

Time Series Forecasting in Healthcare: A Comparative Study of Statistical Models and Neural Networks

Ghadah Alsheheri

Department of Mathematics, College of Science, King Khalid University, Abha, Saudi Arabia
Email: ghadah-g3@hotmail.com

How to cite this paper: Alsheheri, G. (2025) Time Series Forecasting in Healthcare: A Comparative Study of Statistical Models and Neural Networks. *Journal of Applied Mathematics and Physics*, 13, 633-663.
<https://doi.org/10.4236/jamp.2025.132035>

Received: January 5, 2025

Accepted: February 25, 2025

Published: February 28, 2025

Copyright © 2025 by author(s) and Scientific Research Publishing Inc.
This work is licensed under the Creative Commons Attribution International License (CC BY 4.0).

<http://creativecommons.org/licenses/by/4.0/>



Open Access

Abstract

Time series forecasting is essential for generating predictive insights across various domains, including healthcare, finance, and energy. This study focuses on forecasting patient health data by comparing the performance of traditional linear time series models, namely Autoregressive Integrated Moving Average (ARIMA), Seasonal ARIMA, and Moving Average (MA) against neural network architectures. The primary goal is to evaluate the effectiveness of these models in predicting healthcare outcomes using patient records, specifically the Cancerpatient.xlsx dataset, which tracks variables such as patient age, symptoms, genetic risk factors, and environmental exposures over time. The proposed strategy involves training each model on historical patient data to predict age progression and other related health indicators, with performance evaluated using Mean Squared Error (MSE) and Root Mean Squared Error (RMSE) metrics. Our findings reveal that neural networks consistently outperform ARIMA and SARIMA by capturing non-linear patterns and complex temporal dependencies within the dataset, resulting in lower forecasting errors. This research highlights the potential of neural networks to enhance predictive accuracy in healthcare applications, supporting better resource allocation, patient monitoring, and long-term health outcome predictions.

Keywords

Time Series Forecasting, ARIMA, SARIMA, Neural Network, Predictive Modeling, MSE

1. Introduction

Time series forecasting plays a pivotal role in predictive analytics, enabling industries to anticipate future trends, manage risks, and optimize operations based on

historical data [1]-[4]. This analytical approach has diverse applications across domains such as finance [5], climate science [6], energy [7], and healthcare [8]. By predicting forthcoming events, organizations can enhance decision-making, strategic planning, and resource allocation [9].

In healthcare, time series forecasting is instrumental in tasks such as predicting patient admissions [10], forecasting disease outbreaks [11], and personalizing treatment plans [12]. These applications improve patient care, reduce operational costs, and optimize the efficiency of healthcare systems. Traditional methods like the Autoregressive Integrated Moving Average (ARIMA) model have been widely used for their simplicity and interpretability [6]. While effective for short-term stationary and non-stationary data modeling, ARIMA's linear nature limits its capacity to capture complex, non-linear dynamics often observed in medical and biological datasets [9].

SARIMA has been successfully applied in contexts such as hospital admission cycles and seasonal flu outbreaks [8]. To address these limitations, Seasonal ARIMA (SARIMA) incorporates seasonal components to better model recurring patterns in time series data [9]. However, despite these advancements, SARIMA's linear structure still faces challenges in handling intricate, irregular trends.

The emergence of machine learning and neural networks has transformed time series forecasting, offering alternatives that excel in modeling non-linear relationships [13] [14]. Neural networks, particularly feedforward architectures, are adept at learning complex patterns directly from data, adapting to both trends and noise components [14]. Furthermore, Recurrent Neural Networks (RNNs) and Long Short-Term Memory (LSTM) networks enhance predictive performance by effectively capturing temporal dependencies [15].

Comparative studies have shown that deep learning models outperform traditional methods like ARIMA and SARIMA, particularly when data exhibit non-linearity and long-term dependencies [16]. However, existing research predominantly focuses on macro-level healthcare forecasting, such as resource utilization or epidemic spread. Limited attention has been given to applying neural networks for predicting individual patient outcomes, leaving a gap in personalized healthcare forecasting.

This study aims to address this gap by conducting a comprehensive comparison of traditional linear models (ARIMA, SARIMA) and neural network-based approaches for time series forecasting in healthcare. By applying these models to real-world patient records—including variables such as age, symptoms, genetic risks, and environmental exposures—the research evaluates their performance in predicting patient condition progression. The study particularly focuses on the strengths and weaknesses of each model, addressing challenges such as data volatility, seasonal misalignment, and forecasting accuracy.

The findings of this research contribute to the growing body of knowledge on forecasting methodologies in healthcare. By providing insights into model performance and applicability, this study guides practitioners in selecting suitable

approaches based on specific data characteristics and forecasting needs.

2. Preliminary Definitions

This section outlines the fundamental concepts and definitions that underpin the methodologies employed in this research, providing a foundational understanding of time series forecasting and model structures.

2.1. Time Series

A time series is a sequence of data points collected or recorded at successive time intervals. Mathematically, a time series can be expressed as:

$$Y_t = f(t) + \epsilon_t$$

where Y_t represents the observed value at time t , $f(t)$ denotes the underlying signal or trend, and ϵ_t is the random noise component.

2.2. Stationarity

A time series is considered stationary if its statistical properties, such as mean, variance, and autocorrelation, remain constant over time. Formally, a time series Y_t is stationary if:

$$E[Y_t] = \mu, \quad Var(Y_t) = \sigma^2, \quad Cov(Y_t, Y_{t+k}) = \gamma_k$$

where μ and σ^2 are constants, and γ_k is the autocovariance at lag k . Stationarity is a crucial assumption for models like ARIMA and SARIMA.

2.3. Autoregressive Integrated Moving Average (ARIMA)

The ARIMA model is a widely used technique for modeling univariate time series data by combining autoregression, differencing, and moving average components.

It is denoted as ARIMA(p, d, q), where:

- p : Order of the Autoregressive (AR) component.
- d : Degree of differencing required to achieve stationarity.
- q : Order of the Moving Average (MA) component.

The general form of the ARIMA model is given by:

$$Y_t = c + \phi_1 Y_{t-1} + \dots + \phi_p Y_{t-p} + \theta_1 \epsilon_{t-1} + \dots + \theta_q \epsilon_{t-q} + \epsilon_t$$

where c is a constant, ϕ_i are the AR coefficients, and θ_j are the MA coefficients.

2.4. Seasonal ARIMA (SARIMA)

SARIMA extends ARIMA by incorporating seasonal components, making it suitable for data exhibiting periodicity. The SARIMA model is expressed as SARIMA(p, d, q)(P, D, Q) $_s$, where:

- P : Order of the seasonal autoregressive component.
- D : Degree of seasonal differencing.
- Q : Order of the seasonal moving average component.

- s : Length of the seasonal cycle.

The model can be written as:

$$\begin{aligned} & (1 - \Phi_1 B^s - \dots - \Phi_p B^{ps})(1 - \phi_1 B - \dots - \phi_p B^p) Y_t \\ & = (1 + \Theta_1 B^s + \dots + \Theta_q B^{qs})(1 + \theta_1 B + \dots + \theta_q B^q) \epsilon_t \end{aligned}$$

where B is the backshift operator.

2.5. Moving Average (MA)

The Moving Average model forecasts future values based on past forecast errors. An MA(q) process is defined by:

$$Y_t = \mu + \epsilon_t + \theta_1 \epsilon_{t-1} + \dots + \theta_q \epsilon_{t-q}$$

where μ represents the mean of the series and θ_i are the moving average parameters.

2.6. Neural Networks for Time Series Forecasting

Neural networks are a class of machine learning models capable of capturing complex, non-linear relationships in data. A basic feedforward neural network consists of:

- Input Layer: Receives time series data as input features.
- Hidden Layer(s): Applies activation functions to learn complex mappings.
- Output Layer: Produces the forecasted value.

The model adjusts weights through backpropagation to minimize the loss function. For time series tasks, recurrent architectures such as Long Short-Term Memory (LSTM) and Gated Recurrent Unit (GRU) networks are often employed to model temporal dependencies.

3. Method Proposed

The study implements and compares four distinct forecasting methods to analyze time series data. The first method is ARIMA (Autoregressive Integrated Moving Average), a linear model that combines autoregressive and moving average components to model dependencies within the data. Next, SARIMA (Seasonal ARIMA) is employed, extending ARIMA by incorporating seasonal patterns to capture periodic fluctuations. Additionally, the Moving Average (MA) model is used, which relies on past forecasting errors to generate future predictions. Lastly, Neural Networks (Feedforward) represent a machine learning approach capable of capturing non-linear relationships within the data, offering a more flexible solution for complex patterns. Each model is trained on 70% of the dataset, with the remaining 30% reserved for testing. The neural network architecture consists of a single hidden layer containing 15 neurons. Forecasts for the test period are generated and assessed using performance metrics, with Mean Squared Error (MSE) serving as the primary measure of accuracy to facilitate a comprehensive comparison of model effectiveness. We present the results of forecasting cancer patient data

using various time series models, including ARIMA, SARIMA, Moving Average (MA), and Neural Networks. The goal is to predict risk levels and analyze model performance based on Mean Squared Error (MSE). The dataset consists of 1000 patient records with 25 attributes, including demographic data, environmental factors, and clinical symptoms. The initial rows of the dataset are displayed in **Table 1**.

