International Journal of Web of Multidisciplinary Studies
(Peer-Reviewed, Open Access, Fully Refereed International Journal)
website: www.ijwos.com
Vol.02 No.03.
OPEN ACCESS
E-ISSN : 3049-2424
DOI:

Optimizing Inventory Management: Applying Data Mining to Forecast Demand in Retail

Zeeshan Khan*[1]
*[1]*Student, Bundelkhand University, Jhansi, India*
Jasim Ali *[2]
*[2] *Student, Bundelkhand University, Jhansi,, India*

Article Info

Article History:
(Research Article)
Accepted : 18 Mar 2025
Published:28 Mar 2025

Publication Issue:
Volume 2, Issue 3
March-2025

Page Number:
18-26

Corresponding Author:
Zeeshan Khan

Abstract:
The retail industry operates on thin margins, where efficient inventory management is a critical determinant of profitability and customer satisfaction. Traditional inventory control methods often struggle to cope with the volatile and complex demand patterns inherent in the modern retail environment. This paper explores the application of data mining techniques to optimize inventory management by improving demand forecasting accuracy. We leverage a large transactional dataset from a retail company to develop and compare several forecasting models. The study focuses on identifying significant patterns and variables that influence consumer demand, such as seasonality, promotions, and other external factors. Methodologically, we employ time series analysis, regression models, and machine learning algorithms to predict future sales. Specifically, we implement Autoregressive Integrated Moving Average (ARIMA), multiple linear regression, and a Long Short-Term Memory (LSTM) neural network. The performance of these models is rigorously evaluated using standard metrics, including Mean Absolute Error (MAE), Mean Squared Error (MSE), and Root Mean Squared Error (RMSE). The results indicate that the LSTM model outperforms the traditional ARIMA and regression models, demonstrating a superior ability to capture the non-linear and complex dependencies within the sales data. The comparative analysis reveals that while ARIMA provides a solid baseline, the machine learning approaches offer a significant improvement in forecasting precision, leading to a potential reduction in holding costs and stockouts. This research contributes to the existing body of knowledge by providing a practical framework for retailers to implement advanced data-driven forecasting systems. The findings underscore the transformative potential of data mining in shifting inventory management from a reactive to a proactive paradigm, ultimately enhancing operational efficiency and competitive advantage.

Keywords: Inventory management, demand forecasting, data mining, retail analytics, machine learning, time series analysis, ARIMA, LSTM, predictive modeling.

1. Introduction

The contemporary retail landscape is characterized by intense competition, rapidly changing consumer preferences, and the increasing complexity of supply chains. In this dynamic environment, effective inventory management has emerged as a cornerstone of operational success and a key driver of profitability. The fundamental challenge of inventory management lies in balancing the costs associated with holding stock against the risks of stockouts and lost sales. Overstocking leads to increased carrying costs, including storage, insurance, and the risk of obsolescence, thereby tying up capital that could be used elsewhere. Conversely, understocking results in missed sales opportunities, dissatisfied customers, and potential long-term damage to brand loyalty. Achieving the optimal inventory level requires an accurate prediction of future customer demand, a task that has become progressively more difficult due to market volatility, product proliferation, and the multifaceted nature of consumer behavior.

For decades, retailers have relied on traditional forecasting methods, such as moving averages or exponential smoothing, which are primarily based on historical sales data. While these methods can be effective in stable market conditions, they often fall short in today"s fast-paced retail sector. They are generally limited in their ability to incorporate a wide range of influential factors, such as promotional activities, price changes, holidays, weather patterns, and competitor actions. This limitation often leads to significant forecasting errors, resulting in suboptimal inventory decisions. The advent of the digital age and the proliferation of data from various sources, including point-of-sale (POS) systems, customer loyalty programs, e-commerce platforms, and social media, have created an unprecedented opportunity to overcome these challenges. This wealth of data, often referred to as Big Data, holds the key to unlocking deeper insights into the drivers of demand.

This is where data mining comes into play. Data mining is the process of discovering patterns, correlations, and anomalies within large datasets to predict outcomes. By applying sophisticated algorithms and statistical techniques, data mining can uncover hidden relationships that are not apparent through simple analysis. In the context of retail demand forecasting, data mining can analyze vast amounts of historical data to identify complex, non-linear patterns and the subtle interplay of various demand-influencing variables. This enables the development of more robust and accurate predictive models, empowering retailers to make more informed and proactive inventory management decisions. The application of data mining transforms demand forecasting from a purely historical extrapolation exercise into a predictive, forward-looking process that can adapt to changing market dynamics.

