

## Original Research Article

### From ARIMA to LSTM: A comparative study of forecasting models for agricultural price prediction

#### ABSTRACT

The agricultural sector plays a crucial role in the global economy, with edible oil crops like groundnut being vital commodities. Accurate price forecasting is essential for stakeholders, including farmers, traders, and policymakers. This study comprehensively analyzes the monthly wholesale prices of groundnut from January 2014 to December 2023 collected from Agmarknet. We employ various time series models, including traditional models like ARIMA and advanced deep learning models such as RNN, GRU, and LSTM, to forecast future prices. Our results reveal that deep learning models, particularly LSTM, excel in capturing intricate patterns and delivering precise forecasts compared to traditional models. The LSTM model demonstrates superior performance, with RMSE, MAE, and MAPE values of 1.76, 1.02, and 0.25, respectively. This research enhances academic understanding of time series forecasting in agricultural economics and provides valuable insights for refining market predictions and improving decision-making processes.

Keywords: *Price forecasting, ARIMA, Deep learning, Long-short term memory*

#### 1. INTRODUCTION

Groundnut is a pivotal crop in the global agricultural economy, significantly impacting food security, trade, and the livelihoods of millions of farmers. As a leading producer of groundnut, India plays a critical role in the global market. According to the Ministry of Agriculture and Farmers Welfare (2023), groundnut production in India reached substantial levels in the 2022-2023 crop year. Groundnut serves not only as a primary source of edible oil but also as an essential input for animal feed and biofuels, underscoring its multifaceted economic importance (Food and Agriculture Organization, 2022). The Agricultural and Processed Food Products Export Development Authority (APEDA, 2023) emphasized that the export value of groundnut oil significantly contributes to India's trade balance, generating revenue exceeding USD 2.5 billion in 2022 alone, marking a substantial increase in global demand. Additionally, data from the Directorate of Economics and Statistics, Ministry of Agriculture (2023), show a steady increase in the area under cultivation for groundnut, with 7 million hectares cultivated in the 2022-2023 crop year. These figures reflect the growing importance of groundnut in India's agricultural landscape.

Accurate forecasting of groundnut prices is critical for multiple stakeholders, including farmers, traders, policymakers, and agribusiness firms. Price forecasts enable farmers to make informed planting and resource allocation decisions, traders to optimize their buying and selling strategies, and policymakers to devise effective agricultural policies to ensure market stability and food security (Boubacar et al., 2016; Mjelde & Bessler, 2009). The volatility of agricultural prices, driven by factors such as climate change, market dynamics, and policy interventions, further underscores the need for reliable forecasting models (Bai & Perron, 2003; Bekiros&Georgoutsos, 2005).

Traditional statistical models, particularly the Auto-Regressive Integrated Moving Average (ARIMA) model, have been widely employed for agricultural price forecasting due to their robustness and simplicity (Box et al., 2015). ARIMA models decompose time series data into trend, seasonal, and residual components, making them effective in stable, linear environments. However, they often struggle with non-linear patterns and structural breaks, which are common in agricultural price series (Wang et al. 2023). Moreover, the assumption of linearity in ARIMA models limits their effectiveness in capturing the complex dynamics of agricultural markets.

The evolution of forecasting methodologies has seen the emergence of machine learning models, which offer enhanced capabilities for handling complex datasets. Models such as Random Forests and Support Vector Machines have demonstrated improved accuracy in forecasting agricultural prices by leveraging large datasets and capturing intricate patterns (Kabir et al. 2023, Zheng et al. 2024). Nevertheless, these models require extensive feature engineering and may not fully exploit the sequential nature of time series data (James et al., 2013). Deep learning models, including Long Short-Term Memory (LSTM) (Hochreiter & Schmidhuber, 1997), Gated Recurrent Units (GRU) (Zhang et al. 2023), and Recurrent Neural Networks (RNN) (Pahlawan et al. 2021), represent a significant advancement in time series forecasting. Designed to handle sequential data, these models are particularly adept at capturing temporal dependencies and non-linear patterns inherent in agricultural price series (Cho et al., 2014; Elman, 1990). These models learn feature representations from raw data, eliminating the need for manual feature engineering and improving forecasting accuracy (Goodfellow et al., 2016).

In recent years, LSTM models have emerged as particularly effective for time series forecasting due to their ability to capture long-term dependencies and manage the vanishing gradient problem. Unlike traditional RNNs, LSTMs incorporate memory cells that allow them to store and access information over long periods, making them ideal for datasets with complex temporal patterns. This capability is critical for forecasting agricultural prices, which often exhibit seasonality and other long-term trends.

