, it is evident that the time series is not stationary. This is further confirmed by the ADF test, as seen in Table 1. Indices from Jan 2011 to 2020 have been used as Train, and the subsequent series have been used as Test data, as seen in Fig. 1.



Fig. 1. Test-Train Split – NIFTY-50 Index.

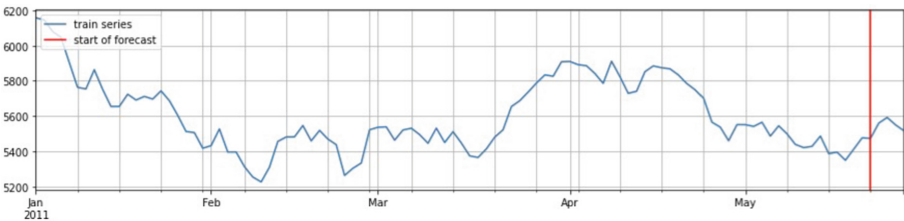


Fig. 2. Ten period Rolling Window – Non-uniform mean & STD – Non-stationary series.

2.2 Forecasting Models

Various models used in the experiments have been elaborated in the succeeding sections. Baseline forecasting models have been built using machine learning models such as ARIMA, SARIMA, and SARIMAX, while deep learning models have been built using DF-RNN, DeepAR, DSSM, and LSTNet.

ARIMA

ARIMA model has been created using the PMDARIMA library and hyper-parameters such as the number of auto-regressive terms (p), number of non-seasonal differences needed for stationarity (d), and number of lagged forecast errors in the prediction equation (q). The hyper parameters passed are shown below.

- $p - 0, 1, 2$
- $q - 0, 1, 2$
- Test to determine 'd' – Augmented Dickey-Fuller Test

The best model hyperparameters for ARIMA are (0,1,1) (Fig. 3).

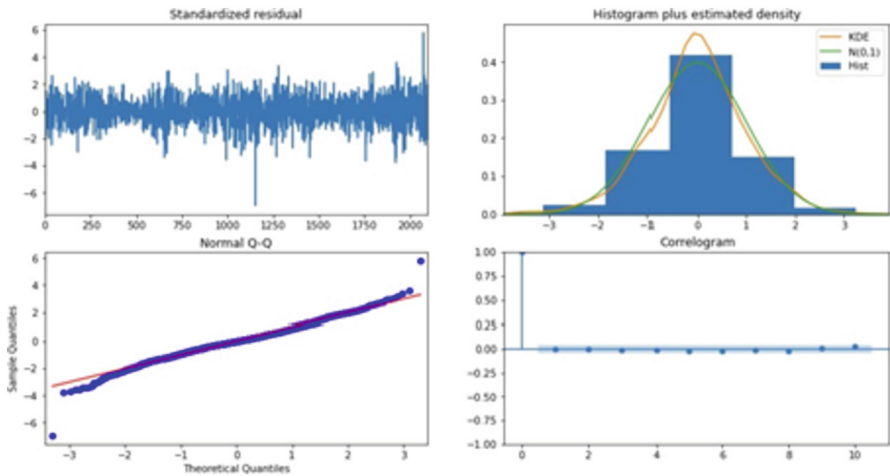


Fig. 3. ARIMA.

SARIMA

The model architecture is like ARIMA, except that SARIMA has an additional seasonal component. In addition to the existing hyperparameters of p , q , and d , seasonal hyperparameters, as mentioned below, are passed into the model. The best model hyperparameters for SARIMA are (0,1,1) (0,0,1) [30], as discussed below.

- $P -$ Seasonal Autoregressive order = 0.
- $D -$ Seasonal Difference order = 0.
- $Q -$ Seasonal Moving Average order = 1.

- m – Number of timesteps for a single seasonal period.

SARIMAX

SARIMAX is similar to SARIMA except for the addition of an exogenous variable which is used as an additional feature in the learning process.

