Journal of Mathematics and Data Science

[ISSN 3005-544X]
Volume: 04 Issue: 01 | February-2026

JMDS

A Comparative Evaluation of Time Series Forecasting Models for Pedestrian Footfall Prediction in Dublin City Centre

Authors

Mercy Jemursoi Koech [(1)]; Nasim Sobhani [(2)]

Main author's email: mercyjemursoi@gmail.com

(1.2) Griffith College Dublin, Ireland.

Cite this article in APA

Koech, M. J., & Sobhani, N. (2026). A comparative evaluation of time series forecasting models for pedestrian footfall prediction in Dublin City Centre. *Journal of mathematics and data science, 4*(1), 1-14. https://doi.org/10.51317/jmds.v4i1.923

OPEN ACCESS A publication of Editon Consortium Publishing (online)

Article history

Received: 2025-12-24
Accepted: 2026-01-18
Published: 2026-02-26

Scan this QR to read the paper online

Abstract

This study conducts an in-depth comparative evaluation of time series forecasting methodologies applied to pedestrian footfall data from Dublin City Centre. Four distinct models are examined: Holt-Winters Exponential Smoothing, Seasonal Autoregressive Integrated Moving Average (SARIMA), Facebook's Prophet and Long Short-Term Memory (LSTM) neural networks. Using Data Collected via PYRO-BOX sensors across the 15 streets from 2022 to 2024, this study conducts a comparative evaluation of time series forecasting models to address limitations in previous studies, which often focus on single methods or fail to capture the non-linear dynamics of urban pedestrian movement. The study implements an 80/20 train-test split and evaluates model performance using Mean Absolute Error (MAE), Root Mean Squared Error (RMSE), Mean Absolute Percentage Error (MAPE), and R-squared metrics. The findings indicate that the LSTM model performs considerably better than conventional methods, achieving the lowest prediction errors (MAE: 26,069.72; MAPE: 8.82%) and an R^2 of 0.7778. The Prophet model emerges as a strong alternative, balancing predictive accuracy with interpretability with a MAPE value of (12.78%). Traditional statistical models show limitations in capturing the non-linear dynamics and irregular patterns common in urban pedestrian movement. These findings provide actionable insights for model selection in real-world forecasting applications and contribute to the growing body of knowledge in urban analytics and smart city infrastructure development.

Key terms: Time Series Forecasting, Pedestrian Footfall, LSTM, Prophet, SARIMA, Holt-Winters.

1.0 INTRODUCTION

Pedestrian footfall, also referred to as foot traffic, represents the number of individuals entering or moving through a shop, area or urban space during a particular period. Accurate prediction of pedestrian movement supports operational decision-making across multiple sectors. In retail environments, footfall is directly related to conversion rates, defined as the proportion of visitors who complete a purchase and is influenced by factors such as product display, store layout and overall customer experience. In urban governance and smart city applications, predictive analytics enables more effective infrastructure planning, transportation management and public safety enhancement. Advances in sensing technologies, including Wi-Fi tracking, CCTV and AI-based video analytics, thermal sensors, PYRO-Box people counters, beacon technology and facial recognition systems, have significantly increased the availability and granularity of pedestrian movement data, creating opportunities for improved forecasting and behavioural analysis.

Pedestrian footfall time series data exhibit characteristics such as sequential dependencies, seasonality, irregular spikes and non-linear patterns which pose challenges for conventional forecasting approaches. Historically, statistical techniques, including ARIMA (Box & Jenkins, 1976) and SARIMA, have been commonly employed due to their strengths in representing temporal relationships and correlated time-series behaviour. Similarly, Holt-Winters' exponential smoothing has demonstrated effectiveness for seasonal demand forecasting (Bermudez et al., 2007). While these models perform well under relatively stable seasonal conditions, their reliance on assumptions of stationarity and regular cyclic behaviour can limit performance in complex urban environments affected by events, holidays and sudden behavioural shifts.

More recent approaches have sought to address these limitations. Prophet, a decomposable time series model, provides flexibility in handling multiple seasonality, trend changes and holiday effects (Rafferty, 2021; Sharma et al., 2022). Advanced deep learning methods, particularly Long Short-Term Memory (LSTM) networks, can model extended temporal patterns and identify complex, non-linear interactions in sequential datasets (Hochreiter & Schmidhuber, 1997; Lindemann et al., 2021). Comparative evaluations of forecasting models commonly employ metrics such as Root Mean Squared Error (RMSE), Mean Absolute Error (MAE), and Mean Absolute Percentage Error (MAPE) to assess predictive performance (Hyndman & Koehler, 2006).

Despite the availability of high-frequency pedestrian data and the growing range of forecasting techniques, uncertainty remains regarding which models provide the most reliable performance across varying urban conditions. Many prior studies focus on single models or limited datasets, and systematic comparisons across multiple urban locations using consistent evaluation frameworks remain relatively scarce.

