Physical Culture and Sport. Studies and Research, 2025, 111, 25–35
https://doi.org/10.2478/pcssr-2025-0027

Time-Series Forecasting in Sports: Using LSTM and GRU for Stadium Attendance Prediction

Authors' contribution:
A) conception and design of the study
B) acquisition of data
C) analysis and interpretation of data
D) manuscript preparation
E) obtaining funding

Yu Pang *A-E

National Research University Higher School of Economics, St Petersburg, Russia

Received: 21.02.2025
Accepted: 11.04.2025

***Correspondence:** Yu Pang, 3 Kantemirovskaya St., Building 1A, St Petersburg, Russia: yupang@hse.ru

Abstract

This study investigates the use of Long Short-Term Memory (LSTM) and Gated Recurrent Unit (GRU) models for predicting stadium attendance in National Football League (NFL) games. Using a comprehensive dataset spanning 26 years and incorporating various game-specific, economic, and temporal features, the performance of LSTM and GRU architectures in forecasting attendance rates is compared. The analysis reveals that both models effectively capture underlying patterns in attendance data, with the LSTM model demonstrating slightly superior predictive accuracy. Optimal configurations are identified through a comparative evaluation of different hidden sizes and layer counts. The best-performing models achieve a Root Mean Squared Error (RMSE) of 7.81% and a Mean Absolute Error (MAE) of 5.62%, representing a significant improvement over previous approaches. While the LSTM model exhibits better adaptability to sudden variations in attendance, the GRU model offers faster convergence and more consistent predictions. To enhance transparency, SHapley Additive exPlanations (SHAP) were used to interpret model outputs. The results revealed seasonality, team identity, and economic indicators as the most influential factors, aligning with domain knowledge and supporting practical applications. These findings contribute to the growing field of AI-driven demand forecasting in sports, offering valuable insights for financial decision-making and operational planning in the sports industry.

Keywords: Stadium attendance prediction, long short-term memory, gated recurrent unit, time-series forecasting, financial decision-making in sports

Introduction

Context and Challenges in Stadium Attendance Forecasting

Stadium attendance plays a crucial role in the financial and operational success of professional sports events (Schreyer & Ansari, 2021). The number of spectators attending games directly affects revenue streams, including ticket sales, concessions, sponsorship value, and overall fan engagement (Borland & MacDonald, 2003). From a management perspective, accurate attendance forecasting is essential to optimize resource allocation, security planning, and infrastructure management. Furthermore, attendance trends provide valuable insights into fan behavior, team performance, and market dynamics, making predictive analytics a powerful tool for sports organizations and policymakers. However, predicting attendance remains a challenging task because of a multitude of influencing factors, such as team performance, scheduling, economic conditions, and external disruptions (Karg et al., 2021; Spenner et al., 2004). Given these complexities, using advanced data-driven approaches to improve forecasting accuracy is of academic and practical significance.

Predicting demand for sports, particularly stadium attendance, has long been a central focus of sports economics research. Traditionally, econometric models have been used to forecast attendance by incorporating factors, such as team performance, ticket prices, weather, and socioeconomic conditions (Baimbridge et al., 1996; Hall et al., 2010; Karg et al., 2021; Paul et al., 2021). For example, previous studies have highlighted the importance of team success, competitive balance, and star players in driving attendance at professional sports events (Borland & MacDonald, 2003). Additionally, traditional regression-based models often assume stationarity and linear relationships, which limits their ability to capture long-term temporal dependencies and dynamic fan behaviors. Consequently, these conventional methods may fail to provide accurate and robust attendance forecasts, particularly in highly volatile sports environments. These challenges underscore the growing need for more advanced, data-driven forecasting techniques capable of modeling the intricate and evolving nature of stadium attendance.

From Traditional Models to Deep Learning: A Methodological Shift

Recent advances in artificial intelligence and machine learning have transformed demand forecasting across industries, introducing architectures that go beyond traditional ML. In addition to ensemble methods such as Random Forest and XGBoost, recent years have witnessed the development of hybrid models that combine Convolutional Neural Networks (CNNs) with LSTM or Transformer-based architectures to capture both spatial and temporal dimensions in complex forecasting tasks (Sun et al., 2023; Yeung et al., 2025). In sports economics, machine learning techniques, such as Support Vector Machines (SVM), random forests, and XGBoost, have been increasingly applied to predict stadium attendance, often outperforming traditional regression-based models. These methods utilize pattern recognition, feature selection, and ensemble learning to enhance predictive accuracy. For instance, these techniques have demonstrated high accuracy in forecasting attendance for the Chinese Football Association Super League (Du et al., 2022), as well as for major American professional sports leagues, such as the NFL, NBA, and MLB (Mueller, 2020; Şahin & Uçar, 2022). In previous work by the authors (Pang & Wang, 2024), a suite of ensemble models—including Random Forest, CatBoost, and XGBoost—was applied to over 5,000 regular-season NFL games. These models achieved strong predictive performance and identified key attendance drivers such as stadium capacity, unemployment rate, and personal income. However, while effective in identifying static associations, the models lacked the ability to capture dynamic temporal structures inherent in attendance trends.