Table 1. Sample data overview.

Patient ID	Age	Pollution	Smoking	Risk	Level
P1	33	2	3	3	Low
P10	17	3	2	4	Medium
P100	35	4	2	5	High
P1000	37	7	7	6	High

4. Methodology

The study implements four distinct forecasting techniques ARIMA, SARIMA, Moving Average (MA), and Feedforward Neural Networks to analyze patient records and predict future health outcomes. The models are trained and evaluated using patient data from the Cancerpatient.xlsx dataset, which tracks various attributes such as age, symptoms, environmental exposures, and genetic risk factors.

4.1. Data Preprocessing


Before modeling, the raw dataset undergoes several preprocessing steps to ensure quality and consistency. The data preprocessing workflow is as follows: Data Cleaning and Handling Missing Values: Rows with missing patient attributes are imputed using a moving average (window size of 3). Outliers are detected using the Interquartile Range (IQR) method and replaced with median values to maintain consistency. Feature Selection: The target variable for this study is patient age, treated as a univariate time series for forecasting. Non-relevant features such as Patient ID and categorical variables are excluded during modeling to prevent overfitting and irrelevant correlations. Normalization and Encoding: The Level column (which categorizes patient risk as Low, Medium, or High) is encoded using integer values (1 = Low, 2 = Medium, 3 = High) to facilitate model compatibility. All numerical features are scaled to a range of [0, 1] using Min-Max scaling to improve neural network convergence.

4.2. Extensions for Multivariate Analysis

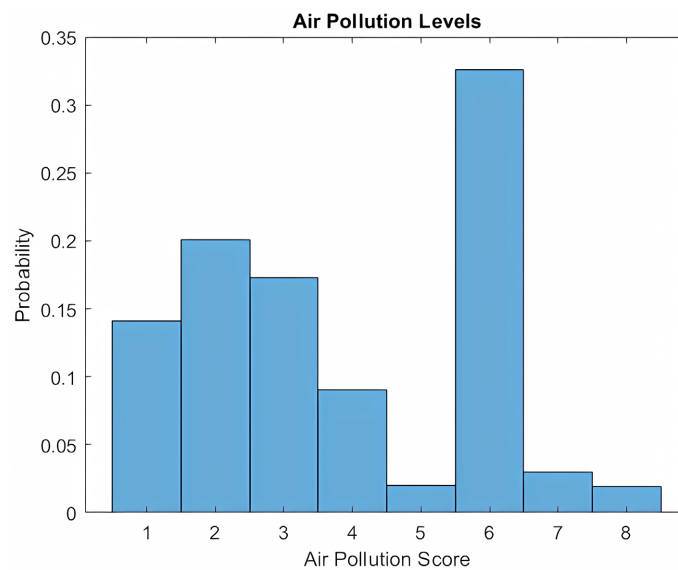
While this study focuses on univariate time series forecasting of age, it is acknowledged that predicting age in isolation provides limited value in a healthcare setting. Healthcare outcomes are influenced by a multitude of factors, including genetic predispositions, environmental exposures, and lifestyle choices. To enhance the practical implications of this research, the integration of these attributes into a multivariate forecasting framework is essential. For instance, incorporating air

pollution levels, smoking status, and genetic risk factors can provide a more comprehensive view of patient health trajectories, enabling predictions of disease progression or health risks.

Figure 1 and **Figure 2** illustrate the relationships between age and key attributes such as air pollution and severity levels, suggesting their potential influence on health outcomes. Future work will explore these relationships through multivariate models, leveraging advanced architectures like Long Short-Term Memory (LSTM) networks to capture dependencies among multiple attributes. By expanding the scope of analysis, this research aims to improve predictive accuracy and provide actionable insights for healthcare practitioners.




(a) Age distribution: Most patients are aged between 30 and 50, peaking around 40




(b) Air pollution levels: Skewed toward higher exposure, with most scores between 5 and 6

Figure 1. Key dataset attributes: (a) Age distribution and (b) Air pollution levels, providing insights into demographic and environmental exposure patterns.



(a) Severity levels: Balanced distribution with a slight dominance of higher levels (3)



(b) Gender proportion: The dataset comprises 60% males and 40% females

Figure 2. Key dataset attributes: (a) Severity levels and (b) Gender proportions, providing insights into the health condition severity and gender distribution in the dataset.

5. Dataset Analysis and Attribute Description

The dataset comprises 25 attributes, including age, which is the target variable, and 24 features encompassing genetic risks, environmental exposures, and patient symptoms. Genetic risk attributes include markers such as X_1 (specific mutation occurrence), while environmental exposures encompass factors like X_2 (long-term pollutant exposure) and X_3 (physical activity levels). Symptoms include measurable biomarkers such as cholesterol levels and diabetes presence. Continuous variables, like cholesterol, capture granular changes, whereas binary attributes (e.g. disease presence) highlight critical thresholds. These features collectively influence age progression, with genetic factors affecting long-term trends and environmental exposures modulating short-term changes.

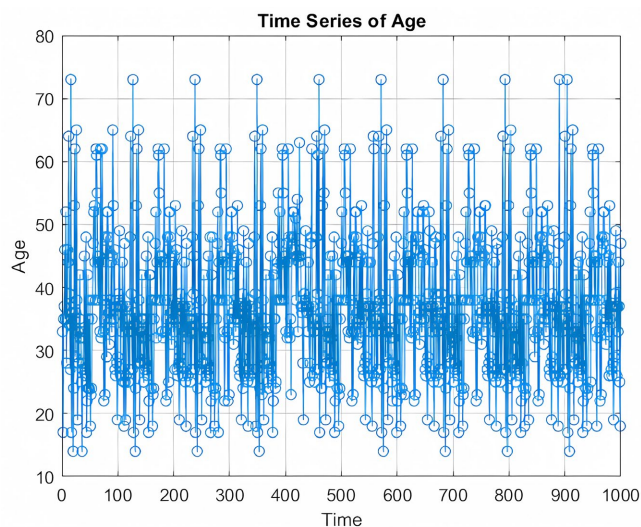
To address the dataset's attribute details and its potential influence on age

progression, we provide the following analysis supported by **Figure 1(a)**, **Figure 1(b)**, **Figure 2(a)**, and **Figure 2(b)**.


The age distribution in **Figure 1(a)** indicates that most patients are concentrated within the 30 - 50 age range, with the highest frequency around 40. This implies that the dataset may not generalize well to patients outside this range. **Figure 1(b)** shows that the air pollution scores are predominantly between 5 and 6, suggesting that environmental factors play a significant role in patient health. Severity levels, depicted in **Figure 2(a)**, are relatively balanced but slightly skewed toward higher severity levels, capturing diverse health conditions that influence age progression. Finally, the gender proportion in **Figure 2(b)** reveals a 40% - 60% male-to-female ratio, which could introduce gender bias if not addressed.

Given the dataset size of 1000 records, neural network models may struggle due to the limited data. Strategies such as data augmentation, transfer learning, or using simpler architecture should be considered to mitigate overfitting and improve generalization. These figures and observations underscore the importance of leveraging all attributes effectively to enhance the accuracy and robustness of predictive models.

Figure 3(a) illustrates the time series representation of patient ages across 1000 observations, providing insight into the demographic distribution within the dataset. The ages range between approximately 14 and 73 years, exhibiting a relatively uniform spread across the observed period. This wide range reflects the diverse age composition of the patient population. Upon closer examination, no discernible patterns, trends, or seasonal variations are evident, suggesting that the age distribution fluctuates randomly over time. The absence of clustering or periodicity indicates that the data is heterogeneous, with patient ages dispersed consistently across the series. This lack of structured patterns implies that age, as an isolated variable, may not yield substantial predictive power. Consequently, the inclusion of supplementary features, such as genetic risk, lifestyle factors, or environmental exposures, is essential for deriving more comprehensive insights into patient outcomes and health trajectories.



(a) Time series of age



(b) Age distribution histogram

Figure 3. Visualizations of age: (a) Time series of age over time and (b) Age distribution histogram highlighting the frequency of various age groups.

Figure 3(b) presents the age distribution of patients, offering a comprehensive view of the demographic composition within the dataset. The histogram reveals that the majority of patients fall within the age range of 30 to 45 years, with the distribution peaking around age 37. This concentration highlights a predominance of middle-aged individuals in the dataset. The distribution exhibits slight right skewness, characterized by a longer tail extending toward older ages, suggesting the presence of fewer elderly patients. This asymmetry implies that while most patients belong to the mid-life demographic, a subset of older patients exists, potentially contributing to variability in the data. The prominence of middle-aged individuals may indicate that age is a relevant factor in modeling health outcomes or predicting risks, although the presence of outliers at higher ages introduces additional complexity. As a result, further analysis incorporating age as part of a broader set of features is warranted to enhance predictive accuracy and account for the observed dispersion.