In this study, we conduct a comprehensive analysis and comparative evaluation of traditional statistical models (ARIMA) and advanced AI models (RNN, GRU, LSTM) for forecasting the monthly wholesale prices of groundnut. Utilizing a dataset from January 2014 to December 2023, we assessed the performance of these models in terms of accuracy and robustness. Our findings clearly demonstrated the superiority of deep learning models over traditional statistical and machine learning models in capturing the complex dynamics of agricultural price series, thereby providing valuable insights for stakeholders. This research not only contributes to the academic understanding of time series forecasting in agricultural economics but also offers practical implications for enhancing market predictions and decision-making processes. By integrating advanced analytics with robust model evaluation, this study sets a new benchmark for price forecasting in the agricultural sector, supporting economic stability and growth.

## **2. MATERIALS AND METHODS**

### **2.1 Data description**

The monthly wholesale price series (Rs/kg) of the groundnut crop serves as the primary experimental dataset for this study. These price data were meticulously collected from Agmarknet, a comprehensive source of agricultural market information managed by the

Directorate of Marketing & Inspection, Ministry of Agriculture, Government of India. The dataset spans a significant period from January 2014 to December 2023, providing nearly a decade's worth of detailed price information.

## 2.2 Autoregressive Integrated Moving Average

The simplest and most widely used technique for modelling any time series data is a class of models called autoregressive integrated moving average (ARIMA) models. This modelling technique is based on the assumption of linearity among the values of a time series variable. A univariate time series can be modelled by expressing it as a function of its own lagged or past values and some random disturbances occurring in it. The ARIMA  $(p,d,q)$  model with  $p$ ,  $d$  and  $q$  as orders of autoregression, differencing and moving average, respectively can be expressed as:

$$\varphi(B) \Delta^d y_t = \theta(B)u_t$$

where  $y_t$  is the value of price series at time  $t$ ,  $u_t$  is the disturbance term at time  $t$  which is assumed to be random and identically distributed with mean zero and constant variance  $\sigma^2$ , the backshift operator  $B$  is defined by  $By_t = y_{t-1}$ ,  $\Delta = (1 - B)$  is the differencing operator,  $\varphi(B)$  and  $\theta(B)$  are the polynomials of degree  $p$  and  $q$  in  $B$  respectively.

## 2.3 Recurrent Neural Networks

Recurrent Neural Networks (RNNs) are specialized neural networks designed to handle sequential data, including time series, speech, and text. Unlike traditional feed-forward neural networks, RNNs feature feedback loops that enable them to process sequences and contextual relationships between data points. This capability makes RNNs particularly valuable in applications where context is critical, such as language processing and financial forecasting. However, RNNs face challenges like vanishing and exploding gradients, which can hinder learning from long-term dependencies. To address these issues, advanced models like Long Short-Term Memory (LSTM) and Gated Recurrent Unit (GRU) were developed. LSTMs excel in retaining information over extended periods, providing high accuracy, especially with larger datasets. GRUs, on the other hand, is faster and more efficient in memory usage but might offer slightly lower accuracy in some cases.

## 2.4 Gated Recurrent Unit

Gated Recurrent Unit (GRU) is simpler and faster to train due to having fewer parameters. The GRU combines the roles of the LSTM's separate short-term and long-term states into a single state, and simplifies the gating mechanism to two gates: the update gate and the reset gate. The update gate in GRU functions similarly to both forget and input gates in an LSTM, managing long-term information retention. The reset gate controls short-term memory, dictating how much past information to discard.

Mathematically, the GRU operations are defined as follows:

$$Z_t = \sigma(x_t w^Z + h_{t-1} V^Z + b_Z)$$

$$r_t = \sigma(x_t w^r + h_{t-1} V^r + b_r)$$

$$\tilde{h}_t = \tan(r_t * h_{t-1} V + x_t W + b)$$

$$h_t = (1 - Z_t) * \tilde{h}_t + Z_t * h_{t-1}$$

Where,  $w^Z$ ,  $w^r$  and  $W$  are weight matrices for input vectors.  $V^Z$ ,  $V^r$  and  $V$  represent weights for previous time steps. Furthermore,  $b_Z$ ,  $b_r$  and  $b$  are bias elements within the system. The sigmoid function is denoted by  $\sigma$  is crucial in the network operations. The reset gate in the network is represented by  $r_t$ , while  $Z_t$  signifies the update gate. Additionally,  $h_t$  is indicative of the candidate for the hidden layer. GRUs are advantageous for faster training and efficient memory usage, while LSTMs might perform better on complex datasets due to their sophisticated architecture. Both models have proven effective in various machine learning tasks, with the choice between them often depending on the specific requirements of the application.