To bridge this gap, the study compares the Holt–Winters, SARIMA, Prophet, and LSTM approaches using pedestrian count data from Dublin City Centre. The case study utilises data collected via a network of PYRO-Box passive infrared (PIR) counters deployed across major streets and pedestrian corridors. These devices detect pedestrians by sensing body heat differentials as individuals pass through their field of view, enabling accurate and anonymous counting without capturing personal data. Counts are recorded at hourly intervals and transmitted to central servers managed by Dublin City Council, where validation, cleaning and aggregation processes ensure data reliability prior to publication on open data platforms such as data.gov.ie. By benchmarking statistical and machine learning approaches within a unified framework, this research provides empirical evidence to support model selection in urban analytics, retail planning and smart city forecasting applications.

2.0 LITERATURE REVIEW

Classical Time Series Models

Research on pedestrian movement initially relied on observational studies and manual data collection. Early foundational work highlighted how spatial design, environmental factors and urban infrastructure influence pedestrian behaviour, including variables such as sunlight, seating availability and street layout (Whyte, 1980; Considine, 1976). Although these studies provided valuable qualitative insights into human movement, they lacked quantitative forecasting frameworks necessary for evaluating predictive model performance.

With advances in statistical modelling and data collection technologies in the late twentieth century, researchers began applying formal time series methods to pedestrian and mobility datasets. The introduction of ARIMA models by Box and Jenkins (1976) enabled structured modelling of temporal dependencies, autocorrelations and trends in footfall data. Seasonal extensions through SARIMA further allowed the incorporation of recurring daily, weekly or monthly patterns which are common in pedestrian datasets. Despite their interpretability and statistical rigour, these classical models rely on assumptions of linearity and stationarity, which may reduce predictive accuracy in dynamic urban environments where pedestrian flows exhibit irregular spikes caused by holidays, weather or events (Milenkovic et al., 2018; Majka, 2024)

Holt-Winters exponential smoothing has also been widely applied to pedestrian time series exhibiting trends and seasonal patterns. By decomposing data into level, trend and seasonal components and assigning greater weight to recent observations, the model can adapt to gradual pattern changes over time (Bermudez et al., 2007). Its computational simplicity and interpretability make it attractive for operational forecasting contexts. However, Holt-Winters assumes that trends and seasonalities evolve linearly and remain relatively stable, which may limit performance in pedestrian datasets displaying non-linear behaviour or irregular seasonal fluctuations due to events, holidays or weather variability (Nurhamidah et al., 2020; Sulaiman et al., 2022). Similarly, SARIMA models incorporate seasonal autoregressive and moving average components, providing statistical rigour and interpretability suitable for understanding temporal dependencies in pedestrian flows. Nevertheless, SARIMA requires careful parameter selection and differencing and may underperform when confronted with non-linear dependencies or sudden external shocks common in urban environments (Punyapornwithaya et al., 2021)

Contemporary Forecasting Models

More recent approaches emphasise flexible model-based and machine learning techniques capable of capturing complex temporal patterns. The Prophet model, developed by Facebook, uses an additive decomposition framework to model trend, seasonality and holiday effects. Its ability to incorporate external events and holidays makes it particularly relevant for pedestrian footfall forecasting, where such factors significantly influence traffic patterns (Rafferty,2021; Sharma et al., 2022). Prophet is scalable and user-friendly, allowing rapid fitting on large datasets, although performance may be sensitive to the completeness and accuracy of holiday and event data.

As a specialised recurrent neural network architecture, Long Short-Term Memory (LSTM) networks effectively manage extended temporal dependencies and learn non-linear relationships within sequential data. The architecture consists of forget, input and output gates that regulate information flow, enabling the model to retain or discard information as necessary. This structure allows LSTM models to represent complex pedestrian behaviours, particularly in high-frequency datasets (Hochreiter & Schmidhuber, 1997;

Lindemann et al., 2021). Empirical studies suggest that LSTM networks often outperform classical statistical models in dynamic pedestrian environments. However, they require substantial computational resources and large labelled datasets, and their “black box” nature can reduce interpretability compared with traditional statistical approaches (Murcio & Wang, 2025)

Forecast Evaluation Metrics

Forecasting performance is commonly evaluated using error-based metrics such as Root Mean Squared Error (RMSE), Mean Absolute Error (MAE) and Mean Absolute Percentage Error (MAPE) (Hyndman & Koehler, 2006). RMSE is useful because it penalises large errors more heavily, making it sensitive to substantial forecasting deviations that may occur during peak pedestrian periods or special events. Mean Absolute Error (MAE) provides a straightforward measure of average absolute error, offering interpretability in terms of actual pedestrian counts, which is useful for operational planning. The Mean Absolute Percentage Error (MAPE) measures prediction errors as percentages, enabling relative comparisons across locations or time periods with varying pedestrian counts. In urban pedestrian forecasting, using multiple metrics provides a more comprehensive evaluation of model robustness, accuracy and practical usability.