This research paper aims to demonstrate the practical application and benefits of using data mining techniques to improve demand forecasting in a retail setting. We seek to move beyond theoretical discussions by implementing and evaluating several data mining models using a real-world transactional dataset. The primary objective is to compare the performance of traditional forecasting methods with more advanced machine learning algorithms. Specifically, we will develop models using Autoregressive Integrated Moving Average (ARIMA), a widely used statistical method for time series forecasting, as a baseline. We will then develop and compare it against multiple linear regression and a Long Short-Term Memory (LSTM) neural network, a type of recurrent neural network well-suited for sequence prediction problems. The central hypothesis of this study is that machine learning models, by virtue of their ability to learn from complex and high-dimensional data, can provide significantly more accurate demand forecasts than traditional statistical methods. The performance of these models will be rigorously assessed using standard accuracy metrics, and the results will be presented and analyzed to highlight the practical implications for inventory optimization. Ultimately, this paper endeavors to provide a comprehensive guide for retailers on how to leverage their data assets to build a more intelligent and responsive inventory management system, thereby enhancing operational efficiency, reducing costs, and improving customer satisfaction in an increasingly competitive marketplace.

2. Literature Review

The optimization of inventory management has been a central theme in operations research and supply chain management for many decades. The foundational models, such as the Economic Order Quantity (EOQ) model developed by Ford W. Harris, provided a mathematical basis for determining the optimal order size to minimize the total cost of ordering and holding inventory [1]. While groundbreaking, these early models were based on deterministic assumptions of constant demand and lead times, which rarely hold true in real-world scenarios. Subsequent research introduced stochastic inventory models that incorporated uncertainty, using probability distributions to represent demand variability [2]. These models, while more realistic, still often relied on simplified assumptions about the underlying demand patterns. The primary limitation of these classical approaches is their dependence on aggregate historical data and their inability to dynamically adapt to the complex factors influencing modern consumer demand.

The practice of demand forecasting, which is the critical input for any inventory model, has evolved in parallel. Traditional forecasting techniques are broadly categorized into qualitative and quantitative methods. Qualitative methods rely on expert opinion and judgment, while quantitative methods use historical data. Within quantitative methods, time series analysis has been the most prevalent approach in retail forecasting [3]. Techniques such as moving averages, exponential smoothing, and the Box-Jenkins methodology, which encompasses Autoregressive (AR), Moving Average (MA), and Autoregressive Integrated Moving Average (ARIMA) models, have been widely applied [4]. The ARIMA model, in particular, has been a benchmark for time series forecasting due to its statistical robustness in modeling data with trends and seasonality. Several studies have demonstrated its effectiveness in various forecasting contexts, including retail sales [5]. However, a significant limitation of univariate ARIMA models is that they only consider the past values of the time series itself, ignoring other external variables that could significantly impact demand. To address this, extensions like ARIMAX models were developed to include exogenous variables, but they often struggle with non-linear relationships [6].

The emergence of Big Data and advancements in computational power have paved the way for the application of data mining and machine learning techniques to demand forecasting, promising a paradigm shift from traditional statistical methods. Data mining, as defined by Han, Pei, and Kamber, is the non-trivial process of identifying valid, novel, potentially useful, and ultimately understandable patterns in data [7]. This approach allows for the analysis of large, multidimensional datasets to uncover relationships that are not specified in advance. Early applications of data mining in retail focused on market basket analysis using association rule mining to understand which products are frequently purchased together [8]. This provided insights for store layout and promotions but did not directly address the forecasting problem.

More recently, the focus has shifted towards using predictive data mining techniques for demand forecasting. Regression analysis is a fundamental statistical method often employed in data mining. Multiple linear regression, for instance, can model the relationship between a dependent variable (sales) and multiple independent variables (price, promotions, day of the week, etc.). Researchers have shown that incorporating such variables can improve forecast accuracy over univariate time series models [9]. However, linear regression assumes a linear relationship between the independent and dependent variables, which may not accurately reflect the complex reality of consumer behavior. To capture these non-linearities, machine learning algorithms have been increasingly adopted.

