

Original Research Article

Enhanced Inflation Forecasting by Using a Novel Hybrid Model: Empirical Evidence from Sri Lanka

Abstract

Aims/ objectives: This study aims to develop and evaluate the forecasting performance of a novel hybrid model (HB) for predicting monthly inflation rates. The primary objectives include assessing the predictive accuracy of the HB model in comparison to established models and demonstrating its adaptability and robustness over an extended period using historical data from Sri Lanka (1988-2021).

Study design: This research adopts a time series forecasting approach, specifically employing a unique hybrid model that combines Autoregressive Integrated Moving Average (ARIMA) and Artificial Neural Networks (ANNs). The model's efficacy is evaluated through an extensive analysis of monthly inflation rates.

Methodology: The hybrid model is applied to predict monthly inflation rates, leveraging the strengths of both ARIMA and ANN components. Historical data from Sri Lanka is utilized, and the forecasting accuracy is compared against other established models. The evaluation considers metrics such as Mean Absolute Percentage Error (MAPE) to quantify the predictive performance of the HB model.

Results: The empirical results reveal the superior forecasting accuracy of the hybrid model (HB) compared to other established models. The MAPE on the validation data is notably reduced to 7.10%, demonstrating the heightened predictive prowess of the HB model. Furthermore, the model exhibits adaptability and robustness over different periods, enhancing its credibility as a reliable forecasting tool.

Conclusion: The innovative fusion of ARIMA and ANN methodologies in the hybrid model represents a significant advancement in time series forecasting. The demonstrated effectiveness and reliability of the HB model suggest its potential applicability in diverse economic and financial forecasting domains. By providing more accurate predictions, the HB model empowers decision-makers to make informed choices, contributing to improved decision-making processes.

Keywords: *Inflation Rate; Hybrid Models; Neural Network; Forecasting; Nonlinear Models*

2010 Mathematics Subject Classification: 53C25; 83C05; 57N16

1 Introduction

The rate of inflation is an essential economic gauge that signifies the proportional fluctuation in the comprehensive valuation of goods and services in a given economy throughout a set interval, typically yearly Mankiw (2014). When the inflation rate is positive, it implies that prices are on an uptrend, thereby eroding the purchasing power of money. In contrast, a negative inflation rate, or deflation, suggests that prices are trending downwards Mankiw (2014). For policy-makers, monitoring the inflation rate is crucial as it informs decisions about monetary policies, such as managing the money supply and adjusting interest rates, with the ultimate goals of maintaining economic stability, fostering growth, and staving off recessions Mishkin (2010).

From a microeconomic perspective, understanding inflation trends is essential as it impacts a range of financial decisions Mankiw (2014). For businesses, insights into the trajectory of inflation can aid in forecasting future costs and setting suitable prices for their goods or services. For investors, the inflation rate has implications for returns on investments, bond yields, and the general health of the stock market. For individuals, inflation affects their cost of living and informs decisions about saving and consumption Mishkin (2010). Given these considerations, accurate forecasting of the inflation rate, the central aim of our study, has substantial benefits for a broad range of stakeholders.

Our investigation delves into the realm of time series forecasting, a branch of data analysis that predicts future data points by examining past observations Hyndman and Athanasopoulos (2014). This prognostication approach has seen extensive use in diverse fields, encompassing economics, finance, agriculture, meteorology, and biomedical research, largely due to the abundant and comprehensive datasets and the evolution in information processing technology Hyndman and Athanasopoulos (2014). A critical element of successful temporal sequence predictions is the detection and understanding of patterns and interconnections present in the data. These insights facilitate well-grounded forecasts about upcoming events.

The Box and Jenkins Box and Jenkins (1976) introduced model, known as the Autoregressive Integrated Moving Average (ARIMA), has emerged as a well-accepted approach in time series forecasting. The attractiveness of ARIMA models is anchored in their simplicity of use and versatility, making them a common selection for these predictive analyses. However, these models assume linearity, which posits a linear relationship among the time series observations Hyndman and Athanasopoulos (2014). This presumption often falls short when dealing with nonlinear data. To overcome this, various nonlinear models like bilinear models, threshold autoregressive models, and autoregressive conditional heteroskedasticity models have been proposed Tong (1990).