Research Gap

Although numerous statistical and machine learning models have been used on pedestrian data, few studies provide a systematic comparison of classical and modern approaches. Classical methods such as Holt-Winters and SARIMA provide interpretability and strong statistical grounding, whereas models like Prophet and LSTM offer flexibility and improved capability to capture non-linear dependencies. However, few studies comprehensively compare these approaches within the same urban context using consistent evaluation metrics.

To address this gap, the study performs a thorough assessment of the predictive performance of Holt-Winters, SARIMA, Prophet and LSTM models on pedestrian footfall data from Dublin City Centre with the objective of identifying which time series model delivers the most reliable and accurate forecasts under varying conditions.

3.0 METHODOLOGY

This study adopted a quantitative, comparative time-series research design to evaluate the predictive performance of selected statistical forecasting models on pedestrian footfall data. Secondary data were obtained from Dublin City Council and the National Transport Authority, collected through PYRO-Box infrared pedestrian sensors installed across multiple streets in Dublin City Centre between January 2022 and December 2024. Data is available on the data.gov.ie portal, which provides free access to publicly available datasets published by government departments, local authorities and public bodies. The portal is maintained by the Irish government as part of its Open Data Initiative, encouraging evidence-based decision-making and public engagement. The dataset comprised hourly pedestrian counts from 15 streets.

Data preprocessing involved column standardisation, removal of redundant directional counts, treatment of missing values using forward-fill imputation, and aggregation of hourly observations into daily totals. Stationarity was assessed using the Augmented Dickey–Fuller test, while seasonal patterns were examined using autocorrelation and partial autocorrelation plots. The dataset was divided chronologically into training (80%) and testing (20%) sets to preserve temporal structure.

Four forecasting models, Holt–Winters Exponential Smoothing, Seasonal Autoregressive Integrated Moving Average (SARIMA), Prophet, and Long Short-Term Memory (LSTM) neural networks, were implemented. Model parameters were selected through grid search and information criteria optimisation. Forecast accuracy was evaluated using Mean Absolute Error (MAE), Root Mean Squared Error (RMSE), Mean Absolute Percentage Error (MAPE), and the coefficient of determination (R^2). Comparative performance analysis determined the most suitable statistical approach for pedestrian footfall forecasting.

4.0 RESULTS AND DISCUSSION

Model Performance Overview

Holt-Winters Exponential Smoothing

Model experimentation revealed that the configuration combining a multiplicative trend with additive seasonality (trend='mul', seasonal='add', seasonal periods=7) provided the most accurate results. This configuration is well-suited to the data's characteristics, effectively modelling both the proportionally increasing long-term trend and the recurring weekly patterns. Its performance was evaluated using common time series criteria, summarised below:

Table 1: Holt-Winters Model Performance Metrics

Metric	Value
Mean Absolute Error (MAE)	41,935.71
Root Mean Squared Error (RMSE)	57,222.81
Mean Absolute Percentage Error (MAPE)	16.32%

The Holt-Winters model achieved a Mean Absolute Error (MAE) of 41,935.71 and Root Mean Squared Error (RMSE) of 57,222.81, indicating a moderate average deviation between the predicted and exact footfall values. The Mean Absolute Percentage Error (MAPE) of (16.32%) suggests that the model's forecast differs by about (16%) from the actual values.

Such a level of accuracy is considered reasonable for high-variability time series data, particularly in scenarios involving human behavioural patterns such as daily footfall, where external factors (e.g., holidays, promotions, weather conditions) can introduce significant fluctuations.

The final model was configured using the optimal smoothing parameters determined through the grid search process:

- Alpha (Smoothing of level): 0.2
- Beta (Smoothing of Trend): 0.4
- Gamma (Smoothing of seasonality): 0.7

When updating the level and trend components, these parameter values show that the model gives recent observations a considerable amount of weight, while assigning relatively higher weight to seasonal patterns. This is appropriate for the footfall data under study, where consistent weekly patterns were observed, and recent fluctuations in the trend were less volatile.

Overall, the model successfully captured both the proportional growth trends and the recurring weekly seasonality in the dataset, resulting in reliable short-term forecasts. These forecasts can be effectively used for operational decision-making and resource planning based on expected footfall patterns.

Residual Analysis and Diagnostic Insights

Residual diagnostics were conducted to evaluate the adequacy of the model and identify any remaining systematic patterns not captured by the forecast.

Visual Inspection

Analysis of residual plots, including time series plots, histograms, and Q-Q plots, indicated that the residuals were generally centred around zero and approximately followed a normal distribution. Nevertheless, certain large residual spikes were observed, suggesting the presence of outliers likely corresponding to exceptional events or anomalies not accounted for by the model.

Autocorrelation Analysis

The combined chart (Figure 3.20) provides important diagnostic insights. The ACF plot shows prominent spikes at regular seasonal lags (7, 14, 21 days, etc.), indicating the persistence of unmodeled weekly seasonality in the residuals despite the inclusion of an additive seasonal component. Simultaneously, the PACF plot exhibits a significant spike at lag 1, suggesting that short-term autocorrelation remains, with residuals still influenced by their immediate past values. These patterns highlight that the current model configuration does not fully capture certain temporal dependencies in the data.