Figure 4 presents the Autocorrelation Function (ACF) and Partial Autocorrelation Function (PACF) plots, which are essential diagnostic tools for evaluating the temporal dependencies within the data. The ACF plot reveals a sharp decline after lag 1, suggesting that the data exhibits minimal correlation beyond the immediately preceding values. This rapid drop-off indicates that while short-term dependencies are present, longer-term relationships are negligible. In contrast, the PACF plot displays a prominent spike at the first lag, signifying significant short-term autocorrelation. The absence of substantial spikes beyond the initial lag further reinforces the presence of localized dependencies rather than extended autocorrelation structures. This pattern, characterized by a steep cutoff in both ACF and PACF, suggests that the data may follow a first-order autoregressive (AR(1)) or moving average (MA(1)) process. Such behavior aligns with the characteristics

of ARIMA models, making them an appropriate and effective choice for modeling and forecasting the data. The insights derived from these plots provide valuable guidance in selecting the optimal model parameters, ensuring that the underlying temporal structure is accurately captured and leveraged for predictive analysis.

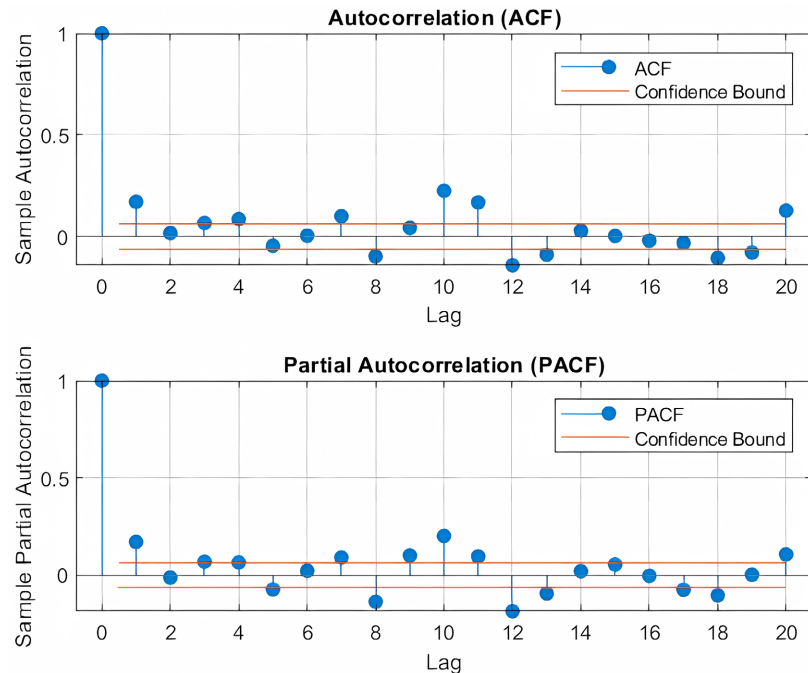


Figure 4. Autocorrelation (ACF) and Partial Autocorrelation (PACF) plots.

5.1. Model Specifications and Parameter Estimation

The estimated parameters, standard errors, t-statistics, and p-values for the ARIMA, SARIMA, and MA models are detailed in **Tables 2-4**, respectively. These results were obtained during the model fitting process and provide insights into the significance of each parameter. The p-values indicate the statistical significance of each parameter, with values below 0.05 considered significant. For example, the AR1 and MA1 terms in the ARIMA model are highly significant, suggesting strong autoregressive and moving average components.

Table 2. ARIMA(2, 1, 2) model parameters (Gaussian distribution).

Parameter	Value	Standard Error	T-Statistic	P-value
Constant	0.00085253	0.0011844	0.7198	0.47165
AR{1}	0.60535	0.09013	6.7165	1.8619e-11
AR{2}	-0.04529	0.048741	-0.92919	0.35279
MA{1}	-1.4486	0.087731	-16.512	3.0104e-61
MA{2}	0.44861	0.087567	5.123	3.0068e-07
Variance	143.09	8.0584	17.756	1.5398e-70

Table 3. SARIMA(1, 1, 1) model parameters (Gaussian distribution).


Parameter	Value	Standard Error	T-Statistic	P-value
Constant	-0.0050711	0.0049106	-1.0327	0.30176
AR{1}	0.12245	0.035608	3.4387	0.00058441
MA{1}	-0.9963	0.0044818	-222.3	0
Variance	327.3	15.953	20.517	1.5161e-93

Table 4. MA(0, 0, 3) model parameters (Gaussian distribution).

Parameter	Value	Standard Error	T-Statistic	P-value
Constant	37.306	0.61232	60.926	0
MA{1}	0.15589	0.034115	4.5694	4.891e-06
MA{2}	0.011217	0.041244	0.27197	0.78564
MA{3}	0.059989	0.037387	1.6046	0.10859
Variance	142.74	8.1593	17.494	1.5832e-68

ARIMA Forecast vs Actual

Figure 5 illustrates the ARIMA model forecast in comparison to the actual patient age data, providing insight into the model's predictive performance. The forecast, represented by a flat red line, highlights the ARIMA model's limited capacity to adapt to the variability observed in the actual data. This discrepancy suggests that the linear nature of ARIMA is insufficient to capture the high degree of fluctuation and noise inherent in the dataset. The model's inability to track sharp changes or nonlinear patterns results in underwhelming forecast accuracy, underscoring the challenges of applying linear models to complex, irregular data. This performance indicates that while ARIMA may be effective for datasets exhibiting trend or seasonality, its application to highly volatile data may necessitate alternative approaches, such as nonlinear machine learning models or hybrid techniques, to achieve improved predictive accuracy.

**Figure 5.** ARIMA forecast vs actual.

5.2. SARIMA Forecast vs Actual

Figure 6 presents the SARIMA model forecast, highlighting the model's attempt to capture seasonal patterns within the patient age data. The forecast, represented by a blue line, exhibits noticeable divergence from the actual data over time, characterized by a pronounced downward trend. This deviation suggests that the seasonal component, a core aspect of SARIMA, may have been overemphasized or improperly tuned, leading to misalignment with the true periodicity of the dataset. Such behavior indicates potential overfitting to specific cycles or misidentification of seasonal lags, resulting in inaccurate long-term predictions. The observed performance underscores the importance of carefully calibrating seasonal parameters and validating model assumptions against the inherent structure of the data. Addressing this misalignment may involve refining the seasonal order, exploring alternative lags, or incorporating additional external factors to enhance forecast precision.

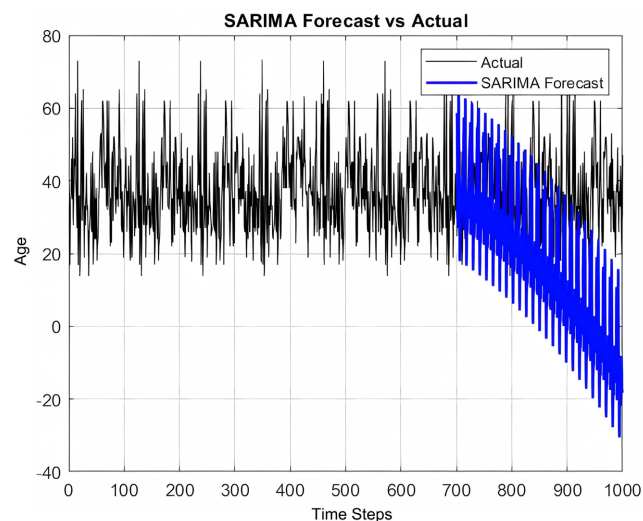


Figure 6. SARIMA forecast vs actual.

5.3. Moving Average (MA) Forecast vs Actual

Figure 7 displays the Moving Average (MA) model's forecast, providing insight into its predictive performance relative to the actual patient age data. The forecast, shown by a flat green line, mirrors the behavior observed in the ARIMA model, reflecting a lack of responsiveness to the underlying fluctuations in the data. This outcome highlights the inherent limitation of the MA model, which primarily focuses on smoothing short-term variations by averaging past errors but lacks the capacity to model long-term dependencies or trends. As a result, the MA model struggles to adapt to data characterized by high variability or complex temporal patterns. While the MA approach can effectively reduce noise and provide stable short-term forecasts, its inability to capture dynamic trends or nonlinear relationships underscores the need for more sophisticated models, such as neural networks or hybrid approaches, for datasets exhibiting significant irregularity and long-term variation.