Moreover, tree-based models often assume independence between observations and lack internal memory, making them ill-suited for capturing carry-over effects or long-term fan behavior trends across seasons (Zhou, 2021). When applied to highly dynamic contexts, such as sports attendance, they tend to underperform compared to models that account for temporal structure. However, traditional econometric approaches often assume linearity and stationarity, which limits their ability to model the nonlinear and evolving nature of stadium attendance. To address these challenges, Artificial Neural Networks (ANNs) have been employed in sports analytics, and studies have demonstrated their effectiveness in forecasting stadium attendance in football matches (Şahin & Erol, 2018). ANN-based models have also been applied to major U.S. leagues, including the NBA, MLB, and NFL, as well as the Korean Baseball League (Park & Park, 2017; Şahin & Uçar, 2022).

However, compared to ANNs, Recurrent Neural Networks (RNNs) have gained prominence in time-series forecasting because of their ability to process sequential data and effectively model long-term dependencies. Among these, Long Short-Term Memory (LSTM) networks, a specialized form of RNN, have been particularly successful in overcoming the vanishing gradient problem, allowing them to learn long-range dependencies that are crucial for attendance prediction.

The application of advanced RNN architectures, such as LSTM and Gated Recurrent Units (GRUs), has proven useful in many areas (Abumohsen et al., 2023; Busari & Lim, 2021; M. Cho et al., 2022). However, it remains an emerging field in sports economics, particularly in predicting stadium attendance. These models offer substantial potential for enhancing prediction accuracy by accommodating sequential data, capturing complex nonlinear relationships, and recognizing intricate temporal patterns. Although research in this field remains limited, studies have demonstrated the value of attention-based LSTM networks in sports analytics, particularly in match-result prediction (Zhang et al., 2022). However, the direct application of these models to stadium attendance forecasting remains underexplored and presents a promising avenue for future research.

This Study: Addressing Gaps with Temporal Deep Learning and Explainability

To address the limitations of traditional and machine learning-based attendance forecasting models, this study applies two deep recurrent architectures, Long Short-Term Memory (LSTM) and Gated Recurrent Units (GRU), to predict stadium attendance using a real-world dataset from the National Football League (NFL). The dataset spans 26 seasons and includes over 7,000 game-level observations with comprehensive team, match, stadium, and economic features.

Unlike prior studies, which have primarily relied on static or shallow models, this research utilizes the sequential learning capabilities of RNNs to capture temporal dependencies that are critical in attendance behavior. To enhance interpretability—a known limitation of deep learning models—SHapley Additive exPlanations (SHAP) are incorporated, allowing detailed assessment of feature contributions in both the LSTM and GRU models.

In addition, this study extends previous machine learning efforts by not only comparing model performance against ensemble-based approaches reported in prior literature (e.g., Pang & Wang, 2024), but also by providing a transparent evaluation of how model decisions are influenced by key inputs. This dual focus on prediction and explainability ensures both empirical robustness and practical interpretability.

To the best of our knowledge, this is the first study to systematically apply LSTM and GRU models to stadium attendance forecasting in a major professional sports context, while integrating model explainability tools. As such, it makes both a methodological and applied contribution to the literature on sports economics, offering new insights into the drivers of fan attendance and the utility of explainable AI in event demand forecasting.