Deep Factor RNN (DF-RNN)

The Deep Factor RNN is a model that incorporates two significant factors, namely global and local fluctuations, to govern the progression of a given time series. These factors are learned using separate RNN models, each controlled by specific hyperparameters that determine the model’s architecture, as illustrated in Table 1.

Table 1. Global vs. Local RNN Hyperparameters of DF-RNN.

Hyper-Parameters	Global RNN Model	Local RNN Model
Number of Hidden Layers	1	5
Number of Neurons in the Hidden Layer	50	5
Type of Cell	LSTM	LSTM

Deep AR

Deep AR, a probabilistic forecasting model, utilizes a global model to learn jointly from multiple related time series using negative binomial likelihood. The model is constructed on a recurrent neural network based on Long Short-Term Memory (LSTM) architecture. Hyperparameters, depicted in Table 2, are used in model building.

Table 2. Deep AR Hyperparameters.

Hyper-Parameters	Deep AR Model
Number of LSTM Layers	3
Number of LSTM Cells	40
Scaling	Enabled
Learning Rate	0.001

Deep State Space Model

The Deep State Space model [29] works on the principles of parametrizing linear state space of individual time series with a jointly learned recurrent neural network.

Deep Renewal

Deep Renewal Processes are a probabilistic intermittent demand forecasting method. This method builds upon Croston's framework and molds its variants into a renewal process. The random variables, M (The demand size at non-zero demand point) and Q (The inter-demand interval) are estimated using a separate RNN (Table 3).

Table 3. Deep AR Hyperparameters.

Hyper-Parameters	Deep Renewal Model
Skip size for skip RNN layer	3
Auto regressive window size – Linear	40
Learning Rate	0.001

3 Performance Measurement

Performance measurement of the forecasting model has been done using evaluation metrics such as MSE, RMSE, MAE, MAPE and custom metrics such as POCID and Theil's U.

3.1 Mean Square Error/Root Mean Square Error

Mean squared error is the mean of the squared error between the target variable (original observation) and the output variable (the predicted variable) in a given time series. Root mean squared error applies a square root to the MSE. The mathematical representation is shown in Eq. 1.

$$\text{MSE} = 1/N \sum_{i=1}^N (\text{target}_i - \text{output}_i)^2 \quad (1)$$

3.2 Mean Absolute Percentage Error

Mean Absolute percentage error defines the percentage difference between the target variable and the output variable w.r.t the output variable. The mathematical representation is shown in Eq. 2. Lower the value of MAPE better the model performance.

$$\text{MAPE} = 100/N \sum_{i=1}^N (\text{target}_i - \text{output}_i)/\text{output}_i \quad (2)$$

3.3 Mean Absolute Error

Mean absolute error is the difference between the target and output variables. The lower the MAE better is the model performance. The mathematical representation is shown in Eq. 3.

$$\text{MAE} = 1/N \sum_{i=1}^N (\text{target}_i - \text{output}_i) \quad (3)$$

3.4 POCID – Prediction of Change in Direction

The prediction of change in direction is the percentage of the number of correct decisions in predicting whether the time series in the next time interval will increase or decrease. The mathematical representation is shown in Eq. 4. The higher the value of POCID, the better the model performance.

$$\text{POCID} = 100 * \frac{\sum_{t=1}^N D_t}{N}$$

$$\text{where } D_t = \begin{cases} 1, & \text{if } (target_t - target_{t-1})(output_t - output_{t-1}) > 0 \\ 0, & \text{otherwise} \end{cases} \quad (4)$$

3.5 Theil's U

Theil's U is similar to the mean squared error except that the error is normalized w.r.t output variable of the previous time interval. U lesser than 1 indicates a better performance of the model. A value equal to 1 indicates a random model and a value greater than 1 indicates a model worse than a random model. Hence it is ideal for achieving a U of value 0. The mathematical representation is shown in Eq. 5.

$$\text{Theil's U (Normalized mean squared error)} = \frac{\sum_{t=1}^N (target_t - output_t)^2}{\sum_{t=1}^N (output_t - output_{t-1})^2} \quad (5)$$

4 Results

As explained in the preceding sections, models have been evaluated using the metrics mentioned above on both machine learning and deep learning models.