Artificial Neural Networks (ANNs), inspired by the structure of the human brain, are powerful tools for modeling complex, non-linear relationships. Numerous studies have demonstrated the superiority of ANNs over traditional statistical methods in various forecasting domains, including retail sales [10],

[11]. ANNs can learn from data without pre-specified model structures, making them highly flexible. However, standard feedforward neural networks do not inherently account for the sequential nature of time series data. This led to the development of Recurrent Neural Networks (RNNs), which have internal memory loops, making them suitable for sequence data. A specific architecture of RNN, the Long Short-Term Memory (LSTM) network, has shown exceptional performance in time series forecasting [12]. LSTMs were designed to overcome the vanishing gradient problem in standard RNNs, allowing them to learn long-term dependencies in the data [13]. Several recent studies have successfully applied LSTM networks to retail demand forecasting, reporting significant improvements in accuracy compared to both traditional methods and standard ANNs [14], [15]. These models have proven adept at capturing complex patterns like seasonality, trends, and the impact of special events from historical sales data.

The integration of these advanced forecasting models into inventory management systems is a growing area of research. Accurate forecasts generated by data mining models can be fed into inventory optimization algorithms to determine optimal reorder points and order quantities. This integration allows for a more dynamic and responsive inventory policy. Research has shown that improved forecast accuracy directly translates into lower inventory levels, reduced holding costs, and fewer stockouts [16]. Furthermore, by analyzing the patterns discovered by data mining models, retailers can gain a deeper strategic understanding of their business. For example, identifying the most influential promotional strategies or understanding how weather affects sales of certain products can inform marketing and operational decisions beyond just inventory control [17]. This body of literature collectively points towards a clear trend: the increasing sophistication of data sources and analytical techniques is enabling a fundamental shift in how inventory is managed. The move is away from static, assumption-laden models and towards dynamic, data-driven systems that can learn and adapt. This paper builds upon this existing work by conducting a rigorous, comparative study of a traditional statistical model (ARIMA) and two data mining approaches (multiple linear regression and LSTM) using a real-world retail dataset, thereby providing empirical evidence of the value of applying advanced data mining for inventory optimization.

3. Methodology

The methodology adopted in this research is designed to provide a structured and rigorous comparison of different data mining techniques for retail demand forecasting. The process can be broken down into several key stages: data acquisition and preparation, feature engineering, model selection and implementation, and model evaluation. The overarching goal is to build predictive models that can accurately forecast daily product sales, thereby enabling optimized inventory management.

The foundation of this study is a large, real-world dataset obtained from a retail chain. The dataset comprises transactional records spanning over a three-year period. It contains granular information at the stock-keeping unit (SKU) level, including the date of sale, the specific store location, the quantity of items sold, and the unit price. To enrich the dataset and provide more predictive power to the models, external data sources were integrated. This included information on promotional activities, such as discounts and marketing campaigns, which were obtained from the company's marketing department. Additionally, public holiday data and special event calendars were incorporated to account for their significant impact on consumer purchasing behavior. The raw data, consisting of millions of transactions, required substantial preprocessing to be suitable for modeling. This initial step involved data cleaning to handle missing values and correct inconsistencies. Missing sales data for certain days were imputed using the average sales of the same day of the week from adjacent weeks. Outliers, identified using the interquartile range (IQR) method, were carefully examined. Those deemed to be data entry errors were corrected, while legitimate extreme values, such as those during major holiday sales, were retained as they contain important information about demand surges. The data was then aggregated to a daily sales level for a specific, high-volume product category to create a consistent time series for forecasting.

Following data cleaning, the next critical step was feature engineering. This process involves creating new input variables (features) from the existing data that can help the models better capture underlying patterns. Based on the literature and domain knowledge, several features were engineered. Time-based features were created to capture seasonality and trends. These included variables for the day of the week, week of the year, month, and year. A binary feature was created to indicate whether a given day was a weekend or a weekday. Similarly, binary flags were generated for public holidays and days when specific promotional campaigns were active. To capture recent trends, we engineered lag features, which are the sales values from previous time steps (e.g., sales from one day ago, seven days ago). We also created rolling window features, such as the moving average of sales over the past seven days, to smooth out short-term fluctuations and highlight longer-term trends. These engineered features, along with the original sales data, formed the final dataset used for training and testing the models. The dataset was then partitioned chronologically. The first 80% of the data was allocated for training the models, and the remaining 20% was reserved as a hold-out test set for evaluating their performance on unseen data. This chronological split is crucial for time series forecasting to ensure that the models are tested on their ability to predict the future, mimicking a real-world deployment scenario.

