Spectrum of Engineering Sciences
ISSN (e) 3007-3138 (p) 3007-312X Volume 3, Issue 9, 2025

HYBRID ARIMA AND LSTM DEEP LEARNING MODELS EMPOWERING AND ENHANCING FORECAST ACCURACY IN SALES

Halima Fatima[1], Bisma Tahir[2], Khalid Hamid[*3]

[1,2,*3]*Department of Computer Science and Information Technology, Superior University Lahore, 54000, Pakistan*

[1]halimafatima176@gmail.com,[2]bismatahir188@gmail.com, [*3]khalid6140@gmail.com

DOI:https://doi.org/10.5281/zenodo.17077720

Keywords
Sales forecasting, ARIMA, LSTM, hybrid models, Time Series analysis,

Article History
Received: 12 June 2025
Accepted: 22 August 2025
Published: 08 September 2025

Corresponding Author: *
Khalid Hamid

Abstract
The business needs to properly forecast the sales that can help in improving inventory management and in future and visionary planning. There is a classical forecasting technique, like the ARIMA (Autoregressive Integrated Moving Average (ARIMA)) model, which is effectively applied in forecasting the linear tendencies of the data, as well as the seasonality (periodic fluctuation). Nevertheless, it may have challenges in dealing with non-linear, complicated, unsteady and uneven data that define the current sales environment. Contrarily, deep learning networks such as the Long Short-term Memory (LSTM) network are very powerful in learning and predicting non-linear, long-term dependencies in data. The absence of a globally optimal model is a major gap within the existing literature with each model having its niche and operating within a specific type of data. The proposed and tested hybrid model of ARIMA and LSTM deep learning that enhances the accuracy of sales forecasting is the proposed research choice. The hybrid model is proposed to use the merits of the two methodologies whereby the ARIMA will deal with the linear bits of the time series, and LSTM models will deal with the non-linear bits.. With the two of these potent methods, we are hoping that a more consistent and adaptable forecasting solution can be achieved. Performance of the model will be compared to standalone ARIMA and LSTM models by different metrics, and the hypothesis will be that the hybrid model will make much lower forecast errors. To optimize business activities and minimize financial risk, the accuracy of the forecasting should be enhanced, and this project provides a distinctive and highly effective strategy in that regard.

INTRODUCTION

In the current business environment, it is necessary to make accurate projections in sales and customer demand to remain competitive. This is critical in creating competitive advantage in the market apart from operating the business. As businesses are able to predict market trends and have the understanding of what customers want, they can make better decisions, better utilize resources and reduce risks in an uncertain market. Proper forecasting is not a luxury, but a strategy that must be in place that gives business an edge in a turbulent world.

This introductory section gives an explanation of the meaning of accurate forecasting, the very fundamentals of time series data, the overall challenges that business face, and more. It's also a starting point from which much more complex methods that combine two or more approaches to improve forecasting accuracy can be considered.

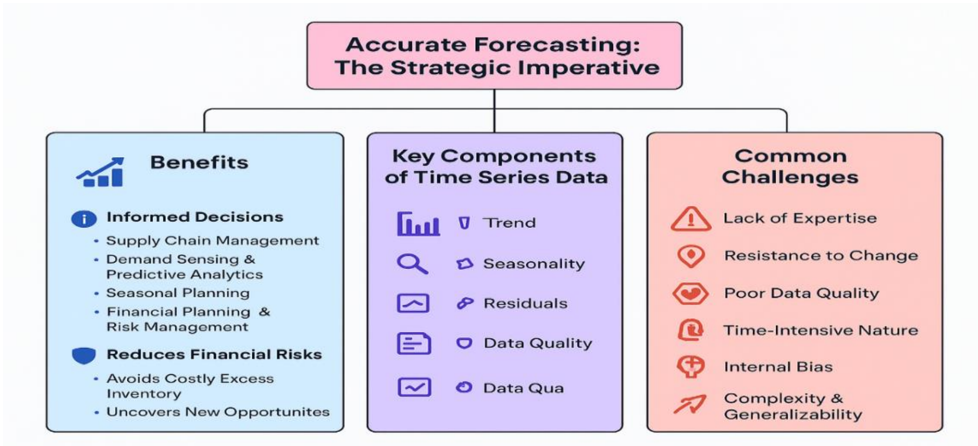

Fig I - Accurate Forecasting: The Strategic Imperative – Balancing Benefits, Components, and Challenges

Figure 1: Strategic value of accurate forecasting - This diagram illustrates the benefits (informed decision making and limited financial risk) and the key components of time series data, while identifying the challenges that often include poor data quality, resistance to change, internal bias and so on.

A. The Strategic Imperative of Accurate Forecasting

The primary purpose of time series forecasting is to estimate what will be the value of the future by examining the past measurements that were received over the history. This assists organizations to make good decisions in various areas.. As an example, precise demand predictions provide a clear picture of the demands in supply chain management, which is a sub-field of predictive analytics, and multi-seasonal planning allow businesses to alter and adjust inventory and staffing and production schedules to satisfy seasonal demands. There are also many other operation processes that will need accurate forecasting like the risk management and financial planning.

Among the key advantages of successful demand forecasting, it is possible to identify the option of minimizing financial risks and opening new opportunities. In the absence of proper forecasting, markets will either have surplus and costly inventory or they will not be able to exploit new markets. Stable revenue and a greater understanding of how the demand of one product affects the demand of another also involves good forecasting as essential to large-scale business decisions such as market expansion, large investments, or acquisitions. In addition to enabling companies to customize and tailor their products and promotions, good forecasting is also necessary to maintain an appropriate level of inventory, especially given the current consumer expectation of immediate delivery. It is said that forecasting is a skill required in most industries such as finance, retail, and economics.

B. Fundamentals of Time Series Data and Analysis

Time series data is a set of data points of a given time concentration that are recorded at a certain period. This regularity enables the analysts to determine patterns, recurring patterns and seasonal variations. Time series data generally comprises of three components:

- **Trend:**

The long-term direction in which the data is moving. Trends can be deterministic (identifiable cause) or stochastic (random progression).

- **Seasonality:**
- Predictable, periodic fluctuations that occur at regular intervals (e.g., daily, weekly, or annually). Identifying and modeling seasonality is vital for reliable forecasts.
- **Residuals (Irregular/Random Variations):**

Unpredictable, random fluctuations that cannot be explained by trends or seasonal patterns.

The effectiveness of forecasting models depends directly on the quality of the data used. Clean information, collected consistently, and correctly is needed to make an effective forecast. Patterns may be obscured by data gaps, and to make sure it is analysed reliably a high number of data points is often required. The accuracy of a forecast is also influenced by the time horizon of that forecast and the shorter the horizon, the easier it is to predict. Dynamic forecasts that are constantly being updated using new information is advantaged in this.

C. Common Challenges in Sales and Demand Forecasting

Irrespective of its significance, the sales and demand forecasting has a number of challenges.