4.1 Forecasting Using ARIMA, SARIMA and SARIMAX

Forecasting has been done over a 36-day horizon. From Table 4, it is evident that there is no significant difference in the model performance among the three baseline models. Furthermore, on analysis of standardized residuals for each forecasted point by all three models, they are uniformly distributed around the mean of zero.

4.2 Forecasting Using Facebook Prophet

The performance of the Facebook Prophet was assessed using cross-validation on the provided time series for various forecast horizons between 36 and 364 days. It was noted that the MSE and MAPE values increase linearly as the forecast horizon increases. This suggests that an increase in forecast horizon leads to a rise in prediction error. Moreover, when compared to the baseline models, namely ARIMA, SARIMA and SARIMAX, Prophet exhibited a relatively lower performance in terms of RMSE and MAPE (Table 5).

Table 4. ARIMA, SARIMA and SARIMAX Metrics.

Models	MSE	MAE	RMSE	MAPE
ARIMA	3003730.628	1387.5576	1733.1274	0.118263
SARIMA	3009965	1391.598	1734.925	0.118803
SARIMAX	3002241	1387.658	1732.698	0.118351

Table 5. Evaluation metrics for Facebook Prophet.

Error	36 Days	364 Days	Average
MSE	203,144	5,000,000	2,522,746
RMSE	Not Done	Not Done	1588.31
MAPE	0.05	0.19	0.12
POCID	Not Done	Not Done	80.48
Theil's U	Not Done	Not Done	494.96

4.3 Forecasting using Deep Learning Models on 100% Train Data

The Table 6 presents the performance evaluation of various deep learning models for time-series forecasting based on five evaluation metrics: MSE, RMSE, POCID, Theil's U, and MAPE. The models include RNN, GRU, LSTM, DF-RNN, DeepAR, DSSM, and Deep Renewal. Among these models, DeepAR performed the best with the lowest values for all evaluation metrics, indicating its high accuracy in forecasting. The LSTM and GRU models also performed well with relatively low values for all metrics. The DF-RNN model performs better than the DSSM and Deep Renewal models but not as well as the others. Therefore, the ranking of the models based on better performance is DeepAR, GRU, LSTM, DF-RNN, RNN, DSSM, and Deep Renewal.

Table 6. Deep Learning Models Performance on 100% train data.

Model	MSE	RMSE	POCID	Theil's U	MAPE
RNN	21341000	4619.6	49	30790	36.1
GRU	53167	230.6	52	892	5.5
LSTM	171396	414	52	898	17.4
DF-RNN	3754898	1937.8	25	1168	0.12
Deep AR	35600	188.7	75	12	0.01
DSSM	25041182	5004	75	7767	0.29
Deep Renewal	246363673	15696	75	77148	1.0

4.4 Forecasting using Deep Learning Models on 50% Train Data

Based on the metrics in the table, the Deep AR model seems to be the best performer (as depicted in Table 7), followed by the LSTM, GRU, and RNN models. The DF-RNN, DSSM, and Deep Renewal models performed poorly than the others.

Table 7. Deep Learning Models Performance on 50% train data.

Model	MSE	RMSE	POCID	Theil's U	MAPE
RNN	1779572	1334	52	883	16.5
GRU	330895	575	52	910	17.3
LSTM	438402	662	51	889	17.1
DF-RNN	56715481	7531	75	1168	0.12
DeepAR	2823	53.1	50	0.83	0.003
DSSM	303172659	17411	75	94228	0.99
Deep Renewal	245658060	15673	50	77269	0.99

4.5 Forecasting using Deep Learning Models on 25% Train Data

The table shows that the DeepAR model performs best with the lowest MSE, RMSE, Theil's U, and MAPE values, even with 25% of the actual train data (as depicted in Table 8). GRU and LSTM models perform similarly with slightly higher values of the evaluation metrics.