Model and Data

Data

The dataset includes an imbalanced panel of 37 NFL teams, covering 7,221 observations from home games played each season over 26 years (1994–2019). Games held in neutral international venues, including Super Bowl Games, were excluded. Attendance and performance data for the National Football League were obtained from the 'Pro Football Reference' website (https://www.pro-football-reference.com/). The target variables are attendance rates, which are between 0 and 1 and align more closely with the characteristics of the neural data (Şahin & Erol, 2018). Given the long time span of the dataset, stadium capacities have changed significantly over the years. To ensure consistency in calculating attendance rates, each stadium's maximum recorded capacity was used as the basis for standardization. However, some attendance rates in the dataset exceeded 1.0. This is because, during high-demand events, standing-room-only areas are opened, allowing attendance to surpass 100% of the nominal capacity. Key match details are encoded as integer values, including the home team, away team, and playoff status. Each record included seasonal and monthly information to provide additional background information on the games. All games are arranged in ascending order by date. Because NFL games typically occur from September to January, the month feature was transformed using sine and cosine functions to handle the cyclical nature of the data and minimize the artificial gap between December and January. To capture the impact of game outcome uncertainty on attendance, two key betting metrics, the Over/Under Line and the Spread, were incorporated into the model (Paul & Weinbach, 2011). Additionally, team performance indicators were also included, specifically the season-to-date winning percentages of both the home and away teams prior to each game. Several key variables were included to capture the influence of stadium characteristics on attendance (DeSchriver et al., 2016; Paul et al., 2021; Wakefield & Sloan, 1995). The stadium name is encoded as an integer to distinguish between the different venues. Each stadium's maximum capacity is included to reflect its size and potential attendance limits, along with its age, to account for variations in popularity, accessibility, and amenities over time. These variables enable the model to consider how stadium-specific factors affect attendance patterns. The distance between the locations of the two teams was included to capture the geographical impact on attendance. Economic conditions are represented by the annual per capita income of the state's Metropolitan Statistical Area (MSA), adjusted to the 1982 Consumer Price Index (CPI) base year (1982=100) as an indicator of real income levels. Additionally, the state unemployment rate is used to reflect the economic fluctuations that may affect discretionary spending on event attendance. This rate was reported on a monthly basis and adjusted for seasonal variations to provide a more accurate measure of economic context. Table 1 summarizes the variables used in the model along with their descriptions and sources.

Table 1 Description of Variables Used in the Model

Variable Name	Description	Source
Attendance Rate	Ratio of reported attendance to maximum stadium capacity	Pro Football Reference
Home Team	Integer-encoded ID of home team	Pro Football Reference
Away Team	Integer-encoded ID of away team	Pro Football Reference
Season (year)	Year of the game season (1994–2019)	Pro Football Reference
Week of Season	Week number of the NFL season (1–17)	Pro Football Reference
Month (sine/cosine)	Cyclical representation of month to capture seasonality	Engineered
Playoff Status	1 = playoff game, 0 = regular season	Pro Football Reference
Over/Under Line	Predicted total score line set by bookmakers	Betting data (Paul & Weinbach)
Point Spread (Favorite)	Handicap for favored team used to balance odds	Betting data
Win Rate (Home)	Win percentage of home team before game	Calculated from match history
Win Rate (Away)	Win percentage of away team before game	Calculated
Stadium Capacity	Maximum recorded capacity of each stadium	Pro Football Reference
Stadium Age	Age of the stadium at time of game	Calculated
Stadium (ID)	Encoded stadium identifier	Engineered
Travel Distance	Distance between home and away team cities	Calculated
Real Income	Per capita income in the MSA, adjusted to 1982 CPI	Bureau of Economic Analysis
Unemployment Rate	Monthly, seasonally adjusted unemployment rate by state	Bureau of Labor Statistics

Long Short-Term Memory (LSTM) and Gated Recurrent Unit (GRU)

LSTM networks are a special type of RNN designed to overcome the vanishing gradient problem, which limits the ability of traditional RNNs to learn the long-term dependencies in time-series data (Hochreiter & Schmidhuber, 1997). The LSTM architecture is characterized by a unique memory cell that can selectively retain or forget information through three types of gates: input, forget, and output.

- Input Gate: Regulates the incorporation of new inputs into the cell state.
- Forget Gate: Determines the degree to which the previous cell state should be discarded or maintained.
- Output Gate: Modulates the cell output to the subsequent hidden state.

These gates allow the LSTM to effectively remember long-term information while discarding irrelevant data. This capability makes LSTM particularly powerful in scenarios where the output at a given time step depends on a combination of short- and long-term factors, such as a series of sports events. In this study, an LSTM model is used to predict future stadium attendance based on past game features, team performance, and external economic indicators. The GRU is a simplified variant of LSTM that combines the forget and input gates into a single gate, referred to as the update gate (K. Cho et al., 2014). GRU has fewer parameters than LSTM, which often makes it computationally more efficient while retaining the ability to capture long-term dependencies in the data. The GRU architecture includes the following.

- Update Gate: Determine how much of the previous information to carry forward and how much new information to incorporate.
- Reset Gate: Controls how much of the previous hidden state should be forgotten when computing a new candidate hidden state.