Limitations of the Holt-Winters Approach

Although the Holt-Winters approach successfully captured primary patterns in the dataset—namely, the overall multiplicative growth trend and a basic weekly seasonality it has inherent limitations: • The model assumes fixed and regular patterns for both trend and seasonality, which may not adequately reflect evolving or irregular behaviours present in the data.

- It lacks the capacity to explicitly model complex dependencies such as autocorrelated residual structures or varying seasonal effects over time.
- The observed residual patterns indicate that certain systematic behaviours remain unexplained, limiting the potential for further accuracy improvements within the Holt-Winters framework.

Inability to Handle Sudden Structural Changes (Level Shifts)

Around mid-2024, there is a clear downward shift in the actual footfall values. The model does not adapt quickly to this sudden change and continues forecasting at a higher level for some time. This indicates that the model struggles with abrupt changes in the time series, which could be caused by major events such as policy changes, economic disruptions, or external shocks.

Over-Smoothing of Extreme Fluctuations

The model fails to fully capture the sharp peaks and troughs present in the actual data, especially during high-variance periods (e.g., mid to late 2023). This suggests that the chosen smoothing parameters prioritise stability over sensitivity, resulting in underestimation during peaks and overestimation during sudden drops.

Fixed Seasonality Assumption

The model assumes a fixed seasonal pattern (weekly seasonality) across the entire forecast horizon. However, the visualised actual values suggest that the strength and impact of seasonality vary over time, particularly after the structural level shift in 2024. Holt-Winters cannot dynamically model changing seasonal patterns.

Assumption of Constant Trend and Seasonality Type

The model is locked into a multiplicative trend and additive seasonality configuration throughout. While this setup performed reasonably well overall, it may not be suitable for all time periods, particularly when volatility increases or trends shift abruptly.

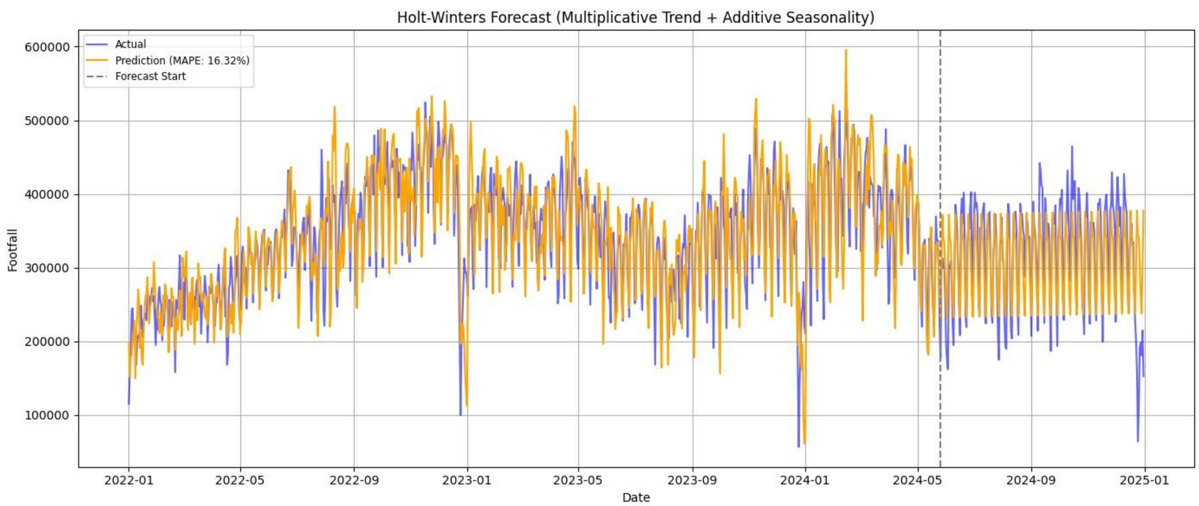

Figure 1: Holt-Winters Model Forecast vs Actual Footfall

Prophet Model

Model Performance Evaluation

The Prophet model, configured with weekly seasonality, was applied to forecast daily footfall patterns using past data.

The model attained a Mean Absolute Error (MAE) of 31,868.63, indicating that the predicted footfall values deviated from the exact observed values by an average of approximately 31,869 footfall units per day. Root Mean Squared Error (RMSE) was computed as 41,487.83, suggesting that while most predictions were close to the actual values, some larger deviations contributed to the overall error. Additionally, the Mean Absolute Percentage Error (MAPE) was recorded at (12.78%), reflecting a reasonably high level of predictive accuracy given the natural variability observed in daily footfall data.

These results indicate that the Prophet model successfully captured the recurring weekly patterns inherent in footfall data. This is largely attributed to the regularity of human activity patterns throughout the week, which are well modelled through the inclusion of weekly seasonality. The model's ability to accurately reproduce historical data patterns further suggests its suitability for short- to medium-term forecasting applications in similar contexts.