Table 8. Deep Learning Models Performance on 25% train data.

Model	MSE	RMSE	POCID	Theil's U	MAPE
RNN	3970916	1993	52	1114	18.2
GRU	252760	503	53	893	17.3
LSTM	2718187	1649	52	848	17.2
DF-RNN	56715481	4985	25	7719	0.32
DeepAR	1886	43	50	0.76	0.002
DSSM	2442782	1563	50	760	0.09
Deep Renewal	245151974	15657	50	77243	0.99

Overall, Deep AR and GRU consistently performed the best across all tables, regardless of the percentage of training data used. Surprisingly these models' performance was consistent despite lowering the train data (Figs. 4 and 5).

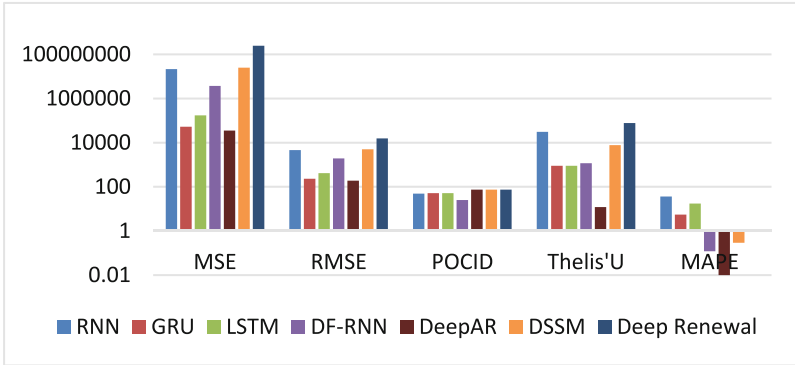


Fig. 4. Evaluation metrics of all models on 100% training data

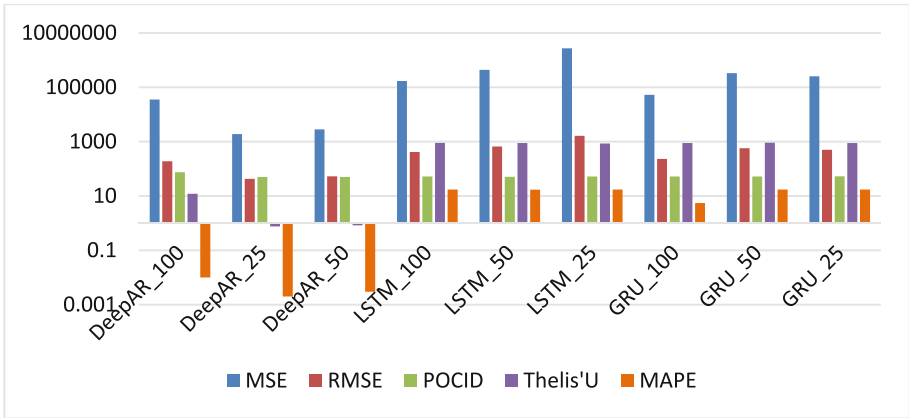


Fig. 5. Evaluation metrics of all models on 100% training data

5 Conclusion

This work compares traditional machine learning models with cutting-edge deep learning architectures for time series forecasting on stock market indices. The study employs several metrics, including MAE, MSE, RMSE, MAPE, POCID, and Theil's U, to evaluate the performance of the models. The results of the experiments demonstrate that state-of-the-art deep neural networks such as DeepAR and GRU outperform traditional forecasting models such as ARIMA, SARIMA, and SARIMAX. Moreover, DeepAR is stable across varying training data sizes and is consistent on all metrics. Furthermore, the study highlights the superiority of recurrent neural networks, their variants, such as LSTMs, for handling stock indices datasets, and their ability to outperform conventional machine learning and statistical-based algorithms. This makes them suitable for deployment in real-world scenarios. However, the study is limited by the use of univariate Stock market datasets. Future research on multivariate datasets could be explored to further establish the superiority of deep learning networks in time series forecasting.