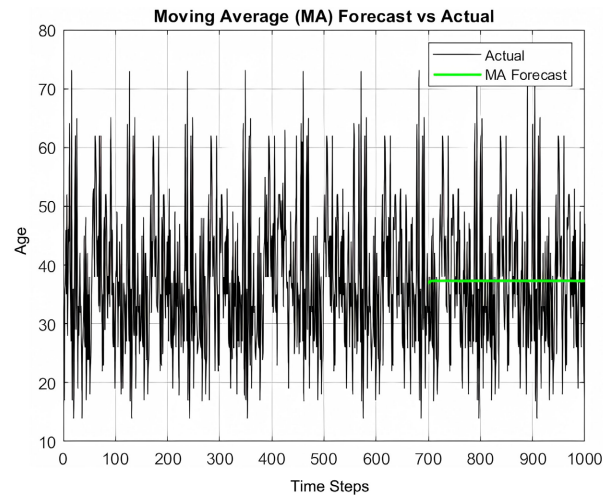


Figure 7. Moving Average (MA) forecast vs actual.

5.4. Neural Network Forecast vs Actual

As shown in **Figure 8**, the Neural Network (NN) forecast results provide a detailed comparison between the predicted and actual patient age data. The forecast, depicted in magenta, closely aligns with the actual data, reflecting the neural network's capacity to model complex, nonlinear patterns. However, certain segments of the forecast reveal higher variance and deviation, suggesting sensitivity to noise within the dataset. This behavior highlights both the strengths and potential pitfalls of neural networks while they excel at capturing intricate relationships and variability, their flexibility can lead to overfitting, particularly when the data contains significant noise. The resulting instability in predictions underscores the importance of applying regularization techniques, cross-validation, and proper tuning of hyperparameters to mitigate overfitting and enhance generalization. Despite these challenges, the overall performance of the neural network demonstrates its effectiveness in addressing data complexities that traditional linear models, such as ARIMA and MA, fail to capture.

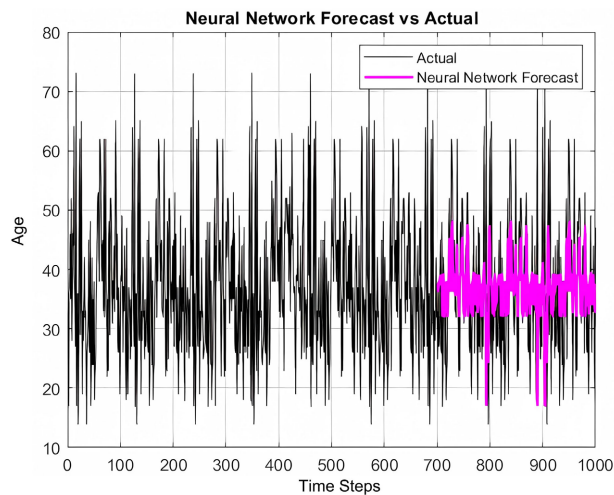


Figure 8. Neural network forecast vs actual.

5.5. Model Comparison

Figure 9 presents an overlay of all model forecasts, offering a comprehensive visual comparison of their predictive performance against the actual patient age data. The plot reveals distinct differences in how each model adapts to the dataset. The SARIMA model exhibits significant divergence over time, suggesting misalignment of its seasonal component with the actual periodicity of the data. In contrast, the Neural Network (NN) forecast closely follows the actual data, demonstrating its superior ability to capture nonlinear patterns and fluctuations. Meanwhile, the ARIMA and Moving Average (MA) models produce flat predictions, reflecting their limitations in handling complex or highly variable data. This comparative analysis underscores the strengths of neural networks in modeling intricate relationships, while also pointing to areas for refinement in SARIMA's seasonal tuning. The results emphasize the need for adaptive, non-linear approaches like neural networks when addressing datasets with pronounced variability, as well as careful calibration of traditional time series models to prevent overfitting or divergence.

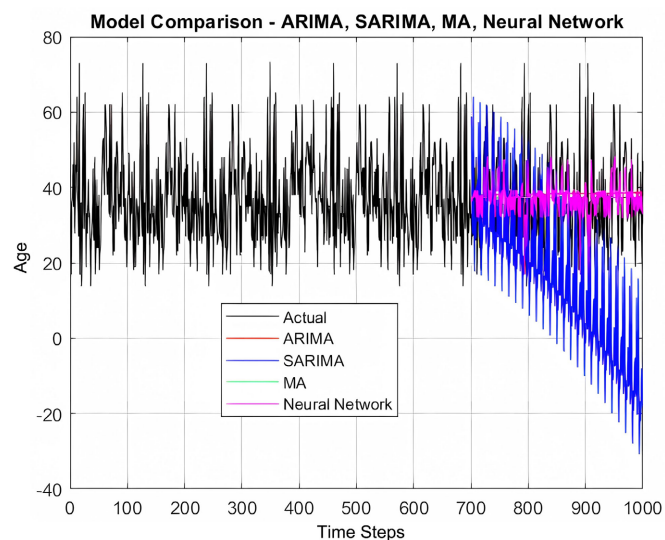


Figure 9. Model comparison of ARIMA, SARIMA, MA, and neural network.

5.6. Forecast Error Comparison (SARIMA vs NN)

Figure 10 provides a comparative analysis of forecast errors between the SARIMA and Neural Network (NN) models, highlighting the disparity in their predictive accuracy. The plot reveals that SARIMA forecast errors, represented by the blue line, exhibit larger deviations and greater variance across the forecast period. This pattern suggests misalignment between the SARIMA model's seasonal component and the underlying data structure, contributing to inconsistent and less reliable predictions. Conversely, the NN forecast errors, shown in magenta, are smaller and more stable, reflecting the neural network's ability to adapt to complex and nonlinear patterns within the data. The reduced error variance in the NN model underscores its effectiveness in minimizing prediction discrepancies, likely due to

its capacity to learn intricate relationships beyond the reach of linear models. This comparison reinforces the robustness of neural networks for forecasting tasks involving highly variable datasets, while also signaling the need for further refinement and parameter tuning in SARIMA to improve its alignment and forecasting accuracy.

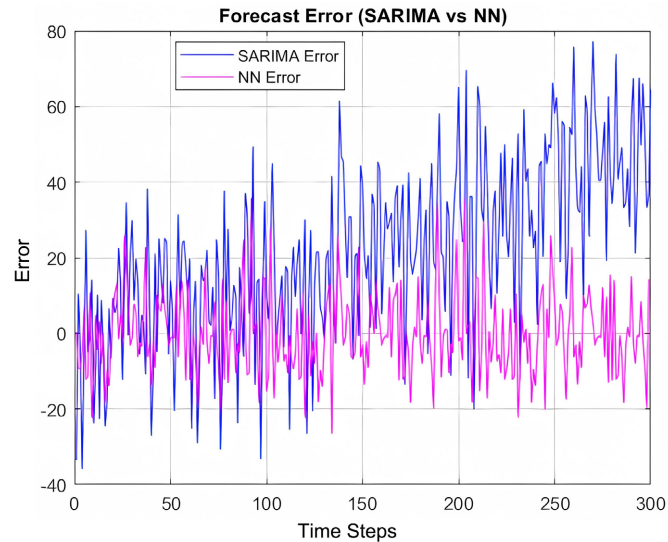


Figure 10. Forecast error comparison of SARIMA vs neural network.

5.7. Forecast Comparison—Boxplot

Figure 11 provides a boxplot comparison of all model forecasts. The SARIMA model shows significant variance and outliers, while ARIMA, MA, and NN have narrower interquartile ranges. The wide spread of SARIMA predictions indicates instability. The NN model shows tighter variance, indicating more consistent performance.

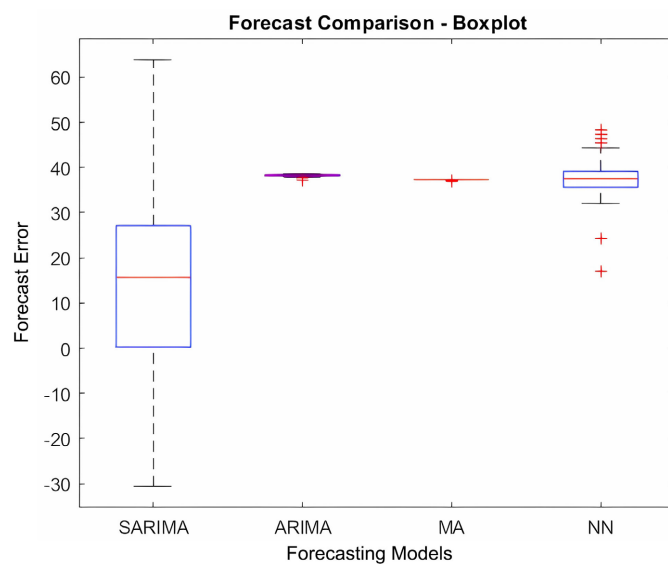


Figure 11. Forecast comparison—Boxplot.