The usage of Artificial Neural Networks (ANNs) has also been observed in the domain of time series prediction Zhang (2003). ANNs offer a flexible computational framework capable of modelling a broad spectrum of nonlinear problems without necessitating prior assumptions about the model's structure Haykin (2009). Although both ARIMA and ANNs have their respective strengths in dealing with linear and nonlinear data, it has proven challenging to identify a single model that can provide accurate predictions for both types of data Zhang (2003). This has led scholars to investigate the possibilities of mixed models that integrate the advantages of both ARIMA and ANNs Zhang (2003).

Several studies have demonstrated the effectiveness of hybrid models. For example Zhang (2003) proposed an ANN-ARIMA hybrid model for forecasting nonlinear time series data, achieving a Root Mean Square Error (RMSE) of 0.0072. Similarly, Cadenas and Rivera Cadenas and Rivera (2009) utilized an ANN-ARIMA hybrid model to forecast the speed of wind in Mexico, with their hybrid model outperforming both standalone ANN and ARIMA models. Another study by Perera et al. Perera et al. (2012) applied an ANN-ARIMA hybrid model to forecast stock data in Colombo, showing that the hybrid model outperformed both ANN and ARIMA models in terms of Mean Absolute Percentage Error (MAPE) and RMSE.

In our study, we aim to assess the performance of a novel hybrid model built upon an auto-tuned algorithm, comparing it against the hybrid model proposed by Zhang Zhang (2003). We use historical inflation data from Sri Lanka, spanning from 1988 to 2021, as the benchmark for evaluating these models of Sri Lanka (2021). The implementation of the models is carried out using the R software package (version 4.1.3) and the Keras library for R Team (2022); Chollet et al. (2018). The performance of our hybrid model will be compared with the Zhang,s model suggested by Zhang Cadenas and Rivera (2009); Zhang (2003) in terms of forecasting accuracy and other relevant performance metrics.

2 Methodology

2.1 Brock- Dechert-Scheinkman Test(BDS Test)

The Brock-Dechert-Scheinkman (BDS) Test, introduced by Brock et al. (1996) . in 1987, is a renowned approach within chaos theory used to identify non linearity. Originally aimed at testing for independence and identical distribution (iid), the BDS test has proven useful for discovering a variety of linear and nonlinear structures. It can also be applied to residuals from a fitted model for conducting portmanteau or misspecification tests.

The underpinning concept of the above test is the integral correlation , which quantifies how often temporal patterns recur in the data. Given a time series X_t , defined as $X_t^m = (x_t, x_{t-1}, \dots, x_{t-m+1})$ and its m -history denoted as $X_t^m = (x_t, x_{t-m+1})$, the m -dimensional correlation integral is expressed as:

$$C_{m,T}(\epsilon) = \sum_{i < s} I_{\epsilon}(X_t^m, X_s^m) * \left\{ \frac{2}{T_m(T_m - 1)} \right\} \quad (2.1)$$

Here, $T_m = T - (m - 1)$ and $I_{X_t^m}, I_{X_s^m}$ are indicator functions. The function values of X_t^m, X_s^m are 1 if the sup norm $\| X_t^m, X_s^m \|$ and 0 in all other cases. Essentially, $C_{m,T}(\epsilon)$ records how frequently m -histories are found in proximity within a hyper-cube of a specific size. In other words, the correlation integral computes the probability that any two m -dimensional points are near each other.

$$P(| X_t - X_s | < \epsilon, | X_{t-1} - X_{s-1} | < \epsilon, \dots, | X_{t-m+1} - X_{s-m+1} | < \epsilon) \quad (2.2)$$

In the limiting case, assuming that the X_t are iid, this probability should equate to the following:

$$C_{1,T}(\epsilon)^m = P(| X_t - X_s | < \epsilon)^m \quad (2.3)$$

Brock et al. (1996) specified the BDS statistics as follows:

$$V_{m\epsilon} = T^{1/2} * \frac{C_{m,T}(\epsilon) - C_{1,T}(\epsilon)^m}{S_{m,T}} \quad (2.4)$$

Where $S_{m,T}$ is the standard deviation, which, according to Brock et al. (1996), can be consistently estimated. The distribution of the BDS statistic approaches a normal distribution, with an average of 0 and a variance of 1.