References

1. Nokeri, T.C.: Forecasting using ARIMA, SARIMA, and the additive model. In: *Implementing Machine Learning for Finance: A Systematic Approach to Predictive Risk and Performance Analysis for Investment Portfolios*, pp. 21–50. Apress, Berkeley (2021)
2. Wang, S., Li, C., Lim, A.: Why are the ARIMA and SARIMA not sufficient. In: *Proceedings of arXiv* (2021). <http://arxiv.org/abs/1904.07632>. Accessed 06 June 2023
3. Sirisha, U.M., Belavagi, M.C., Attigeri, G.: Profit prediction using ARIMA, SARIMA and LSTM models in time series forecasting: a comparison. *IEEE Access* **10**, 124715–124727 (2022)
4. Sen, J., Mehtab, S.: Long-and-short-term memory (LSTM) networks architectures and applications in stock price prediction. In: *Emerging Computing Paradigms*, pp. 143–160. Wiley (2022)
5. Wang, Y., Smola, A., Maddix, D., Gasthaus, J., Foster, D., Januschowski, T.: Deep factors for forecasting. In: *Proceedings of the 36th International Conference on Machine Learning*, pp. 6607–6617. PMLR (2019). <http://proceedings.mlr.press/v97/wang19k.html>. Accessed 09 June 2023
6. Rangapuram, S.S., Seeger, M.W., Gasthaus, J., Stella, L., Wang, Y., Januschowski, T.: Deep state space models for time series forecasting. In: *Advances in Neural Information Processing Systems*, pp. 7785–7794 (2018)
7. Salinas, D., Flunkert, V., Gasthaus, J., Januschowski, T.: DeepAR: probabilistic forecasting with autoregressive recurrent networks. *Int. J. Forecast.* **36**(3), 1181–1191 (2020)
8. Snyder, R.D., Ord, J.K., Beaumont, A.: Forecasting the intermittent demand for slow-moving inventories: a modelling approach. *Int. J. Forecast.* **28**(2), 485–496 (2012)
9. De Gooijer, J.G., Hyndman, R.J.: 25 years of time series forecasting. *Int. J. Forecast.* **22**(3), 443–473 (2006)
10. Newbold, P.: The principles of the Box-Jenkins approach. *Oper. Res. Q.* (1970–1977) **26**(2), 397–412 (1975)
11. Kovvuri, A.R., Uppalapati, P.J., Bonthu, S., Kandula, N.R.: Water level forecasting in reservoirs using time series analysis – auto ARIMA model. In: Gupta, N., Pareek, P., Reis, M. (eds.) *IC4S 2022*. LNICST, vol. 472, pp. 192–200. Springer, Cham (2023). https://doi.org/10.1007/978-3-031-28975-0_16
12. Findley, D.F., Monsell, B.C., Bell, W.R., Otto, M.C., Chen, B.-C.: New capabilities and methods of the X-12-ARIMA seasonal-adjustment program. *J. Bus. Econ. Stat.* **16**(2), 127 (1998)
13. Lo, A.W., Mamaysky, H., Wang, J.: *Foundations of Technical Analysis: Computational Algorithms, Statistical Inference, and Empirical Implementation* (2000). <http://papers.ssrn.com/abstract=228099>. Accessed 21 June 2022
14. Webby, R., O’Connor, M.: Judgemental and statistical time series forecasting: a review of the literature. *Int. J. Forecast.* **12**(1), 91–118 (1996)
15. Sezer, O.B., Gudelek, M.U., Ozbayoglu, A.M.: Financial time series forecasting with deep learning: a systematic literature review: 2005–2019. *Appl. Soft Comput.* **90**, 106181 (2020)
16. Lai, G., Chang, W.-C., Yang, Y., Liu, H.: Modeling long- and short-term temporal patterns with deep neural networks (2018). <http://arxiv.org/abs/1703.07015>. Accessed 21 June 2022
17. Salinas, D., Bohlke-Schneider, M., Callot, L., Medico, R., Gasthaus, J.: High-dimensional multivariate forecasting with low-rank Gaussian Copula Processes (2019). <https://proceedings.neurips.cc/paper/2019/hash/0b105cf1504c4e241fcc6d519ea962fb-Abstract.html>. Accessed 21 June 2022
18. Oreshkin, B.N., Carпов, D., Chapados, N., Bengio, Y.: N-BEATS: neural basis expansion analysis for interpretable time series forecasting (2020). <http://arxiv.org/abs/1905.10437>. Accessed 21 June 2022