Figure 4.2 presents a visual comparison of the forecasted values of the model with the exact footfall data observed. The plot demonstrates that the model successfully tracked the overall trend and seasonality of the data, closely following the fluctuations in daily footfall volumes. While minor deviations between predicted and actual values were observed, particularly during peak demand periods, the model maintained a consistent and reliable predictive performance throughout the evaluation period.

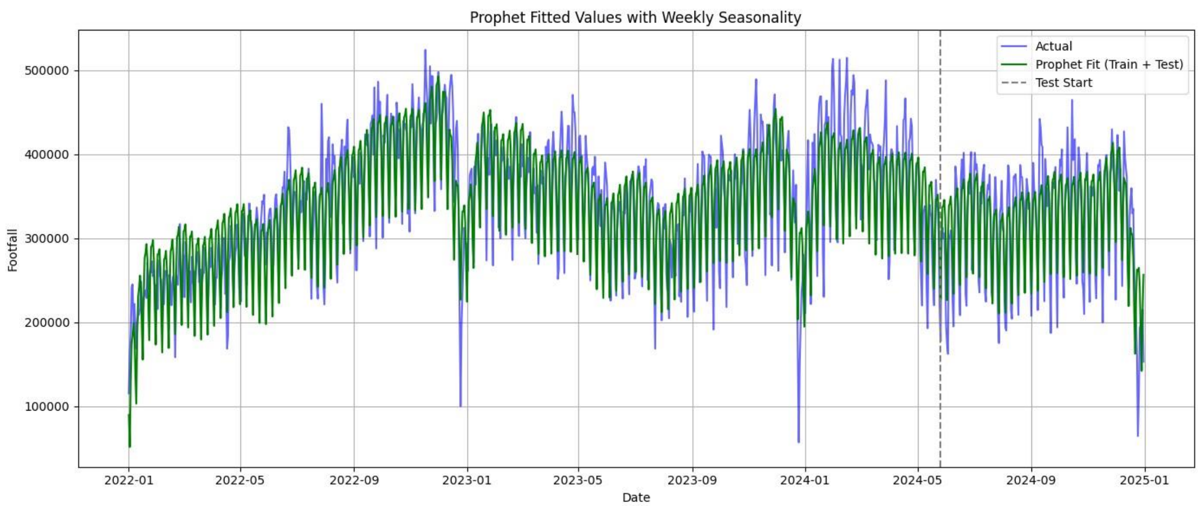

Figure 2: Prophet Model Fitted Values with Weekly Seasonality

In summary, the Prophet model demonstrated satisfactory performance in forecasting daily footfall, accurately capturing both the overall trend and weekly cyclical variations present in the data. These findings validate the applicability of the Prophet model for environments where regular weekly patterns significantly influence footfall behaviour.

Summary of Key Results

The key performance metrics obtained from the evaluation of the Prophet model are summarised in Table 2. These figures provide a quantitative overview of the predictive capabilities of the model.

Table 2: Summary of Prophet Model Performance Metrics

Metric	Value
Mean Absolute Error (MAE)	31,868.63
Root Mean Squared Error (RMSE)	41,487.83
Mean Absolute Percentage Error (MAPE)	12.78%

The table clearly illustrates that the Prophet model is capable of producing reliable forecasts with reasonable error margins, further supporting its suitability for time series forecasting of daily footfall data.

SARIMA Model

Model Performance Evaluation

The SARIMA model configured by parameters order = (1, 0, 1) and seasonal order = (1, 1, 1, 7) was applied to forecast the footfall data. Table 3 summarises the model evaluation results.

Table 3: Evaluation Metrics for SARIMA Model

Metric	Value
Mean Absolute Error (MAE)	35,655.32
Root Mean Squared Error (RMSE)	52,549.05
Mean Absolute Percentage Error (MAPE)	15.28%

The MAE value shows that the forecasts of the model deviated from the actual observed footfall by about 35,655 individuals per day. The RMSE, being higher than the MAE, suggests the presence of some large forecast errors, which may be attributed to occasional extreme fluctuations in footfall data. The MAPE value of (15.28%) reflects a reasonable forecasting accuracy, implying that the mean percentage difference of the predicted and actual values remained within acceptable limits for practical applications.

Figure 4.3 presents a visual comparison of the model forecasts with the actual footfall values. Training data is represented by the blue line, the actual observed values for the test period by the black line, and the model’s forecasts by the red dashed line. Forecast values effectively capture the overall trend and seasonal patterns within the data, particularly the recurring weekly fluctuations. However, some discrepancies between the predicted and actual values are observed, especially during periods of sharp increases or decreases in footfall, which the model was less capable of accurately capturing.