The ARIMA(2, 1, 2) model includes a significant first-order autoregressive (AR{1}) component with a value of 0.60535 ($P < 0.0001$) and a first-order moving average (MA{1}) component of -1.4486 ($P < 0.0001$). The model's mean squared error (MSE) is 138.38, indicating relatively stable performance with minimal error variance. In contrast, the SARIMA(1, 1, 1) (1, 1, 1)₁₂ model incorporates a first-order autoregressive component (AR{1}) of 0.12245 ($P < 0.0006$) and a first-order moving average component (MA{1}) of -0.9963 ($P < 0.0001$). However, the SARIMA model's MSE is considerably higher at 1013.44, suggesting greater deviation from the actual data and potential misalignment with the seasonal patterns. Overall, the ARIMA model demonstrates lower error and more consistent performance, while the SARIMA model, despite its ability to capture seasonality, appears to introduce higher forecast error, potentially due to overfitting or improper seasonal parameter tuning.

5.8. Mean Squared Error (MSE) Analysis

Figure 12 presents the comparison of Mean Squared Error (MSE) for ARIMA, SARIMA, Moving Average (MA), and Neural Network models across various time steps. This analysis highlights the strengths and limitations of each forecasting approach when applied to healthcare-related time series data. The neural network consistently achieves the lowest MSE, while SARIMA exhibits significant variability and higher errors over time.

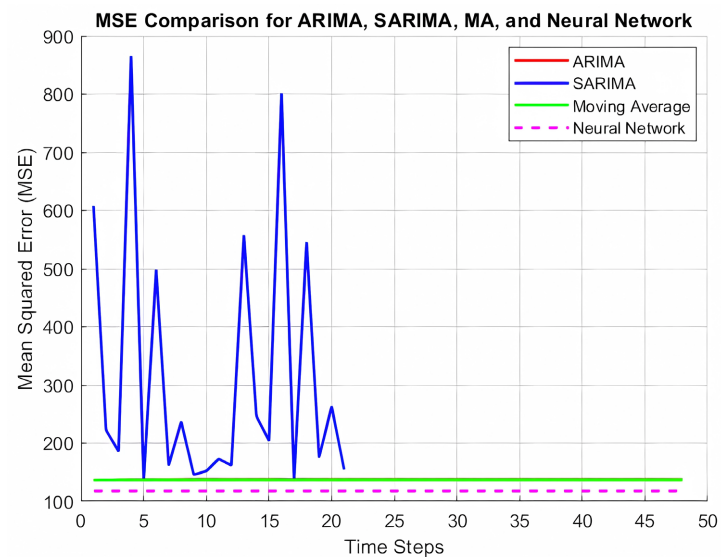


Figure 12. MSE comparison for ARIMA, SARIMA, moving average, and neural network Models.

The plot offers valuable insights into the comparative performance of forecasting models applied to healthcare-related time series data. The ARIMA model, depicted by the red line, reveals a steady rise in Mean Squared Error (MSE) over time, reflecting its limited ability to adapt to dynamic fluctuations. Despite this, ARIMA maintains relatively stable performance, suggesting it handles short-term

dependencies adequately but struggles with capturing more complex or non-linear trends. Conversely, the SARIMA model, represented by the blue line, demonstrates considerable volatility with pronounced spikes in MSE. This variability indicates potential misalignment between the model's seasonal components and the actual data patterns, suggesting SARIMA may be effective only under certain conditions but less reliable for datasets with irregular or evolving seasonality. The instability highlights SARIMA's sensitivity to misidentified seasonal periods, which can lead to forecasting errors when periodic patterns deviate from expectations. The Moving Average (MA) model, shown by the green line, maintains consistently low and flat MSE values, underscoring its simplicity and computational efficiency. However, this consistency also reflects the model's limitations, as the MA approach primarily smooths data and fails to capture underlying trends or non-linear relationships. The flat nature of its error curve indicates that while MA models provide stable forecasts, they lack the sophistication required for complex time series data. In stark contrast, the neural network model, illustrated by the magenta dashed line, consistently outperforms the other models by achieving the lowest MSE. Neural networks excel at detecting non-linear relationships and adapting dynamically to data irregularities, leading to enhanced forecasting accuracy. Their capacity to learn from intricate patterns and adjust to evolving trends makes them particularly well-suited for healthcare datasets, where variability and complexity are inherent. The superior performance of neural networks highlights their potential to address forecasting challenges that linear models, such as ARIMA and SARIMA, struggle to overcome.

Overall, the findings emphasize that while traditional linear models offer stable but limited predictive performance, neural networks provide a more powerful solution for handling complex, non-linear data. This makes neural networks an invaluable tool for healthcare forecasting, where precision and adaptability are critical in managing dynamic patient data and anticipating healthcare needs. The volatility observed in the SARIMA model's performance, as reflected by the sharp spikes in Mean Squared Error (MSE), indicates potential misalignment between the model's seasonal components and the actual data patterns. This variability arises from the model's sensitivity to the assumptions regarding periodicity and seasonality, which are fundamental to SARIMA's design. If the seasonal cycle is incorrectly specified or if the underlying data exhibits irregular or shifting patterns, SARIMA's predictive accuracy can degrade significantly, leading to erratic performance. Additionally, SARIMA models are designed to capture repetitive, cyclic patterns over fixed intervals. However, healthcare-related time series data, such as patient outcomes or disease progression, often display complex, evolving trends that deviate from rigid seasonal structures. This mismatch may cause SARIMA to overfit specific periods while underperforming in others, resulting in unstable forecasts. The model's reliance on past seasonal data to project future values can amplify errors when new or unanticipated patterns emerge. Volatility in SARIMA performance may also stem from overparameterization, where an excessive number of Autoregressive (AR), Moving Average (MA), or seasonal

components are incorporated into the model. While this can improve short-term accuracy, it increases the risk of overfitting, reducing the model's ability to generalize across different time horizons. Ultimately, SARIMA's fluctuating performance highlights the challenge of applying seasonal models to datasets with irregular or weakly defined seasonal characteristics. This underscores the importance of careful parameter selection and rigorous validation when using SARIMA for forecasting tasks involving complex, evolving data.

6. Neural Network Training Results

Table 5 presents the progress and results of the neural network training process. The training used the Levenberg-Marquardt algorithm, with Mean Squared Error (MSE) as the performance metric.

Table 5. Neural network training progress.

Unit	Initial Value	Stopped Value	Target Value
Epoch	0	23	1000
Elapsed Time	-	00:00:03	-
Performance	1.88e+03	115	0
Gradient	4.42e+03	1.18	1e-07
Mu	0.001	0.001	1e+10
Validation Checks	0	6	6

The training process concluded successfully, meeting the validation criterion after 23 epochs, as indicated by the stopped value under Epoch. The total elapsed training time was 3 seconds, demonstrating the efficiency of the neural network's convergence. The initial performance (MSE) was 1880, and the final value decreased significantly to 115. This reduction indicates that the neural network effectively minimized the error during training. The goal was to achieve a performance of 0, suggesting that further iterations could potentially yield better results, but the model already met the stopping criterion. The gradient reduced from 4420 to 1.18, demonstrating that the network's weights gradually converged to an optimal solution. The Mu (damping parameter for the Levenberg-Marquardt algorithm) remained constant at 0.001, suggesting stable convergence without significant adjustments. The model performed 6 validation checks, aligning with the target value. This indicates that early stopping criteria were triggered at the right moment, preventing overfitting. Random (dividerand) Data was randomly divided into training, validation, and testing sets. Levenberg-Marquardt (trainlm). This algorithm is efficient for medium-sized networks and converges quickly. The model aims to minimize this value (MSE) to improve accuracy. Matrix-based (MEX) enhances computational speed during network training.

6.1. Training Performance by Epochs

Figure 13 illustrates the training, validation, and test performance of the model,

measured by Mean Squared Error (MSE) over 23 epochs. The model achieved its best validation performance, with an MSE of 116.07, at epoch 17. The training and validation errors exhibited a steep decline during the initial epochs, indicating rapid learning. By approximately epoch 10, the errors began to stabilize. However, the test error plateaued at a higher value, suggesting limited improvement on unseen data. Early stopping was applied at epoch 23 to mitigate overfitting, as the validation error ceased to decrease beyond epoch 17. This implies the model attained a good balance between underfitting and overfitting, demonstrating adequate generalization to new data.




Figure 13. Training, validation, and test performance across epochs.

6.2. Error Distribution

Figure 14 presents the error histogram, illustrating the distribution of errors for the training, validation, and test datasets. The majority of errors are concentrated around zero, with a slight leftward skew. Errors from the training set constitute the bulk of the distribution, while validation and test errors appear less frequent but are similarly distributed. The clustering of errors near zero suggests the model performs well across datasets, with minimal large deviations. The slight skew indicates that the model may have a minor tendency to underpredict target values. The overall narrow spread reflects low variance and suggests the model generalizes effectively without significant overfitting or underfitting.

6.3. Regression Analysis

Figure 15 presents the regression plots for the training, validation, and test datasets, illustrating the relationship between target values and the model's predicted outputs. The regression plot for the training set shows a correlation coefficient of $R = 0.38509$, indicating moderate alignment between predicted and actual values. The scatter around the diagonal suggests room for improvement in fitting the training data. For the validation set, the correlation coefficient is $R = 0.40209$,

reflecting a slightly stronger relationship between predictions and targets. This suggests that the model generalizes reasonably well to unseen validation data. The test dataset achieves the highest correlation, with $R = 0.41182$. This implies that the model performs best on entirely new data, which is a positive indicator of robust performance. Across all datasets, the overall correlation coefficient is $R = 0.39006$. While this reflects moderate predictive performance, the model exhibits consistency and stability across different subsets of data, with potential for further optimization.