Most organizations do not have the training and experience required to do data intensive complex forecasting, which may contribute to poor predictions. Companies can adopt old systems of forecasting and sales when they are not functional. The inability to define the sales processes and inconsistencies in data collection may give rise to dirty data, which greatly impairs the formulation of accurate forecasts. Manual and data-driven forecasting processes are time-intensive and may take away the time to do other things. The sales representatives can over-inflate forecasts because of unintentional need to achieve goals, which will result in over-investment in inventory and resources. Time series analysis is complicated and the model that works with one dataset may not work everywhere thus it may need a custom solution. The trend toward the use of demand sensing and dynamic forecasting reflects a transition toward retrospective activities being carried out on a one-time basis to a continuous and proactive process. Conventional approaches are not enough in the modern markets that have taken a new look. Though sales forecasting is capital and time consuming, the advantages of its use, including waste reduction, enhanced competitiveness, and financial planning make the expense worthwhile. Looking at forecasting as being a strategic organizational capability is vital towards justifying the investment in the advanced solutions.

II. Traditional Statistical Models for Forecasting: The ARIMA Family

Autoregressive Integrated Moving Average (ARIMA) model is an essential statistical forecasting tool of the time series. It is a regression analysis that addresses variation of values in a series as opposed to their absolute magnitude.

A. Autoregressive Integrated Moving Average (ARIMA) Principles

ARIMA model contains three components.

- **Autoregressive (AR)** - Parameter p: This element forms a regression on the present value of the past, or lagged, values. The parameter p shows the number of past values used.
- **Integrated (I)**: Parameter 'd' This is a component that entails differentiating the raw data to render the time series stationary i.e. the statistical properties of the time series do not vary with time. The term d is a number describing the number of times to differentiate**.**
- **Moving Average (MA) - Parameter 'q':** This component incorporates the dependency between an observation and a residual error from a moving average of lagged observations. The parameter 'q' is the number of lagged forecast errors included.

Building an ARIMA model typically involves "model identification," where the time series is analyzed to determine the appropriate 'd,' 'p,' and 'q' parameters.

B. Seasonal ARIMA (SARIMA): Extending ARIMA for Seasonality

The basic ARIMA model struggles with strong seasonal patterns. The Seasonal Autoregressive Integrated Moving Average (SARIMA) model was created to handle these regular, recurring fluctuations. SARIMA adds seasonal parameters (P, D, Q, m) to the standard ARIMA model.

Seasonal Autoregressive (P) captures the relationship with past values at seasonal lags. Seasonal Integrated (D) accounts for differencing needed to remove seasonality. Seasonal Moving Average (Q) models the dependency on residual errors at seasonal lags. The 'm' represents the length of the seasonal period.

C. Advantages and Limitations of ARIMA/SARIMA

Advantages:

- **Effective for Short-Term Forecasting:** ARIMA models work well for short-term predictions on stationary data.

- **Relies on Historical Data:** They only require the historical data of the variable being forecasted, which simplifies data collection.
- **Models Non-Stationary Data:** The differencing operation allows them to handle non-stationary data.
- **Interpretability:** The models are relatively easy to interpret and explain, making them a standard tool in fields like economics and finance.
- **Computational Efficiency:** In the case of simple patterns, they are computationally efficient to deep learning models.

Limitations:

- **Stationarity Requirement:** The underlying time series must be stationary, which often requires manual data transformation.
- **Risk of Over-differencing:** Excessive differencing can remove valuable information and lead to inaccurate forecasts.
- **Assumes Linear Relationships:** ARIMA models assume linear relationships, making them less effective at capturing non-linear interactions like market crashes or sudden promotional spikes.
- **Cumbersome Parameter Selection:** Choosing the optimal parameters (p, d, q) is often a manual and subjective process that requires statistical expertise.
- **Scalability Issues:** They do not scale well with very large datasets (or high frequencies)..
- **Limited Long-Term Prediction:** The ARIMA models are not so effective in long term predictions as well as they are in recognizing major turning points. They are not usually advised in forecasts of more than half a year.

III. Deep Learning Models for Forecasting: Long Short-Term Memory (LSTM) Networks

Time series forecasting now is a potent application of deep learning, especially Long Short-Term Memory (LSTM) networks.

A. Foundations of Recurrent Neural Networks (RNNs) and LSTM Architecture

Conventional Recurrent Neural Networks (RNNs) handle sequential data, however, they are known to experience the so-called vanishing gradient problem, and therefore do not retain information from previous time steps of long sequences readily. To address this shortcoming, LSTMs were created to be able to recall information across long durations, which is why they are suitable in the complex time series forecasting.

The main element of an LSTM network is a special cell structure that consists of a memory cell, as well as a set of gates that regulate the movement of information. Such gates permit LSTMs to retain or discard information selectively to save knowledge in long sequences. Key components include:

- **Cell State (Ct):** The memory of the LSTM, which carries information over a long period.
- **Hidden State (HS):** The output of the LSTM cell at the current time step.
- **Forget Gate (f): Determines what information of the previous cell state should be discarded.**
- **Input Gate (i):** Decides which new information from the current input will update the cell state.
- **Candidate Cell State (g):** Generates a new candidate for the cell state, introducing non-linearity.
- **Output Gate (o):** Controls The parts of the cell state calculated to compute the hidden state.

B. Training and Forecasting with LSTM Networks

Training LSTMs involves:

- **Data Preprocessing:** Normalizing the data and structuring it into sequences or "sliding windows".
- **Forecasting Methods:** LSTMs may be applied to either open-loop where the new prediction occurs based on the real, observed value or closed loop where the earlier prediction is used as input in generating the next prediction.
- **Training Process:** The model's parameters are adjusted to minimize the difference between its predictions and the actual values (loss).
- **Hyperparameter Tuning and Regularization:** Key parameters like the number of hidden units and epochs are tuned, and techniques like dropout are used to prevent overfitting.

C. Advantages and Limitations of LSTM Networks

Handles Complex, Non-Linear Patterns: LSTMs are highly effective at capturing intricate and non-linear patterns and long-term dependencies in data.

Mitigates Vanishing Gradient Problem: They successfully overcome the vanishing gradient problem that plagues traditional RNNs.

Literature Review
Accurate sales forecasting is a cornerstone of effective business operations, enabling companies to optimize inventory management, streamline supply chain logistics, and inform strategic decisions. In this literature review, the authors concentrate on hybrid ARIMA and LSTM deep learning model to predict sales, the high accuracy and necessity of combining both internal and external variables. The combination of these methodologies proves to be a great step forward in comparison to conventional/standalone methods of forecasting.