The simpler structure of the GRU makes it faster to train, particularly when the dataset is large or when the computational resources are limited. Despite being less complex, GRU has been shown to perform comparably to LSTM in various time-series prediction tasks, including demand forecasting and financial time-series analysis. Given the similarities between the GRU and LSTM, the GRU was used as a baseline to compare its performance against the more complex LSTM model in the context of stadium attendance prediction.

Training and Evaluation Procedure

Both the models were implemented in Python using the PyTorch framework. The data were normalized using MinMax scaling, with 90% of the dataset used for training and the remaining 10% used for testing. The models were trained for 100 epochs with early stopping based on validation loss to prevent overfitting.

- Mean Squared Error (MSE): This measures the average squared difference between predicted and actual attendance, highlighting the model's accuracy and sensitivity to larger errors.
- Root Mean Squared Error (RMSE): This provides an indication of the accuracy of the model by taking the square root of the MSE, bringing the error back to the original scale of attendance.
- Mean Absolute Error (MAE): Calculate the average absolute difference between predicted and actual attendance, providing a straightforward measure of the average error.
- Mean Absolute Percentage Error (MAPE): This reflects the model's prediction accuracy as a percentage, offering an intuitive interpretation of the forecast error relative to the actual values.

By comparing the performance of LSTM and GRU on both the training and test datasets, the goal was to identify the model that provides the best predictive accuracy for stadium attendance. The results of this comparative analysis offer valuable insights into the trade-offs between model complexity and forecasting performance in sports economics.

Empirical Results

In this section, the performance of the LSTM and GRU models in predicting stadium attendance is analyzed. To provide a comprehensive assessment, various configurations, including different hidden sizes and layer counts, were examined and evaluated using the Mean Squared Error (MSE) for both the training and test sets. The results are presented in a series of tables and figures, highlighting the impact of each parameter on model accuracy and the comparative effectiveness of the LSTM and GRU architectures. Insights into the convergence behavior, prediction accuracy, and error distributions of both models are also provided, enabling a detailed understanding of their strengths and limitations in this forecasting task.

Table 2 Performance comparison with varying hidden sizes and layers

Model	Hidden size	Number of layers	Train MSE	Test MSE
LSTM	50	2	0.0090	0.0066
	80		0.0087	0.0070
	100		0.0087	0.0069
	50	3	**0.0087**	**0.0061**
		4	0.0088	0.0061
		5	0.0088	0.0061
GRU	50	2	0.0089	0.0079
	80		0.0084	0.0070
	100		0.0086	0.0071
	50	3	**0.0084**	**0.0069**
		4	0.0085	0.0075
		5	0.0089	0.0069

Table 2 provides a comparative analysis of the LSTM and GRU models for different hidden sizes (50, 80, and 100) and layer counts (2, 3, 4, and 5). For both models, increasing the number of hidden layers tended to slightly improve the test set MSE up to a point before reaching a plateau or even showing a minor increase in error. The LSTM model achieved its best test set performance with a hidden size of 50 and three layers, yielding a test MSE of 0.0061, suggesting that a moderate configuration balances the model complexity and generalization. Larger configurations, such as four or five layers, showed no significant improvement, and may even lead to overfitting, as indicated by the higher training MSE. Similarly, the GRU model achieved its best test performance with a hidden size of 50 and 3 layers, resulting in a test MSE of 0.0069. GRU, despite being computationally less complex than LSTM, demonstrated a similar performance when using smaller configurations. However, as the number of hidden layers and their sizes increased, the GRU exhibited a minor decline in test performance. These results support the notion that less complex configurations tend to yield better generalizations for a particular dataset.

Table 3 presents the performance metrics of the best-performing configurations of the LSTM and GRU models, as listed in Table 2. Each model configuration was selected based on its optimal performance, with both models set to a hidden size of 50, with three layers. The RMSE of 0.0781 for the LSTM model indicated that, on average, the predictions deviated from the actual attendance by approximately 7.81 percentage points. An MAE of 0.0562 provides a straightforward measure of the average absolute error, indicating that the predictions are, on average, off by 5.62 percentage points. A MAPE of 6.42% indicates that the model's predictions deviate, on average, by 6.42% from the actual attendance, offering an intuitive understanding of the relative error. These metrics reveal that the LSTM model achieves a slightly lower RMSE and MAPE, indicating marginally better overall predictive accuracy. In contrast, the GRU model exhibited a lower MAE of 0.0546, suggesting that it may perform slightly better in terms of absolute prediction error.