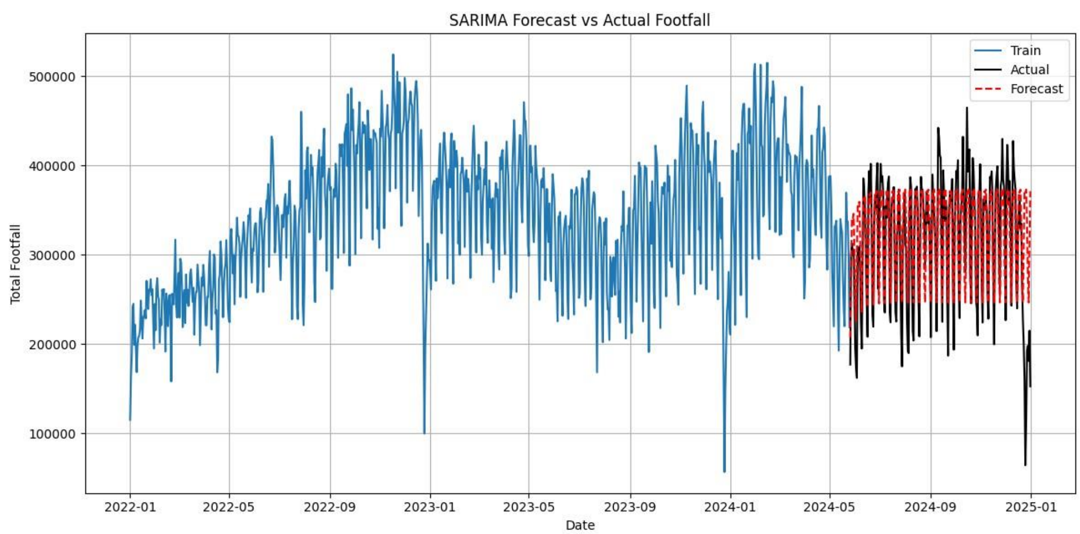

Figure 3: SARIMA Forecast vs Actual Footfall

In conclusion, the SARIMA model demonstrated a satisfactory level of predictive performance for footfall forecasting. Despite some limitations in capturing extreme variations, the model effectively modelled the underlying seasonality and trends present in the dataset, making it an appropriate instrument for forecasting in the short to medium term in similar time series scenarios.

LSTM Model

This section presents the outcomes of the LSTM-based footfall prediction model. The model was developed using a univariate time series composed of historical daily total footfall values. A 30-day lookback window was employed to give the model the potential to learn trends and temporal dependencies. Dropout layers were added to mitigate overfitting, and training was stopped early when validation loss reached a plateau.

4.6.1Quantitative Performance Evaluation

The model's performance was assessed using both absolute and relative error metrics alongside a variance-explained score, with results summarised in Figure 4.

Table 4: Evaluation Metrics for LSTM Model (Combined Train and Test Sets)

Metric	**Value**
Mean Absolute Error (MAE)	26,069.72
Root Mean Squared Error (RMSE)	35,285.11
Mean Absolute Percentage Error (MAPE)	8.82%
R-squared (R2)	0.7778

Mean Absolute Error (MAE): 26,069.72

This shows that, on average, the model's predictions deviate from exact footfall counts by around 26,070. Considering the footfall range (approximately 56,000 to 520,000), this level of error is moderate and acceptable for prediction using time series.

Root Mean Squared Error (RMSE): 35,285.11

The RMSE penalises larger deviations more heavily than MAE. A value of roughly 35,285 implies the model occasionally experiences higher errors during volatile or outlier periods, yet remains effective for general trend tracking.

Mean Absolute Percentage Error (MAPE) (8.82%)

A MAPE under (10%) is considered strong in forecasting contexts. This value reflects that predictions of the model deviate from actual values by less than (9%) on average, demonstrating high relative accuracy across varying scales.

R-squared (R2): 0.7778

The coefficient of determination reveals that the model explains around (77.8%) of the variation in foot traffic. This confirms a strong fit, although with some room for improvement through the inclusion of exogenous variables or advanced model architectures.