3 Data Set

In our research, we examine the data on Sri Lanka's monthly inflation rate from January 1988 through August 2021 TradingView (2023). Figure 01 depicts these inflation rate data.




Figure 1: Monthly mean Inflation rate of Sri Lanka(1988-2021)

The dataset employed in this research comprises a total of 407 data points, with an average of 9.61, variance 29.91554, minimum value of -0.890, and a peak value of 28.310. The data's time series representation shows a stochastic pattern with occasional outlier points, one noteworthy instance being in June 2008, likely associated with the nation's political climate at that moment. We divided the dataset into three segments: training data spanning from January 1988 to February 2020, which aids in checking the precision of the modeled fits; test data extending from March 2020 to February 2021; and validation data ranging from March 2021 to November 2021, which helps evaluate the accuracy of the projected values.

4 Models

4.1 Auto Regressive Integrated Moving Average of Order (p, d, q) , $(ARIMA(p, d, q))$

The ARIMA model Box and Jenkins (1976) is an expansion of the ARMA model that includes scenarios of non-stationarity as well. In ARIMA models, a non-stationary time series is transformed into a stationary one by implementing finite differences to the data points. The mathematical representation of the ARIMA(p, d, q) model using lag polynomials is provided as follows:

$$\left(1 - \sum_{i=0}^p \phi_i L^i\right)(1 - L)^d y_t = \left(1 + \sum_{j=0}^q \theta_j L^j\right) \epsilon_t \quad (4.1)$$

In this context, p, d, and q are integers that are equal to or greater than zero. They denote the order of the autoregressive, integrated, and moving average components of the model, respectively. The integer d determines the degree of differentiation. When d equals 0, the model simplifies to an ARMA(p,q) model. An ARIMA(p,0,0) corresponds to an AR(p) model, whereas ARIMA(0,0,q) aligns with an MA(q) model.

4.2 Artificial Neural Networks(ANN)

The single hidden layer feedforward neural network serves as the most prevalent approach for modeling and forecasting time series Haykin (1999). The architecture encompasses three tiers of simple processing units interconnected by acyclic links. There exists a mathematical connection between the output X_t and the inputs $X_{t-1}, X_{t-2}, \dots, X_{t-p}$:

$$X_t = \alpha_0 + \sum_{j=1}^p \alpha_j H(\beta_0 + \sum_{i=1}^q \beta_{i,j} X_{t-i}) + \epsilon_t \quad (4.2)$$

Here, α and β are the model parameters, p denotes the count of input nodes, and q symbolizes the number of hidden layers. The logistic function $H(t)$, typically employed as the transfer function for the hidden layer, is given as follows:

$$H(X) = \frac{1}{1 + \exp(-X)} \quad (4.3)$$

Based on equation (1) and equation (2), the ANN model essentially carries out a non-linear function that maps previous observations ($X_{t-1}, X_{t-2}, \dots, X_{t-p}$) to future values (X_t). This can be expressed as:

$$X_t = f(X_{t-1}, X_{t-2}, \dots, X_{t-p}, \mathbf{V}) + \epsilon_t \quad (4.4)$$

where V is a vector of all parameters, and f is determined by the network structure and the weights of connections. Therefore, the neural network corresponds to a nonlinear auto-regressive model.

4.3 Reiteration of G.Peter Zhang's Model(PZM)

G. Peter Zhang's hybrid model(PZM)Zhang (2003) harnesses the combined power of ARIMA and ANN models to enhance forecasting performance. The model initially applies an ARIMA model to scrutinize the linear components of the data, then an ANN model is developed to model the residuals generated from the ARIMA model, which encapsulate information about the data's non linearity.