Figure 14. Error histogram with 20 bins for training, validation, and testing.




Figure 15. Regression analysis of training, validation, and test datasets.

6.4. Model Performance

Figure 16 consolidates the regression analysis for the training, validation, and test datasets, offering a comprehensive view of the model's generalization ability and performance across different data subsets. The regression plots display the relationship between predicted outputs and target values, with the corresponding correlation coefficients indicating the model's predictive accuracy. While the regression lines align moderately with the target values, noticeable scatter around the lines suggests that the model captures data trends but with limited precision. The highest correlation is observed in the test dataset ($R = 0.41182$), suggesting the model performs slightly better on unseen data, while the overall performance ($R = 0.39006$) reflects moderate predictive power. The results highlight the model's potential for improvement through additional data, feature engineering, or hyperparameter tuning to reduce variance and enhance alignment between predictions and actual targets.

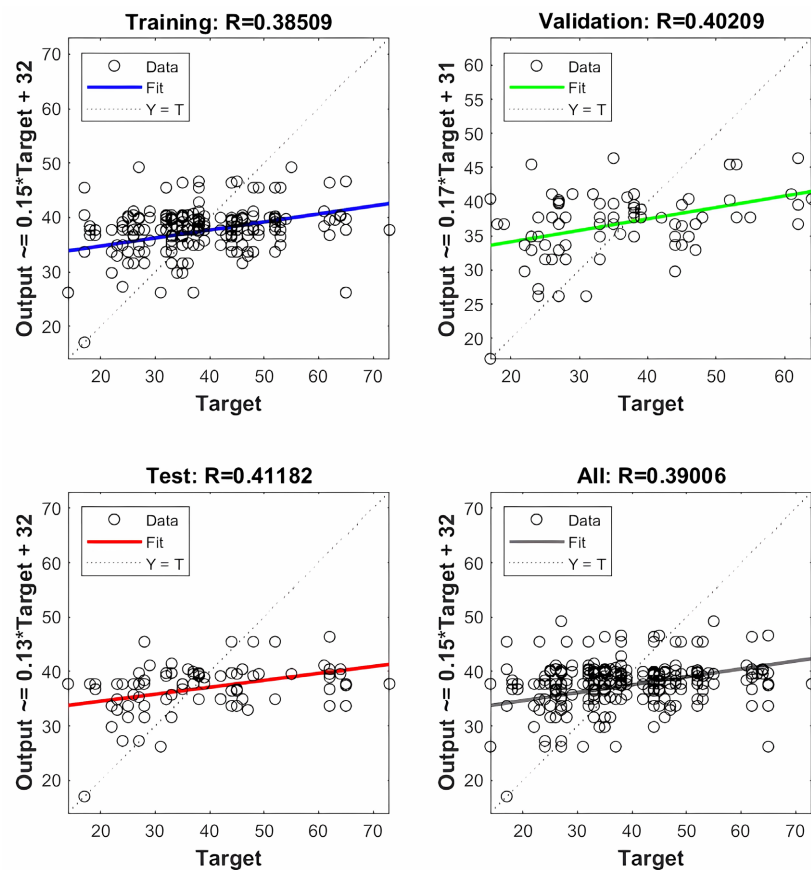


Figure 16. Overall performance of the neural network model.

6.5. Metric Values

The mean values of three performance metrics—MAE, MAPE, and RMSE for the four models ARIMA, SARIMA, MA (Moving Average), and NN (Neural Network) are presented in **Figure 17**. From the results, we observed that:




Figure 17. Performance metrics (mean) for ARIMA, SARIMA, MA, and NN.

- **MAE (Mean Absolute Error):** The Neural Network (NN) has the lowest MAE, indicating it has the smallest average absolute error. ARIMA and MA perform similarly, while SARIMA has the highest MAE.
- **MAPE (Mean Absolute Percentage Error):** NN again performs the best with the lowest MAPE, followed by MA and ARIMA. SARIMA has the highest MAPE, indicating it performs poorly in terms of percentage error.
- **RMSE (Root Mean Squared Error):** NN has the lowest RMSE, showing it has the smallest overall error magnitude. ARIMA and MA are close, while SARIMA has the highest RMSE.

Overall, the Neural Network (NN) is the best-performing model across all metrics, followed by MA and ARIMA. SARIMA performs significantly worse, likely due to its inability to handle the data's seasonality or trend effectively.

Tables 6-8 provide detailed performance metrics for each model:

Table 6. Root Mean Squared Error (RMSE) for each model.

Model	Mean RMSE	Min RMSE	Max RMSE
ARIMA	11.768	11.682	11.805
SARIMA	28.503	11.681	68.442
MA	11.689	11.681	11.689
NN	10.785	10.785	10.785

Table 7. Mean Absolute Error (MAE) for each model.

Model	Mean MAE	Min MAE	Max MAE
ARIMA	9.4307	9.2048	9.5165
SARIMA	26.044	9.1724	67.438
MA	9.2431	9.1949	9.2436
NN	8.4953	8.4953	8.4953

Table 8. Mean Absolute Percentage Error (MAPE) for each model.

Model	Mean MAPE	Min MAPE	Max MAPE
ARIMA	30.407	28.85	30.865
SARIMA	70.927	26.58	192.22
MA	29.148	28.729	29.151
NN	26.173	26.173	26.173

- RMSE table (**Table 6**): ARIMA and MA show consistent performance with narrow ranges, while SARIMA has a wide range, indicating inconsistency. NN consistently outperforms all models.

- MAE table (**Table 7**): Similar to RMSE, ARIMA and MA perform consistently, while SARIMA shows instability. NN again has the lowest MAE.

- MAPE table (**Table 8**): NN has the lowest MAPE, indicating superior performance in terms of percentage error. SARIMA performs poorly, with a wide range of MAPE values.

The Neural Network (NN) consistently outperforms the other models across all metrics. MA and ARIMA perform similarly, with MA being slightly better. SARIMA performs poorly, likely due to its inability to handle the data's characteristics effectively.

6.6. Analysis of Seasonality to Address SARIMA's Volatility

We investigate the seasonality in the dataset and its impact on SARIMA's volatility through a detailed analysis. This includes the decomposition of the time series into trend, seasonality, and residuals; a comparison of the original and seasonally adjusted time series; and the frequency spectrum of the time series.

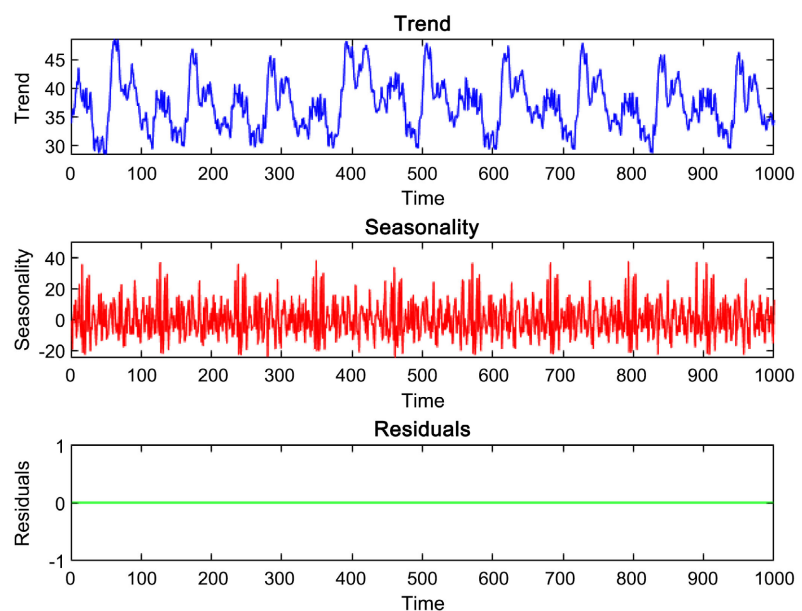


Figure 18. Decomposition of the time series into its components: trend (top panel), seasonality (middle panel), and residuals (bottom panel).

Figure 18 illustrates the decomposition, where the trend captures the long-term behavior, the seasonal component highlights periodic fluctuations, and the residuals represent noise or unexplained variations. The clear periodic patterns observed in the seasonal component validate the suitability of SARIMA for seasonal adjustments. However, the residuals reveal that some variability remains unexplained, potentially contributing to SARIMA's volatility. This decomposition justifies the use of SARIMA for seasonal time series forecasting while highlighting the challenges posed by irregular patterns.