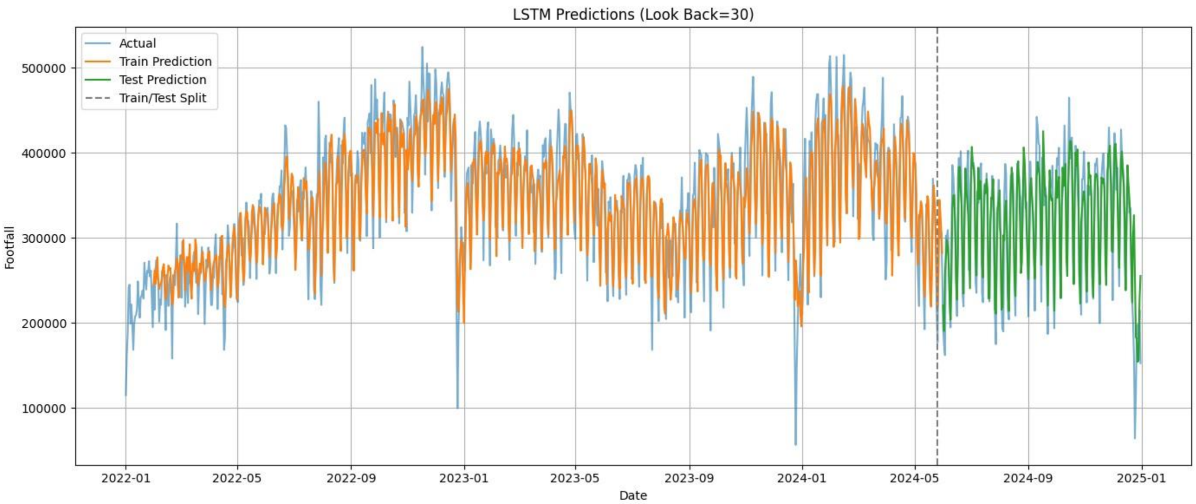

Figure 4: LSTM Predictions vs Actual Daily Footfall (Look Back = 30)

Figure 4 shows that the LSTM model successfully captures both the seasonal patterns and the overall trend in the footfall data. The blue line shows the actual footfall across the entire timeline. The orange line corresponds to predictions during the training period, closely following actual values, especially in capturing regular fluctuations. The green line represents predictions on the testing set, where the model continues to demonstrate good alignment with true values, despite minor deviations during high-variance periods. A vertical dashed line marks the train/test split, highlighting that the model's performance remains stable even on unobserved data. Temporal alignment and preservation of seasonal cycles, particularly weekly patterns, are evident throughout the forecast, further supporting the model's robustness and effectiveness.

Comparative Summary of Model Performance

Table 5: Performance Comparison of Forecasting Models

Model	MAE	RMSE	MAPE (%)	Accuracy (%)	R-squared
Holt-Winters	41,935.71	57,222.81	16.32	83.68	–
Prophet	31,868.63	41,487.83	12.78	87.22	–
SARIMA	35,655.32	52,549.05	15.28	84.72	–
LSTM	26,069.72	35,285.11	8.82	91.18	0.7778

Journal url: https://journals.editononline.com/

Editon Consortium Publishing
Transforming Scholarly Publishing

A thorough analysis of the four prediction models, LSTM, SARIMA, Prophet, and Holt-Winters, was conducted using key performance metrics. Additionally, predictive accuracy was approximated by subtracting MAPE from (100%), providing an intuitive measure of each model's forecasting effectiveness. Results of these metrics are summarised in the comparative Table 5.

Among the models assessed, the LSTM (Long Short-Term Memory) model demonstrated superior performance across all evaluation criteria. It yielded the least MAE, RMSE, and MAPE values, and correspondingly achieved the highest estimated accuracy of approximately (91.18%). This performance indicates that the LSTM model was most effective in minimising prediction error and capturing underlying patterns in the data. Furthermore, the inclusion of the coefficient of determination (R2 = 0.7778) reinforces the LSTM model's robustness, as it suggests that approximately (77.78%) of the variability in the target variable was explained by the model. The model's architecture, which is especially made to identify long-term relationships in sequential data, likely contributed to its superior forecasting capability.

The Prophet model, developed for dealing with time series data with strong seasonality and trend components, ranked second in performance. With an accuracy of (87.22%), it exhibited moderately low error values and provided reasonably reliable forecasts. Although not as precise as the LSTM model, Prophet's ease of use, interpretability, and robustness to missing data make it a viable alternative, particularly in contexts where model transparency is valued.

The SARIMA (Seasonal Autoregressive Integrated Moving Average) model exhibited comparatively weaker performance, with an accuracy of (84.72%). Despite SARIMA's proficiency at identifying seasonal patterns and linear trends, its relatively higher error values suggest limitations in modelling the more complex, non-linear characteristics of the dataset under study. As such, while SARIMA remains a widely used and theoretically grounded statistical model, its predictive efficacy in this context was less favourable.

The Holt-Winters model exhibited the least favourable results among the models evaluated, with the highest MAE, RMSE, and MAPE values, and the lowest estimated accuracy of (83.68%). This model, which relies on exponential smoothing techniques, is insufficient in effectively expressing the complex temporal dynamics and variability found in the data. Its relatively simplistic approach to trend and seasonality modelling likely constrained its forecasting performance.

In conclusion, when it came to forecasting accuracy and error minimisation, the LSTM model performed better than any other model. It is especially well-suited for time series forecasting jobs due to its capacity to represent intricate, sequential patterns involving dynamic and non-linear data structures. The Prophet model served as a competent alternative, offering a balance between interpretability and predictive strength. Both SARIMA and Holt-Winters, while traditionally effective in more stable or linear time series, were comparatively less suited for the characteristics of the dataset used in this study.

5.0 CONCLUSION AND RECOMMENDATIONS

Conclusion: This research conducted a comparative study of time series forecasting models applied to pedestrian footfall data in Dublin City Centre with the objective of evaluating how various models perform in predicting pedestrian traffic patterns. The study implemented and assessed four models: Holt-Winters Exponential Smoothing, SARIMA, Prophet, and LSTM, each representing a distinct class of forecasting methodology, ranging from traditional statistical techniques to deep learning architecture. Model evaluation combined statistical metrics with practical interpretability considerations to determine how effectively each approach represented the intricate temporal dynamics of pedestrian counts.