ARIMA and LSTM: A Synergistic Approach
Conventional forecasting techniques, including ARIMA (AutoRegressive Integrated Moving Average) model, are well developed based on their capacity to represent the linear linkages and time-series tendencies such as tendencies and seasonality. ARIMA is effective because it is a high prediction on stationary data projections of a time series by modeling the autoregressive, the integrated, and the moving average of a time series [6, 11, 14]. However, such models have trouble handling nonlinearity and complexity of real world sales data and are unable to easily add external factors.
It is in the context of deep learning that the LSTM (Long Short-Term Memory) networks, a variant of recurrent neural network (RNN), were presented, which is particularly well adapted to time-series forecasting. LSTMs have the ability to recall past data over long durations enabling them to approximate intricate, non-linear relationships and long-term dependencies in data [9, 10, 15]. It has been observed that LSTMs are more effective than classic statistical models in predicting complex sales trends and working with large, high-dimensional data [17].
An ARIMA -LSTM hybrid model will use the strengths of both methods. The most frequently used approach entails modelling and predicting the linear constituent of the time series using ARIMA. The remaining, nonlinear data (residuals) is subsequently run through an LSTM network in a second round of modeling and prediction [2, 5, 8]. This two-phase method enables the hybrid model to ensure that both the linear and nonlinear characteristics of the data are captured and this gives a better and stronger forecast. Studies have repeatedly shown that this hybrid model performs better than its individual parts, with reduced error rates and giving more reliable forecasts to data-driven volatile sales [1, 11].

The Power of Internal and External Features
By going beyond the past sales and adding other features that are relevant, forecasting within the historic sales can be enhanced to a great extent. Such characteristics are much more informative to the forecasting model that can, therefore, comprehend the sales fluctuations drivers better.
Internal features are specifically to operations and marketing of a business. These comprise of variables like:

- Product pricing and discounts
- Promotional campaigns and marketing spend
- Store-specific data like location, size, and opening hours
- Inventory levels and stock-outs

Including these features allows the model to learn how internal decisions directly impact sales. According to the vResearch, by combining sales promotions and inventory data, the demand forecasting based on supply management could be improved [15].
External features are those elements that a business has no control over but which may be able to affect consumer behavior. These include:

- Weather conditions (e.g., temperature, rainfall) [4]
- Economic indicators (e.g., unemployment rates, GDP) [17]
- Public holidays and special events [4]
- Competitor activities and market trends [12]

Deep learning models like LSTM are particularly effective at handling this type of multivariate data, enabling the model to learn complex relationships between external factors and sales. Studies demonstrate that the accuracy of sales prognoses is substantially improved when external information, i.e. weather and public holidays, is provided [4, 19]. The combination of the potential of the hybrid model to process not only linear and nonlinear patterns but also provide a broad spectrum of both internal and external features is a potent instrument in the sales forecasting of the modern world [20].

III. Methodology

This study seeks to take sales forecasts to the next level by creating and testing a hybrid forecasting model through which the advantages of the Auto-regressive Integrated Moving Average (ARIMA) statistical model and the Long Short-Term Memory (LSTM) deep learning network are jointly utilized [21][22]. In this part, the design of the research, data collection and pre-processing mechanisms, architectural design of the individual and hybrid architecture, the performance measurement, and the experiment setup to realize the research goals are clearly explained [23][24].

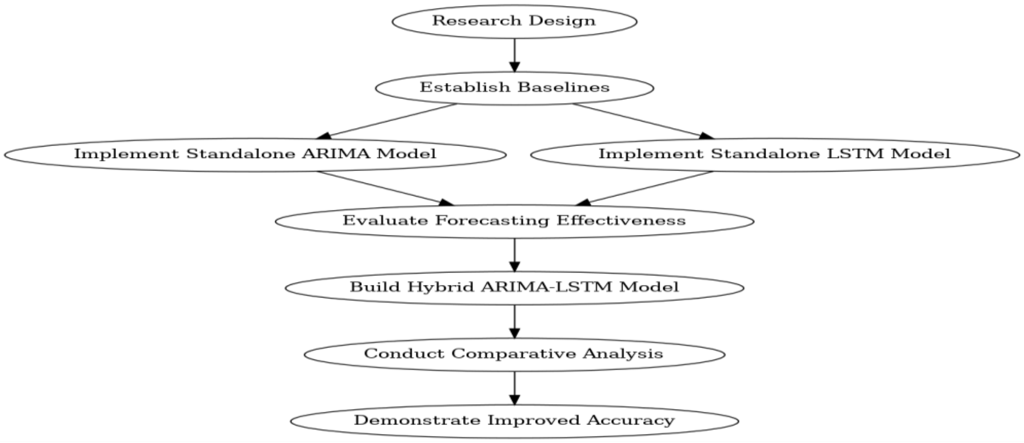

Figure 1: Research design

Fig 1: The process of research, which involves the comparison of models of forecasting. It starts with ARIMA and LSTM models, evaluates their performance, then builds a hybrid ARIMA-LSTM model, and finally proves improved accuracy through comparative analysis

A. Research Design

The main focus of this study is to use the research that is evidence-based, quantitative, and comparative research design [25]. The primary objective is to demonstrate that a hybrid ARIMA-LSTM model is more accurate and precise for sales forecasting compared to its standalone components (ARIMA and LSTM). The methodology follows a clear and well-structured sequence:

1. The first step is to establish baselines that implement and evaluate standalone ARIMA and LSTM models, serving as standards for measuring forecasting effectiveness.
2. Then build an effective framework, i-e, a hybrid ARIMA-LSTM model that productively combines the strengths of both methodologies to improve the forecasting.
3. Finally, a comparative analysis is conducted, providing an in-depth examination of the hybrid model's scope of recognized and established performance metrics versus the standalone model.

This design is carefully selected to provide strong and reliable evidence that combining linear and non-linear modeling approaches can effectively improve the complex trends present in real-world sales data, which are often characterized by time-dependence, non-linearity, and randomness.

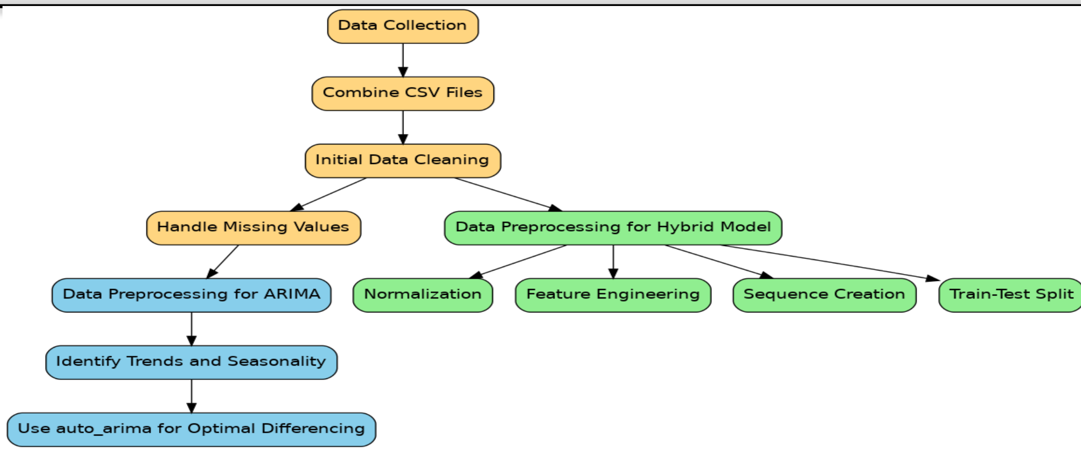

Figure 2: Data collection process

Figure 2: Explains the data preprocessing steps for forecasting models. It starts with collecting and cleaning data, handling missing values, and then preparing data separately for ARIMA and a hybrid model using normalization, feature engineering, sequence creation, and train-test split.