19. Khare, K., Darekar, O., Gupta, P., Attar, V.Z.: Short term stock price prediction using deep learning. In: 2017 2nd IEEE International Conference on Recent Trends in Electronics, Information & Communication Technology (RTEICT), pp. 482–486 (2017)
20. Guo, C., Kang, X., Xiong, J., Wu, J.: A new time series forecasting model based on complete ensemble empirical mode decomposition with adaptive noise and temporal convolutional network. *Neural. Process. Lett.* **55**, 4397–4417 (2022). <https://doi.org/10.1007/s11063-022-11046-7>
21. Cao, J., Li, Z., Li, J.: Financial time series forecasting model based on CEEMDAN and LSTM. *Physica A* **519**, 127–139 (2019)
22. Zhou, F., Zhou, H., Yang, Z., Gu, L.: IF2CNN: towards non-stationary time series feature extraction by integrating iterative filtering and convolutional neural networks. *Expert Syst. Appl.* **170**, 114527 (2021). <https://doi.org/10.1016/j.eswa.2020.114527>
23. Atha, S., Bolla, B.K.: Do deep learning models and news headlines outperform conventional prediction techniques on forex data? In: Rout, R.R., Ghosh, S.K., Jana, P.K., Tripathy, A.K., Sahoo, J.P., Li, K.C. (eds.) *Advances in Distributed Computing and Machine Learning*. LNNS, vol. 427, pp. 413–423. Springer, Singapore (2022). https://doi.org/10.1007/978-981-19-1018-0_35
24. Wen, R., Torkkola, K., Narayanaswamy, B., Madeka, D.: A multi-horizon quantile recurrent forecaster. *arXiv* (2018). <https://doi.org/10.48550/arXiv.1711.11053>
25. Makridakis, S.S., Assimakopoulos, V.: The M4 competition: results, findings, conclusion and way forward. *Int. J. Forecast.* **34**(4), 802–808 (2018). <https://doi.org/10.1016/j.ijforecast.2018.06.001>
26. Dietterich, T.G.: Ensemble methods in machine learning. In: Dietterich, T.G. (ed.) *MCS 2000*. LNCS, vol. 1857, pp. 1–15. Springer, Heidelberg (2000). https://doi.org/10.1007/3-540-45014-9_1
27. Alexandrov, A., et al.: GluonTS: probabilistic time series models in Python. *arXiv* (2019). <http://arxiv.org/abs/1906.05264>. Accessed 21 June 2022
28. A Novel Machine Learning Approach for Predicting the NIFTY50 Index in India | SpringerLink. <https://link-springer-com.ezproxy2.library.arizona.edu/article/10.1007/s11294-022-09861-8>. Accessed. 24 July 2023
29. Rangapuram, S.S., Seeger, M.W., Gasthaus, J., Stella, L., Wang, Y., Januschowski, T.: *Deep State Space Models for Time Series Forecasting* (2018)