Three distinct forecasting models were selected for this comparative study to represent a spectrum of techniques from traditional statistics to advanced machine learning. The first model is the Autoregressive Integrated Moving Average (ARIMA). ARIMA is a widely used statistical model for analyzing and forecasting time series data. It is defined by three parameters: p, d, and q. The 'p' parameter represents the number of lag observations included in the model (the autoregressive part). The 'd' parameter is the number of times the raw observations are differenced to make the time series stationary (the integrated part). The 'q' parameter is the size of the moving average window (the moving average part). The optimal parameters (p,d,q) for our ARIMA model were determined using the Akaike Information Criterion (AIC) and by analyzing the autocorrelation (ACF) and partial autocorrelation (PACF) plots of the time series data.

The second model is Multiple Linear Regression (MLR). This model aims to establish a linear relationship between the dependent variable (daily sales) and a set of independent variables (the engineered features like day of the week, promotions, lag sales, etc.). The model takes the form of $Y = \beta_0 + \beta_1X_1 + \beta_2X_2 + ... + \beta_nX_n + \epsilon$, where Y is the predicted sales, X_i are the independent variables, β_i are the regression coefficients learned from the data, and ϵ is the error term. The MLR model provides a clear and interpretable way to understand the marginal effect of each feature on sales.

The third and most advanced model is the Long Short-Term Memory (LSTM) neural network. LSTMs are a special kind of Recurrent Neural Network (RNN) specifically designed to learn long-term dependencies in sequential data. This makes them exceptionally well-suited for time series forecasting. An LSTM unit is composed of a cell, an input gate, an output gate, and a forget gate. These gates regulate the flow of information, allowing the network to selectively remember or forget information over long periods. Our LSTM model was constructed using a sequential architecture with multiple LSTM layers followed by a dense output layer. The model was trained using the Adam optimizer and the mean squared error as the loss function. The input to the LSTM model was a sequence of past observations and corresponding features, and the output was the sales prediction for the next day. The hyperparameters of the network, such as the number of LSTM units, the sequence length, and the learning rate, were tuned through experimentation on a validation set carved out from the training data.

To ensure a fair and robust comparison, all three models were trained on the same training dataset and evaluated on the same unseen test dataset. The performance of each model was quantified using a set of standard forecasting accuracy metrics. The Mean Absolute Error (MAE), which is the average of the absolute differences between the predicted and actual values, was used to measure the average magnitude of the errors. The Mean Squared Error (MSE), which is the average of the squared errors, was used to penalize larger errors more heavily. Finally, the Root Mean Squared Error (RMSE), the square root of the MSE, was used as it provides an error metric in the same units as the target variable (sales), making it more interpretable. These metrics provide a comprehensive view of the models' predictive accuracy and form the basis for the comparative analysis in the subsequent section.

4. Results and Analysis

Following the implementation of the three forecasting models—ARIMA, Multiple Linear Regression (MLR), and Long Short-Term Memory (LSTM)—on the prepared retail sales dataset, a comprehensive evaluation was conducted to assess and compare their predictive performance. The models were tasked with forecasting daily sales on the hold-out test set, which consisted of data the models had not seen during the training phase. The accuracy of these forecasts was measured using Mean Absolute Error (MAE), Mean Squared Error (MSE), and Root Mean Squared Error (RMSE). The quantitative results of this evaluation are summarized in the comparison table below, providing a clear basis for analyzing the relative strengths and weaknesses of each approach.

Table 1: Comparative Performance of Forecasting Models

Model	Mean Absolute Error (MAE)	Mean Squared Error (MSE)	Root Mean Squared Error (RMSE)
ARIMA	28.45	1358.72	36.86
Multiple Linear Regression (MLR)	23.12	981.45	31.33
Long Short-Term Memory (LSTM)	16.57	496.28	22.28

The results presented in Table 1 clearly indicate a hierarchy of performance among the tested models. The ARIMA model, representing the traditional statistical time series approach, served as our baseline. It achieved an RMSE of 36.86. This level of error, while providing a reasonable approximation of future sales, is the highest among the three models. The performance of the ARIMA model is commendable given its relative simplicity and reliance solely on the past values of the sales data itself. It effectively captures the inherent autocorrelation and seasonality within the time series to a certain extent. However, its primary limitation, as anticipated, is its inability to incorporate external factors. The model's forecasts are a projection of past patterns, and it struggles to adapt to shifts in demand caused by external events like promotions or holidays that are not reflected in its simple temporal structure. A visual analysis of its predictions against the actual sales data showed that the ARIMA model tended to smooth out peaks and troughs, failing to capture the sharp increases in sales during promotional periods or the dips on uneventful weekdays.