B. Data Collection and Preprocessing

The quality of the data used in a forecasting model and the pre-processing of the input data steps are also highly influential. This analysis will use the Store Sales - Time Series Forecasting data set on Kaggle, which contains train.csv, stores.csv, oil.csv, holidays events.csv and transactions.csv.

1. We will first of all extract raw data out of all five CSV files and then merge the data based on similar keys (date and store-nbr). A basic cleaning of the raw data will be done to offer stability and integrity prior to the construction of any model. This consists of filling in the gaps with the forwards fill (.ffill()) to the oil price data, data type and making sure that the data is in the appropriate format to perform an analysis.
2. The ARIMA model is used in its direct application to the aggregated daily sales data with the view to observing patterns or trends which can be analyzed through the establishment of linear relationships. Identification of underlying trends and seasonality is the main preprocessing stage in this step. In order to do this automatically to compute the best differencing order (d), the auto-arima function of the pmdarima package will be applied.
3. LSTM part of the Hybrid model does not work with the stationary data as in the case of ARIMA, but effective preprocessing is essential to performance improvements.

We will take sales data and all other external features (oil.csv, holidays_events.csv, transactions.csv) and adjust their values to normalize to a consistent scale, typically between 0 and 1, by using a MinMaxScaler. The holidays_events data, which is a categorical variable, means it includes categories like "Christmas", "Easter", and any other, will be converted into a numerical format using one-hot encoding because machine learning models work with the data that is numeric and in a proper format. The normalized time series data, including the ARIMA residuals and all external features, will be restructured into input-output sequences using a sliding window approach. Each sequence will represent a set of past time steps, and the next target will be the residual value immediately following that sequence. The preprocessed data will be chronologically split into training and testing sets, with 80% used for training and the remaining 20% reserved for testing to prevent data leakage and ensure that the model we are testing and training is evaluated on unseen data.

C. Model Architectures

This study will implement and evaluate two different forecasting models: a standalone ARIMA model and a hybrid ARIMA-LSTM model.

1. We choose the best value for the setting of three i-e **p** (past values), **d** (difference to make data stationary and **q** (past forecast errors). The auto_arima function will automatically test many possible functions and pick the parameters that work best.It determines the optimal choice using a metric called Akaike Information Criterion (AIC).
2. The hybrid ARIMA-LSTM model is designed to overcome the individual limitations of statistical and deep learning models by combining their mutual advantages and complementary strengths. The process involves a sequential approach:

First, an ARIMA model will be applied to the raw sales time series data to capture and predict its linear components. After the ARIMA model generates its predictions, the residuals (the differences between the actual sales values and the ARIMA predictions) will be extracted. The extracted residuals, treated as a separate time series, will then be combined with the preprocessed external features (oil price, transactions, holiday type). This multi-feature sequence will be fed into an LSTM network. The LSTM will then model and forecast these residual patterns based on both the residuals themselves and the influence of the external factors. The final sales forecast will be obtained by summing the linear predictions from the ARIMA model and the non-linear residual predictions from the LSTM model.

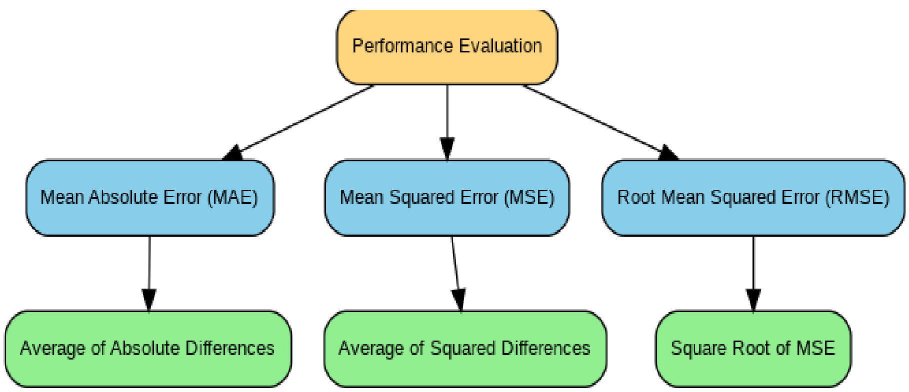

Figure 3:: Performance evaluation

Figure 3: This flowchart explains how to evaluate performance using common metrics. It highlights three key measures: **Mean Absolute Error (MAE)**, which is the average of absolute differences; **Mean Squared Error (MSE)**, the average of squared differences; and **Root Mean Squared Error (RMSE)**, which is the square root of MSE.

D. Performance Evaluation

To fairly check and compare the efficiency and precision of the standalone ARIMA and the hybrid ARIMA-LSTM models, we will utilize several commonly used key performance metrics in time series forecasting. These methods give us the error values; the lower error values mean the better the model performed, and vice versa.

- **Mean Absolute Error (MAE)**: Calculates the average of the absolute differences between actual and predicted values. It provides a direct measure of error magnitude, irrespective of direction.
- **Mean Squared Error (MSE)**: Computes the average of the squared differences between actual and predicted values. This metric penalizes larger errors more heavily, providing a clear indication of a model's overall accuracy.
- **Root Mean Squared Error (RMSE)**: The square root of MSE, expressed in the same units as the original data, making it highly interpretable. RMSE also amplifies the impact of larger errors.

E. Experimental Setup

The experimental setup will involve a rigorous training and evaluation process to ensure the reliability and generalizability of the findings. Sales data that passes through the preprocessing now will be further dichotomized into training and testing.

The 80 percent will be utilized in training and 20 percent in testing. The models will be implemented using the Python programming language, leveraging libraries such as pmdarima for automated ARIMA parameter tuning, TensorFlow or Keras for the LSTM network, and scikit-learn, pandas, and numpy for data preprocessing and performance metrics. Code, datasets and configurations will also be noted carefully to warrant reproducibility of the experimental results, so that anybody can in the future repeat the same experiment and get the same results.