The findings revealed key differences in predictive capability across the models. The LSTM neural network demonstrated the strongest predictive power, achieving an R-squared of 0.7778 and showing a strong ability to learn long-term dependencies and non-linear fluctuations in pedestrian behaviour. The Prophet model, while slightly less accurate on certain measures, performed competitively in terms of RMSE and offered advantages in transparency, modularity, and ease of deployment. In contrast, traditional models such as SARIMA and Holt-Winters, although foundational and interpretable, exhibited limitations due to their reliance on assumptions of linearity and fixed seasonality assumptions that often do not fully capture the dynamic and irregular patterns observed in urban pedestrian movement.

Overall, the comparative evaluation confirms that model selection significantly influences forecasting performance, with advanced machine learning approaches demonstrating enhanced capacity to model complex temporal structures in pedestrian footfall data. From a practical perspective, several actionable recommendations emerge from this study. For urban planners and city authorities, the results suggest that the LSTM model can be valuable for operational decision making where high predictive accuracy is required, such as infrastructure planning, event preparedness and crowd management. Models such as Prophet strike a balance between accuracy and interpretability for routine forecasting tasks in public-sector environments, where transparency and ease of implementation are important. To optimise both usability and predictive accuracy, hybrid models integrating interpretable statistical approaches with advanced machine learning techniques should be explored.

For future research, incorporating external explanatory variables such as weather conditions, tourism patterns and public events could further improve predictive accuracy and provide deeper insight into the drivers of pedestrian behaviour. Evaluating model performance across multiple urban locations and longer time horizons would strengthen generalizability and support the development of scalable smart-city forecasting frameworks.

6.0 REFERENCES

1. Bermudez, J. D., Segura, J. V., & Vercher, E. (2007). Holt-Winters forecasting: An alternative formulation applied to UK air passenger data. *Journal of Applied Statistics, 34*(9), 1075–1090.
2. Box, G. E. P., & Jenkins, G. M. (1976). *Time series analysis: Forecasting and control* (Revised ed.). Holden-Day.
3. Considine, D. M. (Ed.). (1976). *Van Nostrand's scientific encyclopedia* (5th ed.). Van Nostrand Reinhold Company.
4. Hochreiter, S., & Schmidhuber, J. (1997). Long short-term memory. *Neural Computation, 9*(8), 1735–1780.
5. Hyndman, R. J., & Koehler, A. B. (2006). Another look at measures of forecast accuracy. *International Journal of Forecasting, 22*(4), 679–688.
6. Lindemann, B., Müller, T., Vietz, H., Jazdi, N., & Weyrich, M. (2021). A survey on long short-term memory networks for time series prediction. *Procedia CIRP, 99*, 650–655.
7. Majka, M. (2024, September). Seasonal time series analysis: Why SARIMA outshines ARIMA. *Solutio Care.* https://www.researchgate.net/publication/384196885

8. Milenković, M., Švadlenka, L., Melichar, V., Bojović, N., & Avramović, Z. (2018). SARIMA modelling approach for railway passenger flow forecasting. *Transport, 33*(5), 1113–1120.
9. Murcio, R., & Wang, Y. (2025). Identify optimal pedestrian flow forecasting methods in Great Britain retail areas: A comparative study of time series forecasting on a footfall dataset. *ISPRS International Journal of Geo-Information, 14*(2), Article 50. https://doi.org/10.3390/ijgi14020050
10. Nurhamidah, Wan, N., & Faisol, A. (2020). Forecasting seasonal time series data using the Holt-Winters exponential smoothing method of additive models. *Jurnal Matematika Integratif, 16*(2), 151–157.
11. Punyapornwithaya, V., Jampachaisri, K., Klaharn, K., & Sansamur, C. (2021). Forecasting of milk production in northern Thailand using seasonal autoregressive integrated moving average, error trend seasonality, and hybrid models. *Frontiers in Veterinary Science, 8*, 775114.
12. Rafferty, G. (2021). *Forecasting time series data with Facebook Prophet: Build, improve, and optimize time series forecasting models using the advanced forecasting tool.* Packt Publishing Ltd.
13. Sharma, K., Bhalla, R., & Ganesan, G. (2022). Time series forecasting using FB Prophet. In *Proceedings of the Algorithms, Computing and Mathematics Conference 2022* (pp. 59–65). CEUR Workshop Proceedings, 3445. https://ceur-ws.org/Vol3445/PAPER_07.pdf
14. Sulaiman, A. I., Alfa, M. S., & Auwal, A. M. (2022). Application of support vector machines and Holt-Winters exponential smoothing modelling approaches in airlines passengers' time series forecasting. *African Journal of Advances in Science and Technology Research, 6*(1), 38–51.
15. Whyte, W. H. (1980). *The social life of small urban spaces.* Project for Public Spaces.