The Multiple Linear Regression (MLR) model demonstrated a marked improvement over the ARIMA model, achieving an RMSE of 31.33. This represents a performance enhancement of approximately 15% in terms of RMSE. The superior performance of the MLR model can be directly attributed to its ability to incorporate a wide range of engineered features. By including variables for the day of the week, holidays, and promotional activities, the model could establish a more nuanced relationship between these drivers and the resulting sales figures. The coefficients of the trained MLR model provided valuable insights; for instance, the coefficient for the 'promotion' feature was positive and statistically significant, quantitatively confirming that promotions have a strong positive impact on sales. Similarly, weekend days showed higher coefficients than weekdays. Despite this improvement, the MLR model is constrained by its fundamental assumption of linearity. It presumes that the effect of each feature on sales is additive and constant. This is often not the case in reality, where complex interactions between factors exist. For example, the impact of a promotion might be much stronger when it falls on a weekend than on a weekday, a synergistic effect that a standard linear model struggles to capture.

The most compelling results were delivered by the Long Short-Term Memory (LSTM) neural network. The LSTM model achieved an RMSE of 22.28, which is a significant improvement of approximately 39.5% over the ARIMA model and 28.9% over the MLR model. The MAE and MSE metrics show similarly substantial gains, confirming the LSTM's superior predictive power. This outstanding performance stems from the LSTM's sophisticated architecture, which is inherently designed to handle sequential data and capture complex, non-linear patterns over long time spans. Unlike the ARIMA model, the LSTM was able to use the rich set of engineered features, and unlike the MLR model, it was not restricted to linear relationships. The LSTM's internal memory cells and gating mechanisms allowed it to learn the intricate temporal dependencies in the sales data. For example, it could learn that a promotional period not only boosts sales on the day but might also have a lingering effect for several days after. It could also learn complex interaction effects, such as how the combination of a holiday and a promotion leads to an exponential, rather than additive, increase in sales. A graphical plot of the LSTM's predictions versus the actual sales showed a much closer fit compared to the other two models. The LSTM was particularly adept at predicting the magnitude of sales spikes during peak periods and accurately forecasting sales during regular periods, demonstrating its robustness and adaptability.

In conclusion, the analysis of the results reveals a clear and progressive improvement in forecasting accuracy as we move from a traditional univariate statistical model to a multivariate regression model and finally to an advanced deep learning model. The baseline ARIMA model provides a foundational forecast but is limited by its scope. The MLR model shows the significant value of incorporating external variables, offering a more context-aware forecast. However, it is the LSTM network that truly excels, leveraging its ability to learn complex, non-linear, and long-term temporal patterns from the data. The substantial reduction in forecast error offered by the LSTM model has profound practical implications for inventory management. A more accurate forecast directly translates into a more efficient inventory policy, leading to lower safety stock levels, reduced holding costs, minimized risk of obsolescence, and a significant decrease in stockout incidents. This ultimately enhances profitability and improves customer satisfaction, validating the central hypothesis that advanced data mining techniques can fundamentally optimize retail inventory management.

5. Conclusion

This research set out to investigate the application of data mining techniques to enhance the accuracy of demand forecasting within the retail sector, with the ultimate goal of optimizing inventory management. The study conducted a comparative analysis of three distinct forecasting models: a traditional statistical method (ARIMA), a multivariate statistical model (Multiple Linear Regression), and an advanced machine learning algorithm (Long Short-Term Memory network). By leveraging a comprehensive real-world dataset of retail transactions enriched with promotional and calendar information, we have demonstrated a clear and significant progression in predictive accuracy across these models. The empirical results unequivocally support the central thesis of this paper: that data mining, particularly through the use of sophisticated machine learning models like LSTM, offers a powerful tool for overcoming the limitations of traditional forecasting methods.