By separately modeling linear and nonlinear patterns and integrating the forecasts, the hybrid model can bolster overall modeling and forecasting precision. Interestingly, the hybrid method may incorporate sub-optimal models, which paradoxically can enhance the model's utility as it's often more effective to amalgamate forecasts derived from varying information sets.

4.4 The Proposed ARIMA-ANN Hybrid Model (HB)

In accordance with Zhang et al. (2008), it is suggested that a time series is a mixture of a linear component and a non-linear component. Therefore, we can represent a time series X at time t as:

$$X_t = L_t + N_t \quad (4.5)$$

Here, L denotes the linear part of the time series X at time t , while N stands for the non-linear part. In the preliminary application of the ARIMA model to X , the non-linear part of X can be estimated as follows:

$$N_t = X_t - L'_t \tag{4.6}$$

Where L'_t is the estimated linear part of time series X at time t . In this study, 12 data points ahead were chosen as our test data length, and an "Akaike Information Criteria" based auto-tuned ARIMA model was used to forecast ten data points ahead using training data lengths ranging from 12 to 386, leveraging the power of Algorithm 2. The initial ten data points were omitted as the ARIMA model requires some data to make its first prediction. The advantage of this method is that it increases the length of training data, ensuring successful estimation of N_t . As pointed out by Faraway and Chatfield [?], the accuracy of ANN models is heavily dependent on training data length. N_t is estimated at each step as follows:

$$N_t = [X_{t_1}, X_{t_2} \dots X_{t_{12}}] - [L_{t_1}, L_{t_2} \dots L_{t_{12}}] \tag{4.7}$$

Let $X_t - L_t = \epsilon_t$ and denote N_t as:

$$N_{352 \times 10} = \begin{bmatrix} \epsilon_{1,1} & \epsilon_{1,2} & \dots & \epsilon_{1,h} \\ \epsilon_{2,1} & \epsilon_{2,2} & \dots & \epsilon_{2,h} \\ \vdots & \vdots & \ddots & \vdots \\ \epsilon_{L,1} & \epsilon_{L,2} & \dots & \epsilon_{L,h} \end{bmatrix} \tag{4.8}$$

We intend to model N_t based on univariate ANN modeling techniques and transform $N_{386 \times 12}$ to $N_{1 \times 4632}$:

$$N'_t = (N_t)^T \tag{4.9}$$

By treating N'_t as a univariate time series, we can model the non-linear component of X_t . Finally, we can modify [equation 6] as follows:

$$X_t = L'_t + N'_t + e_t \tag{4.10}$$

Here, e_t represents the error term generated from estimating N_t using ANN modeling processes. The objective of making an accurate forecast for this task [Algorithm 2] is to minimize e_t .

4.5 Error Calculation Methods

Let n denote the number of fitted points, Y_t represent the actual value of the response variable Y at time t , and \hat{Y}_t correspond to the predicted value of Y_t .

4.5.1 Mean Absolute Percentage Error (MAPE)

The Mean Absolute Percentage Error (MAPE) can be computed using the formula provided below.

$$MAPE = \frac{1}{n} \sum_{i=1}^n \frac{|Y_t - \hat{Y}_t|}{Y_t}$$

MAPE is the most commonly used measure of the forecast error and it works best in the absence of an extremes values in the data set.