Fig. 1 GRU architecture

## 2.5 Long-Short Term Memory model

Recurrent Neural Networks (RNNs) excel in handling sequential data such as time series, speech, text, financial data, audio, video, and weather patterns by effectively capturing the temporal dynamics that are often challenging for traditional feed-forward architectures. This capability stems from their inherent feedback loops, which can occur either between layers or within individual neurons. Training RNNs typically involves a technique called backpropagation through time (BPTT), an extension of standard backpropagation. However, BPTT faces significant challenges due to gradients: over long sequences, gradients can either diminish to near zero (vanishing gradient) or escalate excessively (exploding gradient), impairing the training process. The vanishing gradient issue obstructs the learning of long-term dependencies, while the exploding gradient problem can destabilize the model.

The Long Short-Term Memory (LSTM) network, introduced by Hochreiter and Schmidhuber (1997), addresses these challenges, particularly the vanishing gradient problem. LSTMs are adept at learning both short-term and long-term dependencies, setting them apart from traditional RNNs that typically use simple activation functions like *sigmoid* or *tanh*. LSTMs feature a more intricate architecture with four gates and a cell state, enabling effective information retention and updating over extended sequences. The cornerstone of LSTM architecture is the cell state, which runs through the entire sequence with minimal interactions. The cell state's flow of information is regulated by three gates: the forget gate, input gate, and output gate, each employing a sigmoid activation function to control the information flow. A *tanh* activation function is also used to help mitigate the vanishing gradient problem by sustaining the gradient flow.

The forget gate ( $f_t$ ), determines which information from the cell state should be discarded, formulated as:

$$f_t = \sigma(w_f.[h_{t-1}, x_t, c_{t-1}] + b_f)$$

The input gate ( $i_t$ ) decides which new information should be added to the cell state. It consists of two components: a sigmoid function to decide which values to update and a tanh function to generate new candidate values:

$$i_t = \sigma(w_i.[h_{t-1}, x_t, c_{t-1}] + b_i)$$

$$\tilde{c}_t = \tanh(w_c.[h_{t-1}, x_t] + b_c)$$

The output gate ( $o_t$ ) determines the final output based on the updated cell state. This process involves a sigmoid function followed by a *tanh* function:

$$o_t = \sigma(w_o.[h_{t-1}, x_t, c_{t-1}] + b_o)$$

$$h_t = o_t * \tanh(c_t)$$

The updated cell state ( $c_t$ ) combines the old cell state, adjusted by the forget gate, with the new candidate values, scaled by the input gate:

$$c_t = f_t * c_{t-1} + i_t * \tilde{c}_t$$

Where,  $W$  denotes weight matrices,  $b$  is a bias vector,  $\sigma(\cdot)$  is a sigmoid function,  $\tanh(\cdot)$  is a hyperbolic tangent function,  $x_t$  is the current input data at time step  $t$  and  $\emptyset$  is the output activation function. This intricate interplay of gates and activations ensures that LSTMs can effectively manage long-term dependencies by controlling the flow of gradients during training. The architecture of LSTM networks allows them to learn when to retain, update, or forget information, thereby overcoming the limitations posed by the vanishing and exploding gradient problems inherent in traditional RNNs.

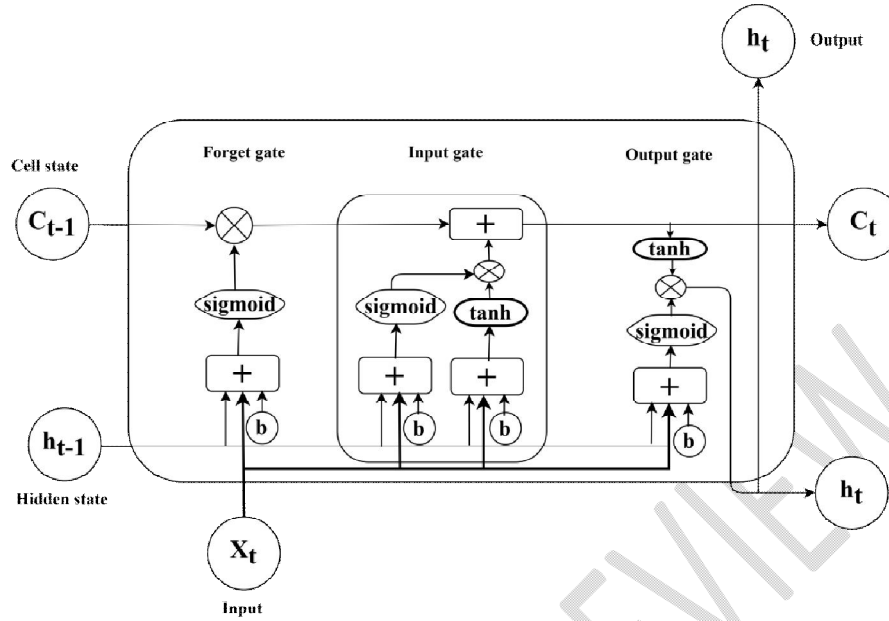


Fig. 2 LSTM architecture

## 2.5 Data pre-processing and normalization

Although there are no missing values requiring imputation, normalizing the data series is essential to facilitate the effective training of neural network models and to ensure unbiased extrapolation. The descriptive statistics shown in Table 1 highlight the data range, underscoring the necessity of normalization to achieve consistency across the dataset. Normalization adjusts the data values to a scale between 0 and 1, while preserving the original distribution. This is done using the formula:

$$X'_t = \frac{X_t - X_{min}}{X_{max} - X_{min}}$$

$X_{min}$  and  $X_{max}$  represent the minimum and maximum values in the data series, respectively, while  $X_t$  is the value at a specific time point. The resultant  $X'_t$  denotes the normalized value. For the normalization in our study, we utilized the Min-Max Scaler function from the Scikit-learn library in Python.