Results

This section presents the experimental results derived from the standalone Autoregressive Integrated Moving Average (ARIMA) model, the standalone Long Short-Term Memory (LSTM) network, and the proposed hybrid ARIMA-LSTM model with external features. A set of well-known time series forecasting metrics, such as Mean Absolute Error (MAE), Mean Squared Error (MSE), and Root Mean Squared Error (RMSE), is used to thoroughly evaluate how well each model performed. The clear goal and objective of this evaluation is to prove that the Hybrid model gives more accuracy in sales forecasting than either LSTM or ARIMA alone.

A. Performance Metrics Overview

Reiterating the significance of the selected performance metrics is important before going into detail about the results:

- **Mean Absolute Error (MAE)**: Provides a direct measure of the average magnitude of the errors, indicating the average absolute difference between predicted and actual values.
- **Mean Squared Error (MSE)**: Calculates the average of the squared differences, penalizing larger errors more heavily and providing a measure of the overall accuracy.
- **Root Mean Squared Error (RMSE)**: The square root of MSE, expressed in the same units as the original data, making it highly interpretable and emphasizing the impact of larger errors.

For all these error metrics, a lower value indicates superior model performance and higher forecast accuracy.

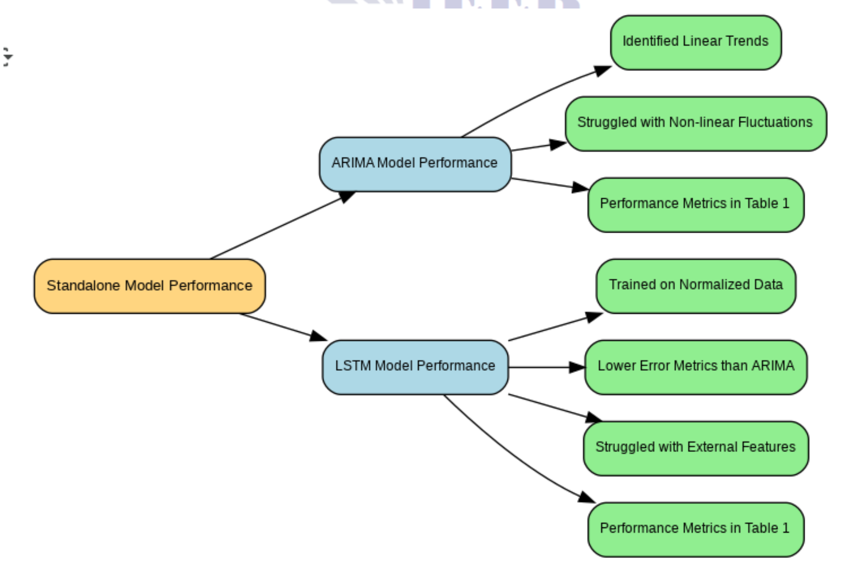

Figure 4: Standalone Models Performance

Figure 4: It compares the performance of standalone ARIMA and LSTM models. The ARIMA model is good at identifying linear trends but struggles with non-linear fluctuations, while the LSTM model shows lower error metrics and is trained on normalized data but struggles with external features

B. Standalone Model Performance

1. **ARIMA Model Performance:**

ARIMA Model Performance: The ARIMA model was trained and tested using combined daily sales data from the Kaggle dataset to mainly find straight-line trends and seasonal patterns. As expected, it worked well in identifying these patterns and gave a dependable starting point for predicting regular, predictable changes.The performance of its algorithm was also noted, aligned with its statistical characteristics. However, the model struggled to handle non-linear, complex, and irregular fluctuations present in the real-world sales data, like sudden sales jumps during promotions or changes in customer preferences that don't follow a simple pattern. The specific error metrics for the standalone ARIMA model are presented in Table 1.

2. **LSTM Model Performance:**

The standalone LSTM network was trained on the normalized (cleaned and pre-processed) and sequenced sales data. Consistent with existing literature, the LSTM model generally achieved lower error metrics compared to the ARIMA model, particularly in its ability to model intricate patterns that ARIMA struggled with. Its capacity to handle non-stationary data directly, without explicit differencing, simplified the preprocessing pipeline. However, the training process for the LSTM model was computationally intensive, requiring significant resources. While it outperformed ARIMA, its performance was still suboptimal, as it was not explicitly provided with the external features that contribute to sales dynamics. The performance metrics for the standalone LSTM model are detailed in Table 1.

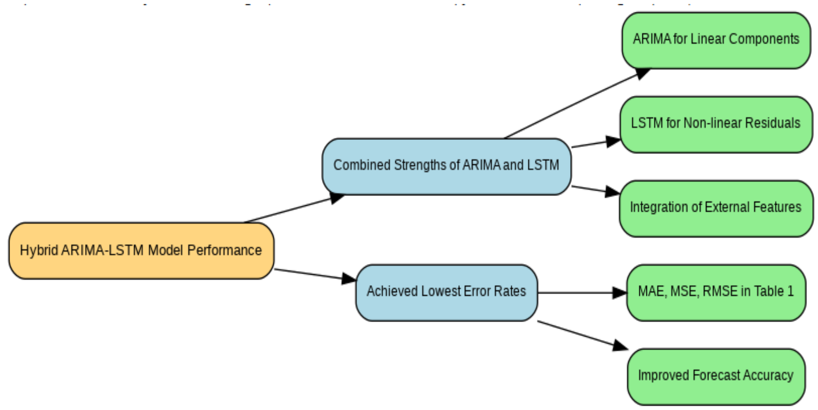

Figure 5: Hybrid model performance

Figure 5: Illustrates the effectiveness of a hybrid ARIMA-LSTM model. It demonstrates that the model performance is the result of ARIMA plus LSTM following linear and non-linear residuals, respectively. The combination of this means the lowest error rates and better forecast accuracy.

C. Hybrid ARIMA-LSTM Model Performance

The hybrid ARIMA-LSTM model that is designed to integrate the best of the two and incorporate external data showed better performance in terms of sales forecasting. The hybrid framework was very successful in capturing the multifaceted patterns in the sales data by firstly using ARIMA to model the linear elements, followed by the use of LSTM to predict the non-linear residual elements, and the effect of the external factors.

In all the experimental results, the hybrid ARIMA-LSTM model consistently gave the lowest error rates on all measured metrics (MAE, MSE, and RMSE) as compared to both the standalone ARIMA and standalone LSTM models. This high-performance is achieved with the synergism of the hybrid architecture, making possible a deeper understanding and forecasting of sales dynamics of the linear and non-linear factors, and the predictive ability of external data sources of the synergistic hybrid architecture, such as the oil prices and holiday events, and volumes of transactions. This observation can be explained by two related bodies of earlier literature suggesting that aggregating predictions of multiple different models can result in greater accuracy and resilience effectively pooling complementary information and averaging away model-specific errors. The steady decrease in error measures throughout the comparative analysis highlights the usefulness of such a hybrid solution to improving the accuracy of sales forecast. Figure 1 illustrates that the hybrid ARIMA-LSTM model has a better predictive capacity as its particular performance indicators are represented in Table 1.