5 Simulation Results

5.1 Pseudo Code for Optimal *ARIMA* Order

Algorithm 1: Pseudo Code for Optimal *ARIMA* Order

Input: X , p_{max} , d_{max} , q_{max} Output: Optimum *ARIMA* Order

- For $p \leftarrow 0$ to p_{max}
 - For $d \leftarrow 0$ to d_{max}
 - For $q \leftarrow 0$ to q_{max}
 - Try:
 - model \leftarrow fit(arima(p,d,q))
 - current_AIC \leftarrow AIC(model)
 - If $current_AIC < AIC$
 - AIC \leftarrow Optimum *ARIMA* Order
 - optimum_model \leftarrow model
 - Cat: error = function(e)
 - End If
 - End For
 - return Optimum *ARIMA* Order(p, d, q)
-

The algorithm detailed here delineates the procedure for determining the ideal order of an ARIMA model, a critical component of time series forecasting. The ARIMA model necessitates three parameters: p (autoregressive part's order), d (degree of first differencing involved), and q (moving average part's order). The algorithm launches a sequence of loops for each parameter, within the bounds of 0 to their respective maximum limits (p_{max} , d_{max} , and q_{max}). It then attempts to fit the ARIMA model with each parameter combination. The Akaike Information Criterion (AIC) for each model is computed and contrasted with the best (lowest) AIC noted so far. If a model results in a lower AIC, it supplants the current optimal model. Ultimately, the algorithm returns the parameters of the ARIMA model that produced the lowest AIC, delivering the optimal order for the ARIMA model for the provided time series data.

5.2 Pseudo code for proposed hybrid model

Algorithm 2: Pseudo code for proposed hybrid model

Input: D-Data set, p_max, d_max, q_max , H-test data length, L-the learning rate,K-Initial index
train *ARIMA* model to , Network structure

Output: Trained model

- For $i \leftarrow K$ to $length(D)$
 - For $p \leftarrow 0$ to p_{max}
 - For $d \leftarrow 0$ to d_{max}
 - For $q \leftarrow 0$ to q_{max}
 - Try:
 - model $\leftarrow fit(arima(p,d,q))$
 - current_AIC $\leftarrow AIC(model)$
 - If $current_AIC < AIC$
 - $AIC \leftarrow Optimum\ ARIMA\ Order$
 - optimum_model $\leftarrow model$
 - Cat: $error \leftarrow function(e)$
 - End If
 - **return** $Optimum\ ARIMA\ Order(p, d, q)$
 - Predict:(ARIMA) $h \leftarrow H$
 - $E \leftarrow D$ -Predict
 - End For
 - Transform(E) \leftarrow Univariate T.S
 - Initialize all weights and biases ;
 - while terminating condition is not satisfied
 - for each training tuple X in D
 - For each input layer unit j
 - $O_j = I_j$;
 - $I_j = \sum_i W_{i,j} O_i + \Theta_j$
 - $O_j = \frac{1}{1+e^{-I_j}}$;
 - For each unit j in the output layer ;
 - $E_j = O_j(1-O_j)(T_j - O_j)$
 - For each unit j in the hidden layers, from the last to the first hidden layer;
 - $E_j = O_j(1-O_j) \sum_k E_k W_{j,k}$
 - For each weight $W_{i,j}$ in network
 - $\Delta W_{i,j} = L(E_j O_i)$;
 - $W_{i,j} = W_{i,j} + \Delta W_{i,j}$;
 - For each bias Θ_j in network
 - $\Delta \Theta_j = L(E_j)$;
 - $\Theta_j = \Theta_j + \Delta \Theta_j$;
-

Algorithm 2 shows that the pseudo code for proposed model. Here H is the test data length and K is the initial index of univariate data set and integer H needs minimum values to train the *ARIMA* model. In this case we selected 10 data points to train the *ARIMA* model.

5.3 BDS Test Results

m(Embedding dimension) = 2, 3, 4
 Epsilon for close points = 2.7348, 5.4695, 8.2043, 10.9390

Table 1: Results of BDS test

m	[2.7348]	[5.4695]	[8.2043]	[10.939]
[2]	$p < 2.2e-16$	$p < 2.2e-16$	$p < 2.2e-16$	$p < 2.2e-16$
[3]	$p < 2.2e-16$	$p < 2.2e-16$	$p < 2.2e-16$	$p < 2.2e-16$
[4]	$p < 2.2e-16$	$p < 2.2e-16$	$p < 2.2e-16$	$p < 2.2e-16$

Table 1 shows that results of BDS test. According to those results our data set is not a indentation identical distribution (*iid*). Because P-values of close points are close to zero in first four Embedding dimension and we can reject null hypothesis.