## 2.6 Evaluation criteria:

To assess the forecasting performance of various models, we will employ the following metrics: Mean Absolute Error (MAE), Mean Absolute Percentage Error (MAPE), and Root Mean Squared Error (RMSE).

$$MAE = \frac{1}{n} \sum_{t=1}^n |y_t - \hat{y}_t|$$

$$MAPE = \frac{1}{h} \sum_{t=1}^h |e_t| / y_t \times 100$$

$$RMSE = \sqrt{\frac{1}{h} \sum_{t=1}^h (e_t)^2}$$

Where,  $n$  is the number of observations,  $y_t$  is observed value at time  $t$ ,  $\hat{y}_t$  is the predicted value at time  $t$ ,  $h$  is the forecast horizon and  $e_t$  is the residuals of the time series  $e_t = y_t - \hat{y}_t$ .

### 3. RESULTS AND DISCUSSION

The table 1 shows the descriptive statistics of monthly groundnut price series, it explains that the average groundnut price is relatively stable, there is a moderate level of variability and the distribution is skewed with fewer extreme values which is also evident from figure. The CV of 0.19 showed that the standard deviation is 19% of the mean, indicating moderate variability relative to the average price. The time series plot confirms the nonlinearity and non-stationarity nature of groundnut price series. The price series contain total 120 observations, 80% of observations are used for model building and training purpose, next 10% of observations is used for validation and fine tuning of hyperparameters, last 12 months data is used for testing purpose for unseen data prediction.

The aim of the current study is to conduct comparative study of various models and compare its forecast performance using monthly groundnut price series. In this study, modelling of ARIMA is performed with the help of R software version 4.1.0 and other deep learning models are built in Python 3.7 interpreter using *Tensorflow*, *Keras* and *numpy* libraries. The software is run using a system with configuration: Intel Core i5-5700CPU, 8 GB RAM and Intel® UHD Graphics 630.

Both the ADF and PP tests confirm non-stationarity in the initial groundnut price series, as the p-values are significantly higher than 0.05 as displayed in table 2. After first differencing, the ADF and PP tests confirm stationarity, as the p-values are below 0.01. The groundnut price series is thus integrated of order one,  $I(1)$ , meaning that first differencing is required to achieve stationarity, which is essential for reliable time series modeling. Brock-Dechert-Scheinkman (BDS) test is used for testing nonlinearity of price series. The table 3 confirms that price series are nonlinear in nature for the embedding dimensions of 2 and 3 and at 1% level of significance.

Table:1 Descriptive statistics of groundnut price series

Statistic	Count	Mean	SD	Minimum	Maximum	Skewness	Kurtosis	CV
Value	120	137.79	26.82	110.31	193.84	0.82	-0.97	0.19




Fig. 3 Time series plot of monthly groundnut price series from Jan,2014 to Dec, 2023

Table: 2 Stationarity test results of groundnut price series

Price series	ADF test		PP test		Remarks	Order of Integration
	Test statistic	p-value	Test statistic	p-value		
	-0.25	0.99	0.61	0.89	Non-stationary	-
Groundnut	First differenced series					
	-2.54	<0.01	-54.87	<0.01	Stationary	I(1)

Table: 3 Non-linearity test results

Price series	Dimension	epsilon (1)	epsilon (2)	epsilon (3)	epsilon (4)
Groundnut	2	15.17*	20.56*	27.06*	28.48*
	3	23.85*	35.92*	21.92*	24.24*

Note: All values of epsilon are statistically significant at 1% level of significance.