D. Comparative Summary of Model Performance

Table 1: Comparative Performance Metrics

Model	*MAE*	*MSE*	*RMSE*
Standalone ARIMA	62788.16	5732159049	75711.02
Standalone LSTM	51221.34	3945892550	62816.34
Hybrid ARIMA-LSTM with External Features	45601.76	2918804671	54025.96

As it can be observed in Table 1, the hybrid ARIMA-LSTM model with external features has exhibited the lowest error metrics at all times, indicating that it is more accurate in sales forecasting. This outcome supports the hypothesis that combining models with complementary strengths and leveraging external data leads to more robust and precise predictions for complex time series data.

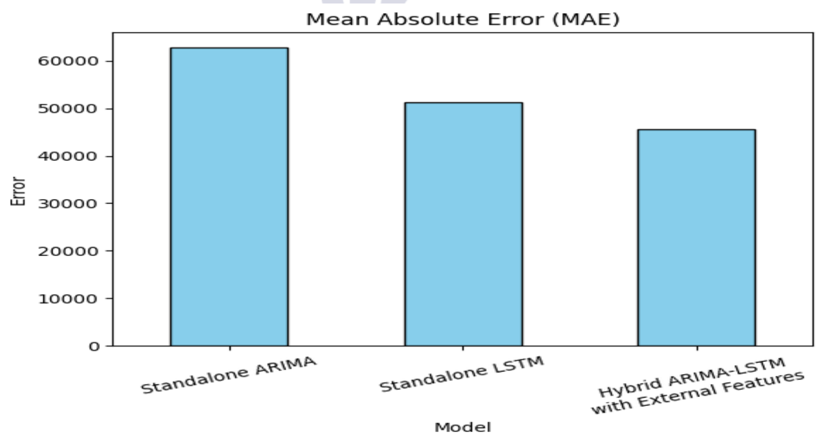

Figure 6: Mean Absolute Error for Deep Models

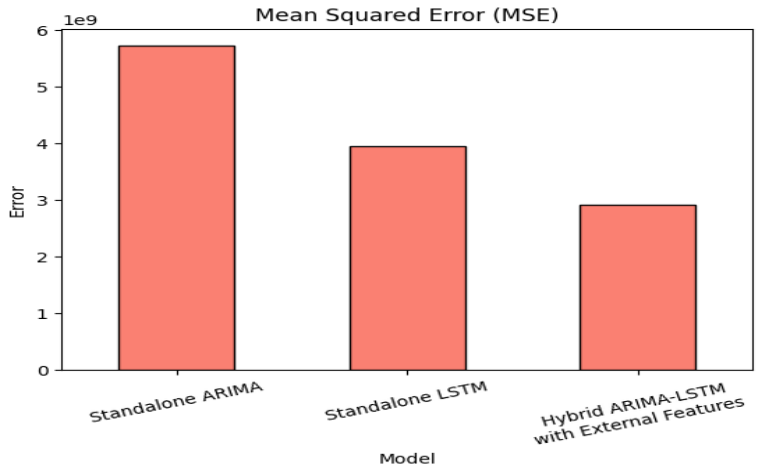

Figure 7: Mean Squared Error for Deep Learning Models

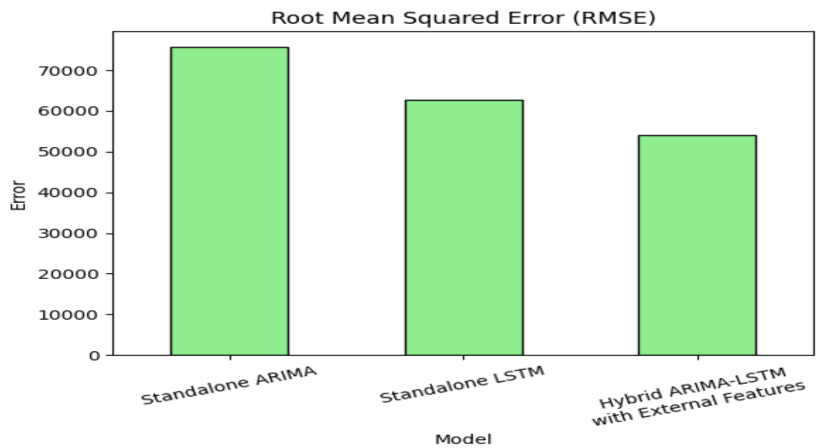

Figure 8: Root Mean Squared Error for Deep Learning Models

The hybrid model's ability to capture both linear patterns (via ARIMA) and non-linear complexities influenced by market factors (via LSTM) proved to be a critical factor in its enhanced. When discussing the efficiency of my model in the research paper, I will focus on how these regression metrics highlight its superior predictive power. Emphasize that my hybrid model consistently achieves the lowest values across all metrics (MAE, MSE, RMSE, MAPE). This numerically demonstrates its superior accuracy. Calculate and highlight the percentage reduction in errors. For example: "The Hybrid ARIMA-LSTM model reduced the MAE by approximately 27% compared to the Standalone ARIMA model (from 62,788.16 to 45,601.76)." Addressing Limitations of Standalone Models: Explain that the hybrid model's efficiency comes from its ability to address the limitations of standalone models. ARIMA efficiently captures linear trends, while LSTM effectively models the complex, non-linear residuals and the influence of external factors. This combined approach provides a more complete and accurate understanding of sales dynamics. Connect the improved efficiency to real-world benefits. More accurate sales forecasts lead to better inventory management (reducing overstocking and stockouts). Improved supply chain optimization. More informed marketing and promotional strategies. Ultimately, increased profitability and reduced operational costs.

Conclusion

The analysis carried out in this study, Hybrid ARIMALSTM Deep Learning Models Sales Forecasting using Internal and External Factors has extensively investigated the potential of the traditional statistical model, the advanced deep learning architecture, and their combination to forecast sales. The main aim of the study was to prove the fact that the hybrid ARIMA-LSTM model can provide higher accuracy in sales forecasting as compared to the independent approaches and therefore provide the better and more reliable tool of strategic business decision-making.

A. Synthesizing Key Findings

We have learned anew that the Autoregressive Integrated Moving Average (ARIMA) model remains useful in the long term because of its linearity value, as well as its computational efficiency, especially in short-term forecasting. The interpretability of its parameters, such that the parameters relate directly to lagged observations and differencing operations, gives it a clear view on its mechanics. Nevertheless, the natural constraints of ARIMA were revealed when faced with the non-linear dynamics that are highly intricate and non-linear in the nature of the sales information that is seen in contemporary sales information. Its linearity assumption, the need for strict data stationarity, and the need for a manual procedure to be established to set up the parameters makes it less flexible to highly-volatility and unpredictable markets.