5.4 Predicted values for test data




Figure 2: Predicted data comparison with Test Data

The displayed plot(Figure 02) is a multi-line chart representing the Inflation Rate predictions made by four different models – Test, HB and PZM, over a period from March 2020 to February 2021. Each line, distinguished by a unique color, represents a specific model's prediction trajectory over the specified

period. By observing the line trends, one can gain insights into the behavior of each model. For instance, a line consistently above others indicates a model with a higher inflation rate prediction. Conversely, a jagged line with noticeable peaks and troughs suggests volatile model predictions with significant fluctuations over time. Thus, this plot offers a comprehensive visual depiction of the models' performances, allowing for an easy comparison and identification of patterns in their respective predictions.

Table 2: Test data comparison of predicted data

Date	Test	HB	PZM
Mar-2020	5.4	5.44	5.96
Apr-2020	5.2	5.15	5.29
May-2020	4	3.97	4.22
Jun-2020	3.9	3.98	3.50
Jul-2020	4.2	4.04	4.38
Aug-2020	4.1	4.08	4.16
Sep-2020	4	3.86	4.29
Oct-2020	4	4.07	4.19
Nov-2020	4.1	4.06	3.23
Dec-/2020	4.2	4.15	4.41
Jan-2021	3	2.90	3.43
Feb-2021	3.3	3.29	3.53
MAPE	-	1.60	7.83

The Table 02 presents the predicted inflation rates(test data) by HB and PZM, across a time frame from March 2020 to February 2021. Each row corresponds to a specific month, while each column signifies the inflation rate predicted by a particular model for that month. It allows for a thorough examination of the individual model's performance over time by comparing the data vertically within each column.

Most notably, the table includes the Mean Absolute Percentage Error (MAPE) at the end, a key indicator of the accuracy of the model's predictions. The MAPE for HB, and PZM, is provided as 1.60% and 7.83%, respectively. These MAPE values serve as benchmarks for model comparison, with lower values signifying higher predictive accuracy. For instance, with a MAPE of 1.60%, the HB model exhibits the most accurate predictions among the two models in this case.

In summary, this table not only provides a granular view of each model's monthly predictions but also quantitatively compares their performance via the MAPE metric. This approach, in conjunction with the visual comparison offered by the line graph, facilitates a comprehensive evaluation of the models' inflation rate predictions.

5.5 Forecasted values for validation data




Figure 3: Validation Data with Forecasted values

The Figure 03 depicted here showcases the inflation rate forecasts from March 2021 to November 2021 using two different models: HB and PZM. The blue line represents the validation data, which serves as the actual observed data against which the forecasts are validated. The other lines - red for HB, green for PZM, the forecasts made by each respective model. Over the given time period, it is evident that all two models exhibit varying degrees of deviation from the validation data. For instance, the PZM model (Green line) demonstrates a higher deviation in later months, whereas the HB model (Red line) appears to track relatively closer to the actual data.

Table 3: Validation data comparison of forecasted data

Date	Validation	HB	PZM
Mar-2021	4.1	3.93	4.88
Apr-2021	3.9	3.84	4.37
May-2021	4.5	5.10	5.12
June-2021	5.2	5.30	4.44
Jul-2021	5.7	5.62	4.60
Aug-2021	6	4.94	3.02
Sep-2021	5.7	4.92	4.10
Oct-2021	7.6	7.30	3.74
Nov-2021	9.9	9.23	4.40
MAPE	-	7.10	27.17

This table presents a comparison of two forecasting models (HB, PZM) against the actual validation data for a period from March 2021 to November 2021. For each model, the Mean Absolute Percentage Error (MAPE) is also provided, which serves as an indicator of the model's accuracy.