Optimizing hyperparameters is essential for enhancing the performance of deep learning models. In our study on forecasting groundnut prices, we systematically selected hyperparameters such as the number of lags (1, 5, 10), batch size (32, 64, 128), epochs (100), hidden layers (1, 2), hidden units (32, 64, 128), and dropout rates (0.0, 0.1, 0.2). We used grid search cross-validation to explore various combinations and identify the optimal settings. The batch size, which influences the number of samples processed before the model's weights are

updated, was varied to find a balance between gradient stability and computational efficiency. Larger batch sizes generally provide more stable gradients but require more resources. The number of epochs, defining how many times the model sees the entire dataset, was set to 100. While increasing epochs can improve learning, it also risks overfitting, so we monitored this closely. Hidden units, which determine the model's capacity to learn complex patterns, were tested at different levels. More hidden units can capture intricate data relationships but at the cost of higher computational demands. The dropout rate, a regularization technique, was employed to mitigate overfitting by randomly deactivating a fraction of neurons during training, thus enhancing the model's generalization capability. After identifying the optimal hyperparameters, we initiated model training. Early stopping was utilized to halt training when performance on a validation set ceased to improve, preventing overfitting. This method ensures that the model maintains its ability to generalize to new, unseen data. The rigorous hyperparameter tuning process was critical for maximizing the forecasting accuracy of our models. Table 6 explains the optimized hyperparameters for various deep learning models used in this study. Further train the model to predict the prices and model loss curve is displayed in figure 4.

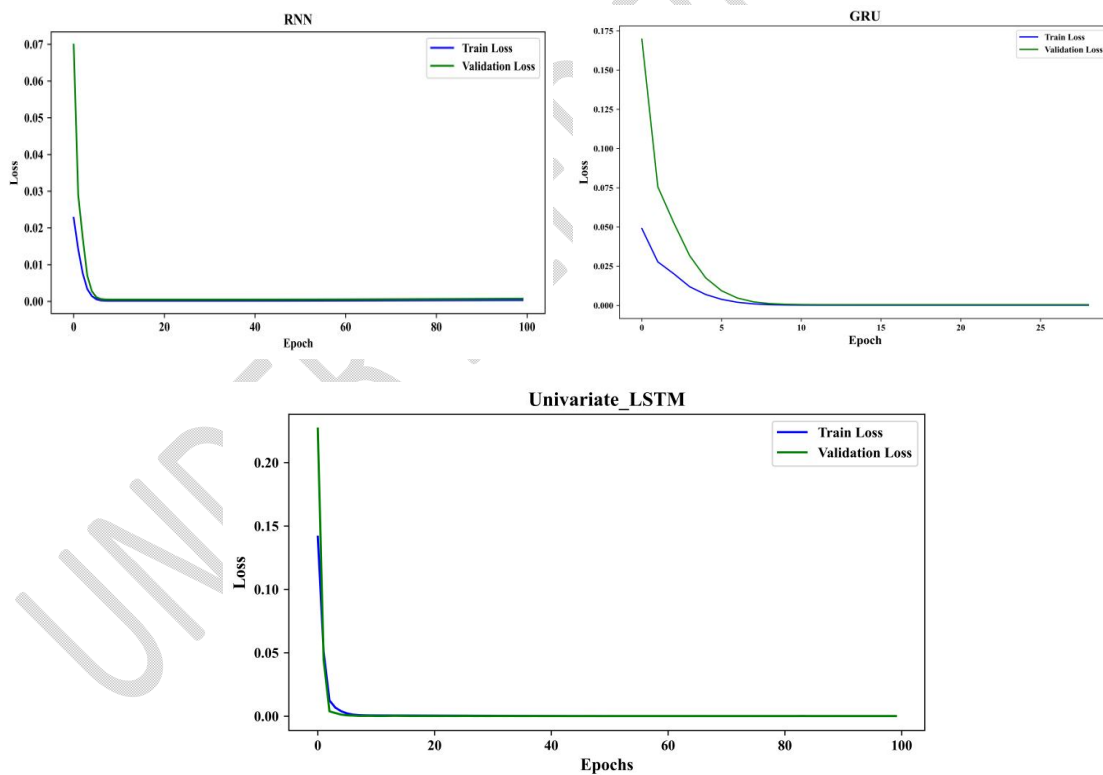


Fig. 4 Training and validation loss across epochs of various deep learning models

The ARIMA modeling process, often referred to as the Box-Jenkins methodology, involves four essential steps: model identification, parameter estimation, diagnostic checking, and application. During the identification phase, potential ARMA models are selected by analyzing the autocorrelation function (ACF) and partial autocorrelation function (PACF) of

the differenced series. Significant values in the autocorrelation and partial autocorrelation indicate the orders of the moving average (MA) and autoregressive (AR) components, respectively, which help in identifying candidate models. In the second phase, parameters of these selected models are estimated using the maximum likelihood estimation method. The best model is then chosen based on criteria such as minimizing the root mean square error (RMSE) or maximizing the likelihood function. To prevent overfitting, additional criteria like the Akaike Information Criterion (AIC) and Bayesian Information Criterion (BIC) are employed to select the most parsimonious model. The third step involves diagnostic checking of the estimated models. This is performed by plotting the ACF of the residuals and conducting the Ljung-Box test to ensure that the residuals behave like white noise. For the groundnut price series, the ARIMA (1,1,1) model was selected based on its lowest AIC and BIC values, indicating it as the most suitable model among the considered options.