The findings indicate that while ARIMA models provide a useful baseline, their inability to incorporate external explanatory variables renders them inadequate for the complexities of the modern retail environment. The Multiple Linear Regression model showed a notable improvement by including these variables, highlighting the importance of a context-aware approach to forecasting. However, the standout performance was achieved by the LSTM network, which reduced the forecast error by a substantial margin compared to both other models. The LSTM's success is attributable to its inherent ability to capture complex non-linear relationships and long-term temporal dependencies within the data, which are characteristic of retail sales patterns. The practical implications of this enhanced accuracy are profound. For a retailer, a more precise forecast allows for a direct reduction in inventory holding costs and the frequency of stockouts. This leads to a more efficient allocation of capital, increased profitability, and higher levels of customer satisfaction and loyalty. The adoption of such data-driven forecasting methods facilitates a strategic shift from reactive inventory replenishment to a proactive, predictive, and optimized supply chain.

Despite the promising results, this study is not without its limitations. The research was conducted on a dataset from a single product category within one retail chain. The performance of the models might vary across different product types with different demand characteristics (e.g., fast-moving vs. slow-moving goods, perishable vs. non-perishable items) and in different retail segments. Furthermore, while we included several key features, other external factors such as competitor pricing, local economic indicators, or even weather data could potentially further improve model accuracy. The implementation of an LSTM model also requires greater computational resources and technical expertise compared to simpler models, which could be a barrier to adoption for smaller retailers.

Future research should aim to address these limitations. It would be valuable to extend this comparative analysis across a wider range of product categories and retail environments to validate the generalizability of the findings. Future work could also explore the integration of more diverse data sources, including unstructured data from social media or customer reviews, using natural language processing techniques to extract sentiment and trends. From a modeling perspective, exploring more advanced deep learning architectures, such as hybrid models combining convolutional neural networks (CNNs) for feature extraction with LSTMs, or incorporating attention mechanisms to help the model focus on the most relevant historical information, could yield further improvements. Finally, a crucial next step would be to conduct a simulation study that directly links the improved forecast accuracy from these models to quantifiable inventory cost savings and service level improvements, providing a clear business case for the adoption of these advanced analytical techniques. In conclusion, this research reinforces the transformative potential of data mining and machine learning in reshaping retail operations. As data becomes ever more abundant, the ability to effectively mine it for predictive insights will increasingly become a critical source of competitive advantage.

References

1. F. W. Harris, "How Many Parts to Make at Once," Factory, The Magazine of Management, vol. 10, no. 2, pp. 135-136, 152, 1913.
2. K. J. Arrow, T. Harris, and J. Marschak, "Optimal Inventory Policy," Econometrica, vol. 19, no. 3, pp. 250-272, 1951.

3. S. Makridakis, S. C. Wheelwright, and R. J. Hyndman, Forecasting: Methods and Applications, 3rd ed. New York, NY, USA: John Wiley & Sons, 1998.
4. G. E. P. Box, G. M. Jenkins, G. C. Reinsel, and G. M. Ljung, Time Series Analysis: Forecasting and Control, 5th ed. Hoboken, NJ, USA: John Wiley & Sons, 2015.
5. S. Kumar and D. P. Singh, "Demand Forecasting for a Retail Store Using ARIMA Model," International Journal of Innovative Research in Science, Engineering and Technology, vol. 5, no. 7, pp. 13243-13249, 2016.
6. R. S. Pindyck and D. L. Rubinfeld, Econometric Models and Economic Forecasts, 4th ed. Singapore: McGraw-Hill, 1998.
7. J. Han, M. Kamber, and J. Pei, Data Mining: Concepts and Techniques, 3rd ed. Waltham, MA, USA: Morgan Kaufmann, 2012.
8. R. Agrawal and R. Srikant, "Fast algorithms for mining association rules," in Proc. 20th Int. Conf. Very Large Data Bases (VLDB '94), Santiago, Chile, 1994, pp. 487–499.
9. K. K. F. Yuen, "The effect of promotional activities on the sales of a beverage product," International Journal of Production Economics, vol. 84, no. 1, pp. 79-88, 2003.
10. G. P. Zhang, B. E. Patuwo, and M. Y. Hu, "Forecasting with artificial neural networks: The state of the art," International Journal of Forecasting, vol. 14, no. 1, pp. 35-62, 1998.