The HB model has the lowest MAPE at 7.10, which indicates that on average, the model's predictions are approximately 7.10% off from the actual values. The PZM model's MAPE is 27.17, suggesting that its predictions deviate by about 27.17% from the actual values. These MAPE values can be useful for determining which model provides the most accurate predictions for the given data. Based on this table, the HB model appears to provide the most accurate forecasts, with the lowest average deviation from the actual validation data.




Figure 4: RMSE Change with forecast horizon -Validation data

The RMSE plot further substantiates the assessment of model performance by visualizing the variation of error as a function of the forecast horizon. As illustrated in the plot, the RMSE of all models tends to increase with forecast length, indicating a reduced predictive accuracy over time. The PZM model (Blue line) shows the highest increase, corroborating the comparatively higher MAPE in the earlier table. Conversely, the HB model (Green line) presents a more steady error increase, affirming its superior forecast precision. This graphical representation thus provides an intuitive means to compare forecast model robustness over an expanding forecast horizon.

6 Conclusion and Discussion

This study presented a comparative evaluation of two forecasting models (HB, PZM) using real-world validation data. The performance of these models was assessed by their Mean Absolute Percentage Error (MAPE), a commonly used metric for evaluating the accuracy of forecasting models. The results of the study revealed significant differences in the forecasting performances of the models.

The proposed model (HB) for this study, demonstrated the best performance with the lowest MAPE 7.10% meaning it had the smallest average deviation from the actual validation data. The other model, PZM showed larger deviations from the actual values, with MAPEs of 27.17% . This indicates that the HB model was more accurate and reliable in forecasting the data in this specific instance.

It is essential to note that while the HB model performed the best in this context, the optimal forecasting model can vary depending on the specific data set and the characteristics of the data. For instance, different models may perform better for data with different levels of noise, seasonality, or trend patterns. Therefore, the choice of the best forecasting model should always be informed by a thorough understanding of the data and the specific requirements of the forecasting task.

This study provides a good basis for further research in improving forecasting accuracy. Future work could explore the use of more advanced techniques, such as machine learning or deep learning models, or the combination of multiple forecasting models to enhance prediction performance. Moreover, further validation of the HB model using different data sets would also be valuable to confirm its robustness and generalizability.

References

- Box, G. E. and Jenkins, G. M. (1976). *Time series analysis: forecasting and control*. Holden-Day.
- Brock, W. A., Dechert, W. D., and LeBaron, B. (1996). The bds statistic as a test for the adequacy of a linear time series model. *Journal of Time Series Analysis*, 17:215–235.
- Cadenas, E. and Rivera, W. (2009). Wind speed forecasting in three different regions of mexico, using a hybrid arima-ann model. *Renewable Energy*, 34:2732–2738.
- Chollet, F. et al. (2018). Keras: The python deep learning library.
- Haykin, S. (1999). *Neural Networks: A Comprehensive Foundation*. Prentice Hall.
- Haykin, S. (2009). *Neural Networks and learning machines*. Pearson Education, 3 edition.
- Hyndman, R. J. and Athanasopoulos, G. (2014). *Forecasting: principles and practice*. OTexts.
- Mankiw, N. G. (2014). *Principles of Economics*. South-Western Cengage Learning.

-
- Mishkin, F. S. (2010). *The Economics of Money, Banking, and Financial Markets*. Pearson.
- of Sri Lanka, C. B. (2021). Central bank of sri lanka statistical information.
- Perera, N., Senadeera, M., and Sivakumaran, N. (2012). Stock market forecasting using arima-based neural networks: Evidence from sri lanka. *International Journal of Business and Economics Research*, 1:53–59.
- Team, R. C. (2022). R: A language and environment for statistical computing.
- Tong, H. (1990). *Non-linear time series: a dynamical system approach*. Oxford University Press.
- TradingView (2023). Economics-Ikiryy.
- Zhang, G., Patuwo, B. E., and Hu, M. Y. (2008). Forecasting with artificial neural networks: The state of the art. *International Journal of Forecasting*, 14:35–62.
- Zhang, G. P. (2003). Time series forecasting using a hybrid arima and neural network model. *Neurocomputing*, 50:159–175.