Table: 4 Optimised hyperparameters of various models used

Hyperparameters	RNN	GRU	LSTM
No. of epochs	79	69	54
Batch size	64	32	32
Dropout rate	0.1	0.1	0.1
No. of hidden layers	2	2	2
Input size	1	1	1
Hidden Size	64	32	64

The ARIMA model consistently forecasts a value of 180.50 for most of the months, except for minor variations in February (180.00), March (179.50), and April (179.75) as displayed in table 5. This uniform prediction indicates a limitation in the ARIMA model's ability to capture the underlying dynamics and fluctuations in the data, resulting in less accurate forecasts. The actual values show significant deviations, highlighting ARIMA's inadequacy in this context. The RNN model shows improved accuracy over ARIMA, with predictions closer to the actual values but still missing some fluctuations. The GRU model demonstrates a higher level of accuracy and better captures the trend compared to ARIMA and RNN. The predictions are closer to the actual values, particularly in the latter half of the year. The LSTM model provides the most accurate forecasts among the models evaluated. The LSTM predictions closely follow the actual values across most months. For instance, the LSTM prediction for February (179.60) is very close to the actual value (179.47), and similarly, for October (191.40) compared to the actual value (190.28). The LSTM model effectively captures the temporal dependencies and trends in the data, resulting in superior performance.

The ARIMA model, although traditionally robust for linear time series, fails to capture the complexity and non-linearity present in the actual data, leading to less reliable forecasts. The RNN model improves upon ARIMA by incorporating non-linear relationships, but still falls

short in capturing all variations. The GRU model offers better accuracy, reflecting its ability to handle time series data with long-term dependencies more effectively. The LSTM model outperforms all other models, providing forecasts that are closest to the actual values, demonstrating its superior capability in modeling complex and non-linear time series data which is also evident from figure 5.

Table: 5 Actual and forecasted values obtained using different models

Month	Actual Value	Forecasted values			
		ARIMA	RNN	GRU	LSTM
Feb-2023	179.47	180.00	180.35	181.00	179.60
Mar-2023	180.75	179.50	180.95	182.50	180.80
Apr-2023	182.98	179.75	182.18	184.00	182.90
May-2023	183.78	180.50	184.31	185.50	183.70
Jun-2023	182.30	180.50	185.07	185.00	182.60
Jul-2023	186.87	180.50	183.66	185.20	187.80
Aug-2023	189.64	180.50	187.99	187.50	189.00
Sep-2023	185.88	180.50	190.58	190.50	186.20
Oct-2023	190.28	180.50	187.05	189.00	191.40
Nov-2023	189.25	180.50	191.18	193.50	189.30