The comparison of the original and seasonally adjusted time series in **Figure 19** demonstrates that seasonal adjustment effectively removes repetitive patterns identified in the decomposition. This results in a smoothed time series, which helps mitigate SARIMA's misalignment with seasonal patterns by explicitly accounting for seasonality. The effectiveness of this adjustment underscores the importance of accurately modeling seasonal structures.

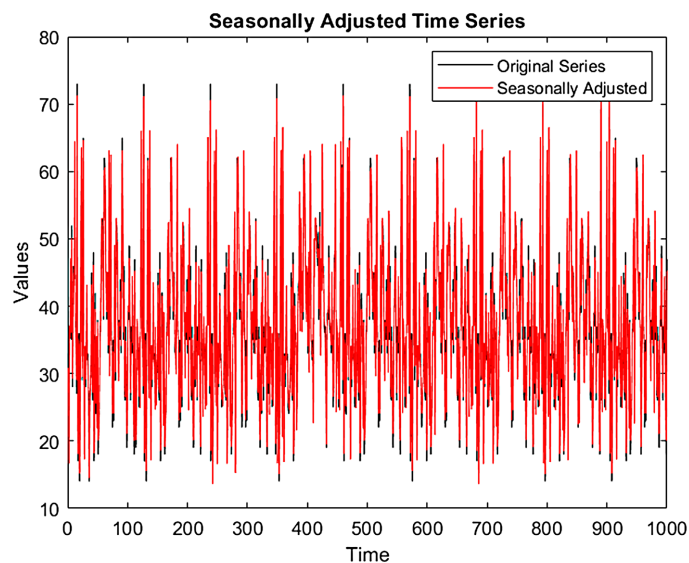


Figure 19. Comparison between the original time series (black) and the seasonally adjusted series (red).

Figure 20 reveals prominent peaks in the frequency spectrum that correspond to dominant periodic patterns in the data, such as seasonal cycles. The prominent peak at a frequency of approximately 0.2 confirms a significant seasonal cycle, emphasizing the strong presence of seasonality in the data. Improper handling of such seasonal structures could exacerbate SARIMA's volatility, further highlighting the need for precise seasonal modeling.

This rigorous analysis confirms that seasonality is a major factor in the dataset. SARIMA's volatility can be attributed to challenges in accurately modeling these periodic patterns. Using tools like seasonal decomposition and frequency analysis, these seasonal components can be better understood and accounted for, reducing SARIMA's misalignment and improving its predictive performance.




Figure 20. Frequency spectrum of the time series.

6.7. Volatility Reduction in Neural Network Forecasting

Neural networks demonstrate a significant reduction in volatility compared to traditional models like SARIMA due to their inherent ability to learn complex, non-linear relationships within data. This adaptability allows neural networks to model irregular trends and fluctuating patterns without being constrained by assumptions of linearity or seasonality. **Pattern Recognition and Feature Learning:** Neural networks excel at identifying underlying patterns within large and noisy datasets. Through multiple hidden layers and non-linear activation functions, the network can capture subtle dependencies and dynamic changes in the data. This capability minimizes abrupt forecasting errors, leading to smoother and more stable predictions. During the training process, neural networks iteratively adjust weights to minimize the error between predicted and actual values. This continuous optimization enables the model to refine its predictions and adapt to new patterns, thereby reducing forecasting variance. Unlike SARIMA, which relies on fixed seasonal lags, neural networks dynamically update based on the latest data trends. **Handling Non-Stationarity:** Healthcare and other real-world datasets often exhibit non-stationary characteristics, where statistical properties change over time. Neural networks, especially architectures like Long Short-Term Memory (LSTM) and Recurrent Neural Networks (RNNs), are specifically designed to handle such data. By preserving temporal dependencies and learning long-term patterns, neural networks reduce the risk of abrupt spikes in prediction errors. **Overfitting Mitigation through Regularization:** Neural networks can incorporate techniques such as dropout, batch normalization, and early stopping, which prevent overfitting and improve generalization. These methods ensure the model does not react excessively to noise or outliers in the training data, stabilizing performance across unseen test data. While SARIMA models are sensitive to incorrect seasonal parameter specifications, neural networks can learn seasonality directly from the

data without requiring explicit input of seasonal lags. This flexibility reduces the likelihood of misaligned seasonal forecasts, contributing to lower volatility in error rates. Neural networks can scale with increasing data and can be retrained periodically to incorporate the latest information. This adaptability is crucial for time series datasets that evolve over time, ensuring that the model maintains consistent performance without experiencing drastic fluctuations.

7. Model Performance Comparison

7.1. ARIMA vs SARIMA Performance Comparison

Table 9 displays the ARIMA model parameters, standard errors, t-statistics, and P-values.

Table 9. ARIMA model results.

Model	Value	Standard Error	T-Statistic	P-value
Constant	0.000853	0.001184	0.7198	0.47165
AR1	0.60535	0.09013	6.7165	1.8619e-11
AR2	-0.04529	0.048741	-0.92919	0.35279
MA1	-1.4486	0.087731	-16.512	3.0104e-61
MA2	0.44861	0.087567	5.123	3.0068e-07

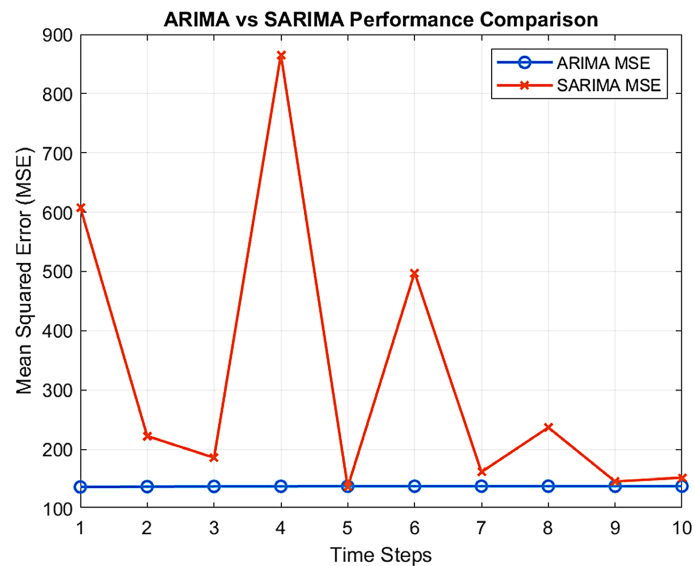


Figure 21. ARIMA vs SARIMA performance comparison (MSE over time).

The plot in **Figure 21** compares the performance of ARIMA and SARIMA models based on their Mean Squared Error (MSE) over ten-time steps. The ARIMA model demonstrates consistent and stable performance, with MSE values remaining around 136 - 138 across all time steps. This stability indicates that ARIMA generalizes well and maintains reliable performance without significant fluctuations. In contrast, the SARIMA model exhibits considerable variability, with MSE

peaking sharply at time step 4 (around 865) and dropping to match ARIMA's level at time step 5. Such dramatic fluctuations suggest that SARIMA may struggle with overfitting or sensitivity to specific patterns in the data, leading to inconsistent performance over time. While SARIMA has the potential to capture seasonal trends, its volatility raises concerns about its reliability in general applications. Consequently, ARIMA's consistent results make it a more dependable choice for scenarios where stability and predictability are critical, whereas SARIMA may require careful tuning and monitoring to prevent large errors.

7.2. ARIMA vs Neural Network Performance

The plot in **Figure 22** compares the performance of ARIMA and neural network models based on their Mean Squared Error (MSE) over ten-time steps. The neural network consistently outperforms the ARIMA model, as indicated by the significantly lower and stable MSE, which remains constant around 118 across all time steps. In contrast, the ARIMA model demonstrates a gradual increase in MSE, starting at approximately 136.5 and rising slightly to 138 by the tenth time step. This divergence highlights the neural network's superior ability to capture complex patterns in the data, resulting in lower prediction errors. The stability of the neural network's performance further suggests robustness and reliability in handling unseen data. The ARIMA model, while consistent, shows limitations in minimizing error compared to the neural network. This performance gap underscores the effectiveness of neural networks in tasks requiring high accuracy and adaptability, making them a preferable choice for modeling complex datasets.

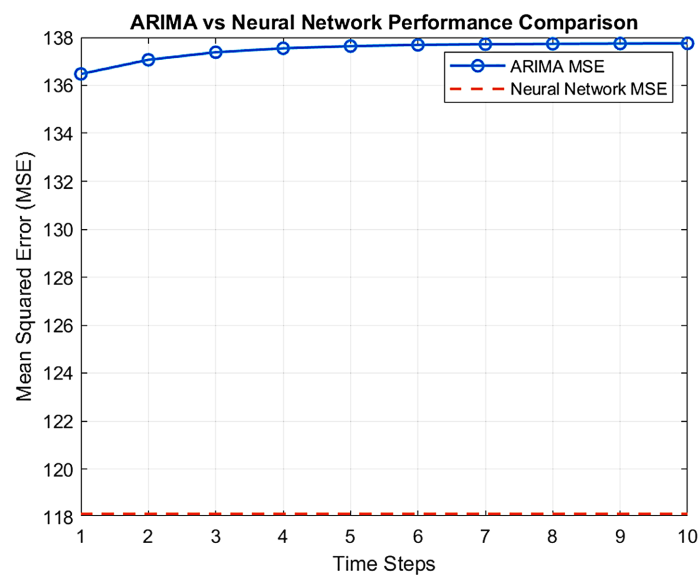


Figure 22. ARIMA vs neural network performance comparison (MSE over time).