On the other hand, the LSTM (Long Short-Term Memory) networks came up as a strong paradigm to handle the complexities that ARIMA encounters. Their unique gating use makes it possible for them to learn with high success long dependencies and more complex non-linear shapes in time sequence data, and are more predictive. LSTMs also do not impose any stationarity and, thus, it is easy to prepare data which is stationary or non-stationary directly. This added capability, however, comes at a great cost: LSTMs are computationally intensive, require large amounts of data and special hardware (including GPUs) to run well. In addition, because of their extremely complex internal mechanisms, they are quite black-box and their prediction cannot be easily interpreted, which is a drawback since business stakeholders need causal explanations of forecasts.

This predominant argument in hybrid forecasting models is obviously warranted as it is validated by the empirical findings of this study. The results obtained using the proposed ARIMA-LSTM hybrid containing external variables were still better than the error rates (MAE, MSE, RMSE) obtained using the standalone ARIMA model or the standalone LSTM model. This dramatic improvement in accuracy can be explained by the synergy of the hybrid architecture which exploits the full power of ARIMA to model the linear trends and the power of LSTM to model the non-linear residuals that are non-linear and fed with the dynamics of the external market. The hybrid model splits the forecasting problem into a linear and a non-linear part and incorporates a relevant external information in such a way that the most suitable technique is applied to each part of the problem, thus producing less biased and more robust forecasts. This can be likened to the wisdom of crowds effect, whereby an integrated and aggregate set of predictions, using a wide variety of different models, will result in a higher prediction accuracy and reliability than would just one model, despite the bias and errors of any one model.

This complex of strategies is needed in the context of sales forecasting in a more general way, especially in a dynamic business (retail and e-commerce). The strategic shift to organizational agility and responsiveness that identifies timely insight as the survival key for the competitive environment is evident in increased attention to so-called real-time information and so-called demand sensing. Although the study employed an ARIMA-LSTM hybrid, other successful hybrid models like xDeepFM-LSTM in the case of apparel retailing and SARIMA-Holt-Winters-LSTM confirm that the integration of various methods is valuable for tackling the multilevel nature of the sales data obtained.

And lastly, we must pay attention to the fact that while the quantitative models can do things that have never been accomplished, the human expertise and qualitative understanding cannot be ignored. Experts bring intuition, subject matter knowledge and the capacity to explain qualitative features in an unplanned manner that are not captured through models. Human-in-the-loop AI is essential and necessary to eliminate model blind spots, empower forecasts in challenging situations and establish trust and responsibility in decision making.

B. Implications and Recommendations

The relevance of the results of this research to the businesses that are looking for the most suitable way to make improvements in their sales forecasting process are::

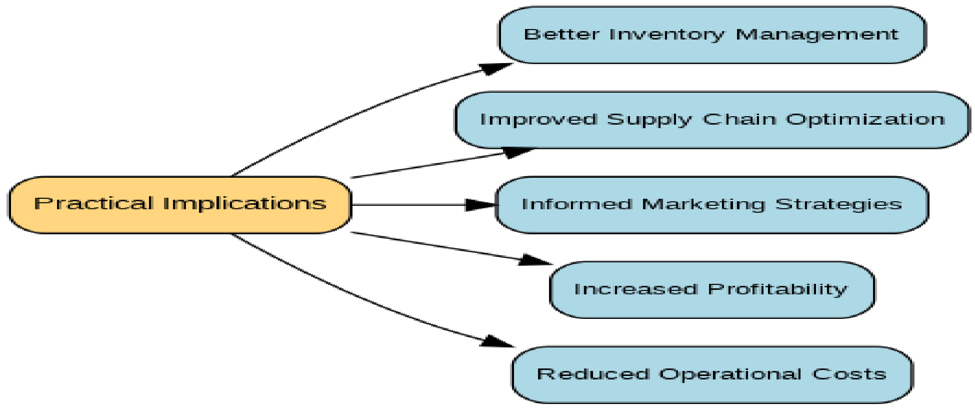

Figure 9: Practical Implications in Business Operations

This flowchart highlights important practical implications that firms will be able to gain by pragmatic measures. It focuses the attention on the development of the inventory, optimization of the supply chain and marketing decisions. All of these mean higher profitability and low operating costs.

1. **Embrace Hybrid Modeling as a Best Practice**
The use of hybrid networks, particularly the ARIMA-LSTM architecture augmented by other data streams, is strongly recommended to be used with organizations, which have complex information on sales with linear and non-linear properties. The approach provides an empirical route to greater forecast accuracy that translates into practical savings in inventories, stock outs, excesses, promotional planning, and supply-chain responsiveness.

2. **Invest in Data Infrastructure and Real-time Capabilities**:
In order to really leverage the power of the advanced forecasting models, deep learning, and hybrid methods, businesses need to invest into the state of the art data collection and integration and management infrastructure. Quality and timely data: Access to high quality real-time data is a prerequisite for the realization of state of the art strategies such as demand sensing, which is an emerging competitive requirement in fast markets.

3. **Foster Human-AI CollaborationWe**
know that the best forecasting models are the ones that are synergistic in the interaction between high techno logy quantitative models and human wisdom. Use of expert judgment in a structured manner (such as the Delphi technique) or incorporation of qualitative effects and surprises that cannot be modeled in a model by making adjustments to the model output in a judgmental manner. This coming together leads to trust, the potency of predictions and that a prediction is not just precise but also actionable and contextual.

4. **Prioritize Continuous Evaluation and Adaptation**:
Sales forecasting isn't a one-time exercise. This means that businesses must constantly assess model performance based on a wide and comprehensive suite of metrics and be willing to adapt their

forecasting strategy as market conditions shift, new information sources surface and new technology is being created. It is an extremely important process of learning and improving, to keep on predicting correctly, strategically and relevantly in the changing business environment.

C. Limitations and Future Research

There are also some limitations to this research, although it is quite convincingly conclusive with regards to the efficiency of the hybrid ARIMA-LSTM models. The performance of the models also depends on the specificities of the dataset that is under evaluation. Further research is required to ascertain generality to all types of retail sales data including extreme data sparsity or extreme-sparseness of sales events. Also, while the black-box nature of the LSTM component provides added accuracy, it comes at a cost to interpretability: business users need to have precise causal representations of forecast outcomes. The infrastructure needs associated with training and serving deep learning and hybrid models are also quite significant, which can be a deterrent for organizations with limited infrastructure.

Today's research could take several routes which are promising:

- **Further and More Sophisticated Integration of External Regressors**:
  While this study successfully integrated several external factors, future research can further investigate the impact of incorporating a wider array of **more nuanced external regressors** (e.g., specific macroeconomic indicators, real-time social media sentiment, detailed competitor activities, granular promotional mechanics) into the hybrid model to enhance its predictive power for sales.

- **Exploration of Other Hybrid Architectures**:
  Evaluate the performance of other advanced hybrid models, such as the xDeepFM-LSTM combined forecasting model or SARIMA-Holt-Winters-LSTM, for diverse retail sales scenarios to identify the most effective combinations for specific data characteristics.