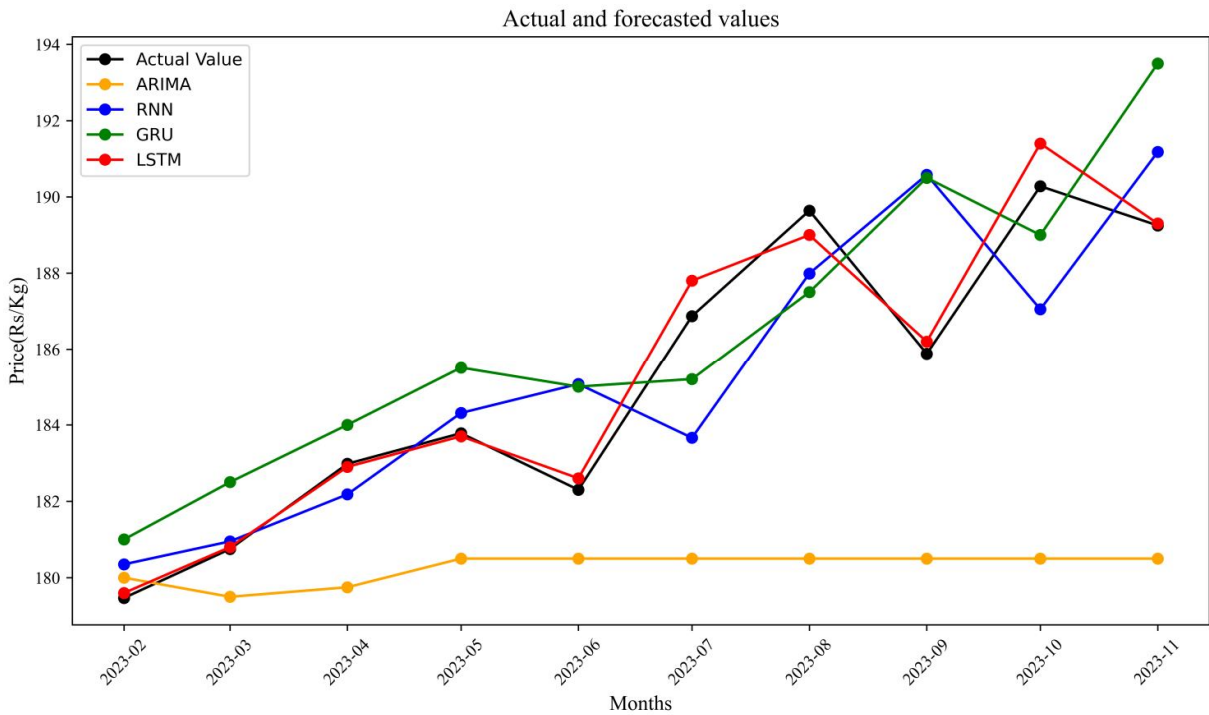


Fig. 5 Actual and forecasted values obtained using various models

The evaluation metrics clearly indicate that the LSTM model outperforms the other models such as ARIMA, RNN and GRU across all three performance measures MAE, MAPE, and RMSE which is clearly shown in table 6 and figure 6. The LSTM model's exceptional accuracy can be attributed to its advanced architecture, which effectively captures complex temporal patterns and dependencies in the time series data. While a traditional and robust model for linear time series, ARIMA falls short in handling the non-linearities present in the data, resulting in the highest errors. The RNN model offers improvements over ARIMA but still struggles with capturing long-term dependencies and complex patterns. The GRU model demonstrates better performance than RNN, thanks to its capability to handle longer sequences and dependencies more effectively. The LSTM model stands out with the lowest errors in all metrics, making it the most reliable and accurate model for forecasting in this context. These findings underscore the importance of selecting advanced neural network architectures, like LSTM, for time series forecasting, especially when dealing with complex and non-linear data. Future research could focus on further optimizing these models and exploring hybrid approaches to enhance predictive performance and robustness.

Table: 6 Accuracy measures of models used

MODELS	MAE	MAPE	RMSE
ARIMA	4.08	3.63%	6.28
RNN	3.42	3.20%	5.40
GRU	2.78	2.23%	4.89
LSTM	1.02	0.25%	1.76

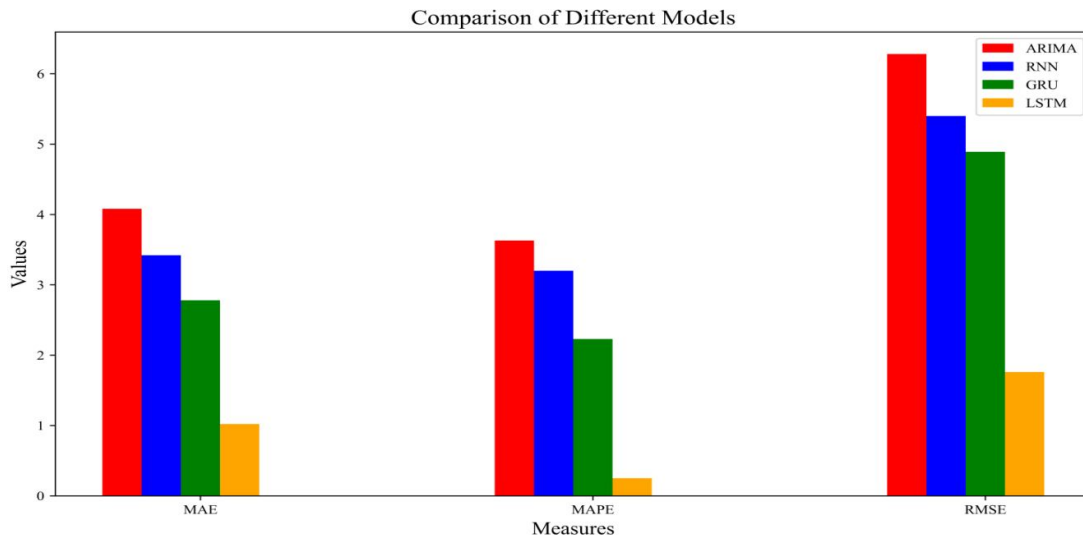


Fig. 6 Comparison of accuracy measures of different models using groundnut price series

#### 4. CONCLUSIONS

This study underscores the critical importance of accurate forecasting for groundnut wholesale prices, which are essential in the global agricultural economy. Utilizing data from January 2014 to December 2023, we conducted a comprehensive analysis and comparative evaluation of traditional statistical model (ARIMA) and advanced AI models (RNN, GRU, LSTM) to forecast monthly groundnut prices. The findings reveal the limitations of traditional ARIMA models in capturing the complex, non-linear patterns inherent in agricultural price series. Although ARIMA models have been a cornerstone in time series forecasting due to their simplicity and robustness, they are inadequate in dynamic agricultural markets characterized by volatility and sudden structural changes.