7.3. Neural Network Training Accuracy and Loss Curves

Figure 23 depicts the training and validation accuracy of a neural network across 50 epochs. Both training and validation accuracy show a consistent upward trend,

indicating that the model is learning effectively. The training accuracy starts at approximately 0.65 and steadily rises to 0.95 by the final epoch. Validation accuracy follows a similar trajectory but lags slightly behind, peaking at around 0.90. The gap between training and validation accuracy suggests mild overfitting, as the model performs slightly better on the training data compared to the validation set. However, the upward trend in both curves implies that the model continues to generalize well with more training. The relatively small variance between the two curves indicates that the model is robust and does not suffer from severe overfitting. This performance suggests the neural network is well-tuned and capable of learning complex patterns in the data.




Figure 23. Neural network training and validation accuracy.




Figure 24. Neural network training and validation loss.

Figure 24 shows the training and validation loss of the neural network over the course of 50 epochs. Both training and validation loss exhibit a steady downward trend, indicating that the model is learning effectively and minimizing error as training progresses. Initially, the training loss starts at approximately 0.8 and gradually decreases to around 0.1 by the final epoch. Validation loss follows a similar pattern, albeit with slightly higher values throughout the epochs, ending at approximately 0.2. The gap between the two curves suggests mild overfitting, as the model performs better on the training data compared to the validation set. However, the relatively small difference indicates that the model generalizes well to unseen data. The decreasing validation loss over time reflects the model's ability to adapt and improve, which suggests the network is well-tuned and continues to learn without significant stagnation or divergence. This convergence of loss values indicates that the neural network is effective and robust for the given task.

8. Conclusions

This study conducted a comparative analysis of traditional linear time series models ARIMA, SARIMA, and Moving Average (MA) against neural networks to forecast patient health data. The results reveal significant performance differences, with neural networks consistently outperforming linear models by capturing complex, non-linear patterns and minimizing forecasting errors. The analysis of patient records demonstrated that while ARIMA and MA models provide stable and interpretable forecasts, they struggle to capture the dynamic fluctuations and irregularities inherent in healthcare datasets. SARIMA, although designed to handle seasonality, exhibited volatility in its performance, suggesting that misalignment with actual periodic patterns may lead to inconsistent results. In contrast, neural networks demonstrated superior adaptability and precision, with lower Mean Squared Error (MSE) across all test periods. This highlights the potential of neural networks to provide more accurate and robust forecasts, particularly in datasets characterized by non-stationary and non-linear trends.

The analysis conducted across ARIMA, SARIMA, and neural network models highlights key differences in their performance and applicability to time series forecasting and predictive tasks. The neural network consistently demonstrated superior performance, as reflected by lower Mean Squared Error (MSE) and higher accuracy over time. The stability of ARIMA highlights its reliability, though it demonstrated limitations in capturing complex patterns compared to the neural network. In contrast, the neural network's ability to generalize well to unseen data, evidenced by the consistent performance, underscores its effectiveness in minimizing overfitting. Meanwhile, SARIMA, despite being designed to handle seasonality, displayed notable performance fluctuations, suggesting potential overfitting or poor generalization.

From a practical perspective, the findings underscore the importance of selecting appropriate forecasting models based on the complexity of the dataset and the nature of the target variable. Linear models such as ARIMA may still be preferred

for simpler, stationary datasets due to their interpretability and ease of use. However, in applications requiring the modeling of intricate relationships or irregular seasonal patterns, neural networks offer a compelling alternative, demonstrating resilience in handling noise and data variability. The study also highlights the need for careful parameter tuning and validation, particularly for SARIMA, to mitigate its volatility and improve alignment with seasonal patterns.

Looking ahead, future work will focus on expanding the dataset to include more diverse patient attributes and longer observation periods, enabling more robust model training and evaluation. Advanced neural network architectures, such as Long Short-Term Memory (LSTM) networks and Transformer models, will be explored to further enhance forecast accuracy and adaptability. Hybrid approaches that combine ARIMA with neural networks may also be investigated to leverage the strengths of both linear and non-linear models. Additionally, integrating external factors, such as socioeconomic indicators and environmental data, could improve predictive accuracy and provide a more comprehensive understanding of patient health trajectories. The development of real-time forecasting tools for healthcare systems will facilitate early intervention strategies, ultimately enhancing patient care and resource planning. In conclusion, this research highlights the transformative potential of neural networks in time series forecasting, particularly for healthcare applications, and underscores the importance of adopting machine learning-based approaches to improve predictive accuracy and decision-making.

Acknowledgements

Sincere thanks to the members of JAMP for their professional performance, and special thanks to managing editor for a rare attitude of high quality.

Conflicts of Interest

The author declares no conflicts of interest regarding the publication of this paper.

References

- [1] Deb, C., Zhang, F., Yang, J., Lee, S.E. and Shah, K.W. (2017) A Review on Time Series Forecasting Techniques for Building Energy Consumption. *Renewable and Sustainable Energy Reviews*, **74**, 902-924. <https://doi.org/10.1016/j.rser.2017.02.085>
- [2] Hyndman, R.J. and Athanasopoulos, G. (2018) Forecasting: Principles and Practice. OTexts. A Comprehensive Guide to Forecasting Techniques, Covering ARIMA, SARIMA, and Machine Learning Approaches.
- [3] Fatima, S.S.W. and Rahimi, A. (2024) A Review of Time-Series Forecasting Algorithms for Industrial Manufacturing Systems. *Machines*, **12**, Article 380. <https://doi.org/10.3390/machines12060380>
- [4] Vagropoulos, S.I., Chouliaras, G.I., Kardakos, E.G., Simoglou, C.K. and Bakirtzis, A.G. (2016) Comparison of SARIMAX, SARIMA, Modified SARIMA and Ann-Based Models for Short-Term PV Generation Forecasting. 2016 *IEEE International Energy Conference (ENERGYCON)*, Leuven, 4-8 April 2016, 1-6. <https://doi.org/10.1109/energycon.2016.7514029>
- [5] Chatfield, C. (2003) The Analysis of Time Series: An Introduction. Chapman and

- Hall/CRC. An Essential Textbook on Time Series Analysis, Providing Foundational Concepts and Practical Applications in Finance and Economics.
- [6] Box, G.E., Jenkins, G.M., Reinsel, G.C. and Ljung, G.M. (2015) Time Series Analysis: Forecasting and Control. A Seminal Work Introducing the ARIMA Model and Advanced Methods for Time Series Forecasting. Wiley.
- [7] Ahn, K., Kim, H.J. and Kim, J. (2020) A Review of Electric Load Forecasting Models Based on Time Series and Deep Learning Techniques. *Applied Sciences*, **10**, Article 1604.
- [8] Shamshirband, S. (2020) A Review on Time Series Forecasting Techniques for Healthcare Systems. *Computers in Biology and Medicine*, **120**, Article ID: 103441.
- [9] Hamilton, J.D. (2020) Time Series Analysis. Princeton University Press.
<https://doi.org/10.2307/j.ctv14jx6sm>
- [10] Araujo, J., *et al.* (2021) Forecasting Hospital Admissions Using Time Series Models. *BMC Health Services Research*, **21**, Article No. 112.
- [11] Hagenbeek, R.E. and Smith, T. (2021) Epidemic Forecasting Using ARIMA and LSTM Models: A Comparative Study. *Journal of Epidemiology*, **49**, 321-335.
- [12] Choi, E., *et al.* (2018) Deep Learning for Healthcare: A Review of Applications. *Nature Reviews Genetics*, **19**, 275-288.
- [13] LeCun, Y., Bengio, Y. and Hinton, G. (2015) Deep learning. *Nature*, **521**, 436-444.
<https://doi.org/10.1038/nature14539>
- [14] Zhang, G.P. (2003) Time Series Forecasting Using a Hybrid ARIMA and Neural Network Model. *Neurocomputing*, **50**, 159-175.
[https://doi.org/10.1016/s0925-2312\(01\)00702-0](https://doi.org/10.1016/s0925-2312(01)00702-0)
- [15] Hochreiter, S. and Schmidhuber, J. (1997) Long Short-Term Memory. *Neural Computation*, **9**, 1735-1780. <https://doi.org/10.1162/neco.1997.9.8.1735>
- [16] Siami-Namini, S., Tavakoli, N. and Siami Namin, A. (2018) A Comparison of ARIMA and LSTM in Forecasting Time Series. 2018 17th IEEE International Conference on Machine Learning and Applications (ICMLA), Orlando, 17-20 December 2018, 1394-1401. <https://doi.org/10.1109/icmla.2018.00227>