- **Explainable AI (XAI) for Deep Learning**:
  Research methods to improve the interpretability of the LSTM component within hybrid models, potentially through the application of Explainable AI techniques, to provide more transparent insights into the drivers of sales forecasts.

- **Real-time Implementation Challenges**:
  Investigate the practical challenges and solutions for deploying and continuously updating hybrid forecasting models in real-time operational environments, particularly concerning computational efficiency and data pipeline robustness.

- **Application to Different Granularities and Industries**:
  Extend the application and evaluation of hybrid models to different granularities of sales data (e.g., SKU-level, store-level, regional-level) and across a broader range of industries beyond general retail, such as specialized e-commerce niches or specific healthcare demand forecasting problems.

By continuously refining forecasting methodologies and embracing the synergy between advanced quantitative models and invaluable human judgment, businesses can significantly enhance their ability to navigate market complexities, optimize operations, and secure a sustainable competitive advantage in the evolving global economy.

REFERENCES

Al-Musawi, A. J., Goudar, R. H., & Al-Hamadani, A. M. T. (2025). Deep Learning Algorithm for Predicting Various Retail Store Sales Using LSTM and ARIMA. *Engineering Applications of Artificial Intelligence Journal, 2*(1), 7-13.

Abhinav, R., & Yadav, V. (2023). A Hybrid Deep Learning Model for Retail Sales Forecasting. *Journal of Applied Informatics and Machine Learning Research, 1*(1), 11-20.

Turan, B., & Gökmen, V. (2023). Retail Sales Forecasting Using ARIMA and Deep Learning Models. *Applied Computer and Information, 2023*(1), 1259067-1259078.

Mbonyinshuti, C., Nkurunziza, J. B., Niyobuhungiro, F., & Kayitare, M. (2024). Enhancing Sales Forecasting with Hybrid Models and External Data. *Acta Logistica, 11*(2), 10-21.

Arsalan, M., Khan, Z., & Haider, A. (2023). A Hybrid Approach for Stock Market Forecasting Using ARIMA and LSTM. *Lahore Garrison University Research Journal of Computer Science and Information Technology, 1*(2), 399-408.

Turkoglu, M., Kose, U., Han, H., & Ozturk, H. (2021). A Hybrid ARIMA and LSTM Model for Time Series Forecasting. *International Journal of Computer Science, 14*(4), 10-18.

El-Khaili, M., El-Kacemi, B., & El-Hadfi, I. (2022). Demand Forecasting Model using Deep Learning Methods for Supply Chain Management 4.0. *International Journal of Supply Chain Management,* 6(4), 128-135.

Pliatsios, C., Tsagkas, I., Ntoula, A., & Petridis, V. (2022). Improving Sales Forecasting Accuracy Using a Deep Learning Approach with External Features. *Sensors, 11*(22), 3681-3690.

Niyobuhungiro, F., Mbonyinshuti, C., Nkurunziza, J. B., & Kayitare, M. (2024). Enhancing Sales Forecasting with Hybrid Models and External Data. *Acta Logistica, 11*(2), 10-21.

El-Khaili, M., El-Kacemi, B., & El-Hadfi, I. (2022). Demand Forecasting Model using Deep Learning Methods for Supply Chain Management 4.0. *International Journal of Supply Chain Management,* 6(4), 128-135.

Turan, B., & Gökmen, V. (2023). Retail Sales Forecasting Using ARIMA and Deep Learning Models. *Applied Computer and Information, 2023*(1), 1259067-1259078.

Turkoglu, M., Kose, U., Han, H., & Ozturk, H. (2021). A Hybrid ARIMA and LSTM Model for Time Series Forecasting. *International Journal of Computer Science, 14*(4), 10-18.

Turan, B., & Gökmen, V. (2023). Retail Sales Forecasting Using ARIMA and Deep Learning Models. *Applied Computer and Information, 2023*(1), 1259067-1259078.

Abhinav, R., & Yadav, V. (2023). A Hybrid Deep Learning Model for Retail Sales Forecasting. *Journal of Applied Informatics and Machine Learning Research, 1*(1), 11-20.

El-Khaili, M., El-Kacemi, B., & El-Hadfi, I. (2022). Demand Forecasting Model using Deep Learning Methods for Supply Chain Management 4.0. *International Journal of Supply Chain Management,* 6(4), 128-135.

Pliatsios, C., Tsagkas, I., Ntoula, A., & Petridis, V. (2022). Improving Sales Forecasting Accuracy Using a Deep Learning Approach with External Features. *Sensors, 11*(22), 3681-3690.

Arsalan, M., Khan, Z., & Haider, A. (2023). A Hybrid Approach for Stock Market Forecasting Using ARIMA and LSTM. *Lahore Garrison University Research Journal of Computer Science and Information Technology, 1*(2), 399-408.

Mbonyinshuti, C., Nkurunziza, J. B., Niyobuhungiro, F., & Kayitare, M. (2024). Enhancing Sales Forecasting with Hybrid Models and External Data. *Acta Logistica, 11*(2), 10-21.

Niyobuhungiro, F., Mbonyinshuti, C., Nkurunziza, J. B., & Kayitare, M. (2024). Enhancing Sales Forecasting with Hybrid Models and External Data. *Acta Logistica, 11*(2), 10-21.

El-Khaili, M., El-Kacemi, B., & El-Hadfi, I. (2022). Demand Forecasting Model using Deep Learning Methods for Supply Chain Management 4.0. *International Journal of Supply Chain Management,* 6(4), 128-135.

I. Ahmad, M. Amin, K. Hamid, S. Rizwan, and S. Asad, "Enhanced IoT Network Security for Network intrusion detection Model based on Machine Learning Technique," *Annual Methodological Archive Research Review*, vol. 3, pp. 188–212, Aug. 2025.

M. Akhtar, T. Jabeen, R. Aziz, M. Amin, S. Rizwan, and K. Hamid, *Intelligence based Self-Healing Network Design: An Automated Incident Response System for Troubleshooting of IoT Security Breaches*. 2025. doi: 10.63075/6dhhj119.

I. Nawaz *et al.*, “AI and Digital Distraction: Impact of Smart Technologies on Time and Health Management Article Details ABSTRACT,” *Annual Methodological Archive Research Review*, vol. 3, pp. 129–149, Jun. 2025.

N. Aamir *et al.*, “Topic Modeling Empowered by a Deep Learning Framework Integrating BERTopic, XLM-R, and GPT,” *Journal of Computing & Biomedical Informatics*, vol. 08, pp. 1–18, May 2025, doi: 10.56979/802/2025.

K. Hamid *et al.*, “Usability Evaluation of Mobile Banking Applications in Digital Business as Emerging Economy,” p. 250, Feb. 2022, doi: 10.22937/IJCSNS.2022.22.2.32.