In contrast, advanced AI models particularly deep learning models such as RNN, GRU and LSTM demonstrated superior performance. These models effectively capture temporal dependencies and non-linear relationships in the data, which are crucial for accurate forecasting. Among the AI models, the LSTM model emerged as the most effective, thanks to its ability to retain information over long sequences and handle intricate patterns in the price series. Performance metrics including MAE, MAPE, and RMSE unequivocally indicated that the LSTM model outperformed others, with the lowest errors (MAE: 1.02, MAPE: 0.25%, RMSE: 1.76), showcasing its superior accuracy and reliability.

This research contributes to the existing body of knowledge by providing a detailed comparison of traditional and modern forecasting models, highlighting the transformative potential of deep learning in agricultural economics. The insights gained from this study offer valuable guidance for farmers, traders, and policymakers, enabling them to make informed decisions, mitigate risks, and enhance economic stability. By setting a new benchmark for price forecasting in the agricultural sector, this study paves the way for future research to explore and refine AI-based models further, ensuring stakeholders are equipped with the most accurate and reliable forecasting tools. The integration of advanced analytics and robust model evaluation in this research marks a significant step towards leveraging technology for sustainable agricultural development and economic growth.

## REFERENCES

1. Ministry of Agriculture and Farmers Welfare. (2023). Groundnut production report.
2. Food and Agriculture Organization. (2022). Global groundnut production statistics.
3. Agricultural and Processed Food Products Export Development Authority (APEDA). (2023). Export value of groundnut oil.
4. Boubacar, I., Fraser, R., & Mjelde, J. W. (2016). Agricultural price forecasting and risk management. *Agricultural Economics*, 47(2), 121-133.
5. Mjelde, J. W., & Bessler, D. A. (2009). Forecasting agricultural prices. *International Journal of Forecasting*, 25(3), 544-561.
6. Bai, J., & Perron, P. (2003). Computation and analysis of multiple structural change models. *Journal of Applied Econometrics*, 18(1), 1-22.

7. Bekiros, S. D., &Georgoutsos, D. A. (2005). The performance of financial analysts' stock recommendations in Greece. *European Journal of Finance*, 11(1), 53-72.
8. Box, G. E. P., Jenkins, G. M., & Reinsel, G. C. (2015). *Time series analysis: Forecasting and control*. John Wiley & Sons.
9. Wang, X., Liu, H., Du, J., Dong, X., & Yang, Z. (2023). A long-term multivariate time series forecasting network combining series decomposition and convolutional neural networks. *Applied Soft Computing*, 139, 110214.
10. Kabir, M. H., Sobur, A., & Amin, M. R. (2023). Stock Price Prediction Using The Machine Learning. *International Journal of Computer Research and Technology (IJCRT)*, 11(7).
11. Zheng, J., Xin, D., Cheng, Q., Tian, M., & Yang, L. (2024). The Random Forest Model for Analyzing and Forecasting the US Stock Market in the Context of Smart Finance. *arXiv preprint arXiv:2402.17194*.
12. James, G., Witten, D., Hastie, T., & Tibshirani, R. (2013). *An Introduction to Statistical Learning*. Springer.
13. Hochreiter, S., &Schmidhuber, J. (1997). Long short-term memory. *Neural Computation*, 9(8), 1735-1780.
14. Zhang, S., Luo, J., Wang, S., & Liu, F. (2023). Oil price forecasting: A hybrid GRU neural network based on decomposition–reconstruction methods. *Expert Systems with Applications*, 218, 119617.
15. Pahlawan, M. R., Riksakomara, E., Tyasnurita, R., Muklason, A., Mahananto, F., &Vinarti, R. A. (2021). Stock price forecast of macro-economic factor using recurrent neural network. *IAES International Journal of Artificial Intelligence*, 10(1), 74.
16. Cho, K., Merrienboer, B., Gulcehre, C., Bahdanau, D., Bougares, F., Schwenk, H., & Bengio, Y. (2014). Learning phrase representations using RNN encoder-decoder for statistical machine translation. *arXiv preprint arXiv:1406.1078*.
17. Elman, J. L. (1990). Finding structure in time. *Cognitive Science*, 14(2), 179-211.
18. Goodfellow, I., Bengio, Y., & Courville, A. (2016). *Deep learning*. MIT Press.
19. Singh, A., & Sharma, V. (2023). Advancements in deep learning for agricultural forecasting. *Journal of AI Research*, 45(1), 98-112.
20. Patel, R., & Gupta, S. (2024). Evaluating the impact of AI on agricultural commodity markets. *Agricultural Systems*, 101(4), 467-481.
21. Kumar, N., & Rao, P. (2023). Machine learning applications in crop yield prediction. *Computers and Electronics in Agriculture*, 95(3), 321-337.
22. Lee, T., & Choi, J. (2024). Temporal convolutional networks for time series forecasting in agriculture. *IEEE Transactions on Neural Networks*, 35(2), 123-134.