Table 3. Comparative performance of LSTM and GRU models on the test set

Model configuration	Test RMSE	Test MAE	Test MAPE
LSTM (hidden size = 50; number of layers = 3)	0.0781	0.0562	6.42%
GRU (hidden size = 50; number of layers = 3)	0.0832	0.0546	6.43%

Figure 1 presents a comparison of the actual stadium attendance and predicted values from the LSTM and GRU models for the test set. The blue line represents the actual attendance, whereas the orange and green lines represent the predictions made by the LSTM and GRU models, respectively. Both models demonstrated strong performance in capturing the general trend of actual attendance values, with slight deviations. The LSTM model exhibited better adaptability for high-variance events and accurately captured peaks and troughs in the attendance data. By contrast, the GRU model produces smoother predictions, making it more consistent but less responsive to sharp fluctuations in attendance. These results suggest that the LSTM model is better suited for scenarios with high variability in attendance, offering nuanced handling of sudden changes. In contrast, the GRU model, while ideal for stable patterns, tends to be slightly underpredicted during highly variable events. This aligns with its inherent design, which favors less variability in predictions and a more stable output.

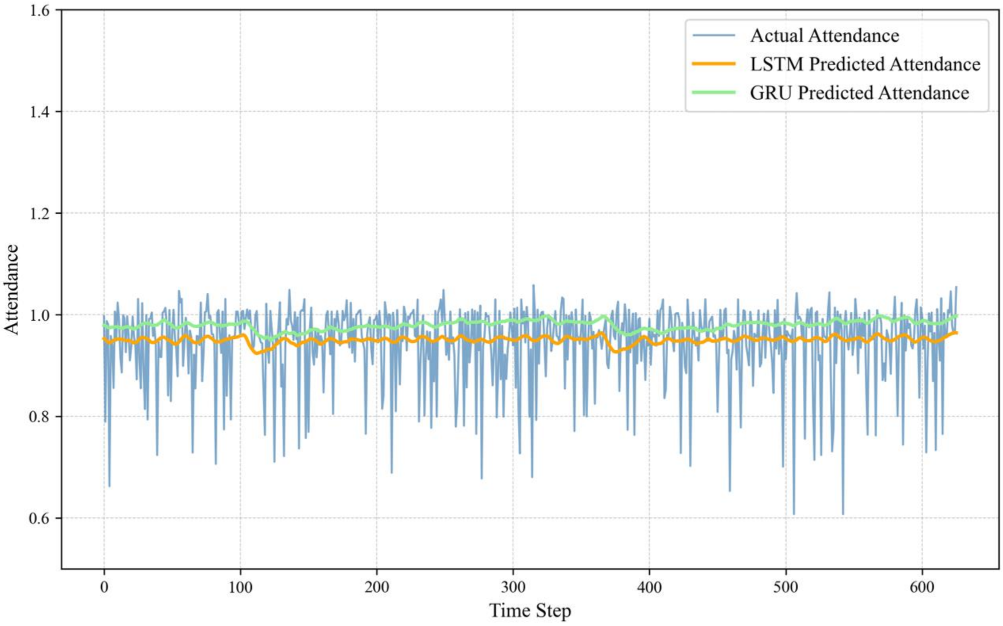

Figure 1. Predicted vs. actual stadium attendance for both models on the test set

Figure 2 illustrates the error density distribution for the LSTM and GRU models, reflecting the frequencies of the different prediction errors in the test set. The orange curve represents the error density of the LSTM model, and the green curve corresponds to the GRU model. Both distributions peaked near zero, indicating that the majority of the predictions for both models had low errors. The GRU model exhibited a narrower and taller peak, suggesting more consistent performance with smaller and less variable errors. By contrast, the LSTM model displayed a broader distribution, capturing a wider range of errors, including more extreme values. These findings imply that the GRU model is better suited to scenarios requiring consistent predictions with minimal variability, whereas the LSTM model, with its broader error distribution, may be better equipped to handle high-variance scenarios and capture extreme fluctuations.

Figure 3 illustrates the training loss versus epochs for the LSTM and GRU models. Both models exhibited rapid reductions in training loss during the initial epochs, highlighting the efficient learning in the early stages of training. The GRU model demonstrated a faster initial convergence, with a steep decline in loss within the first 10 epochs, reaching near-zero loss slightly earlier than the LSTM model. By approximately 20 epochs, both models stabilized, showing minimal fluctuations in the training loss. This comparison indicates that, while the GRU model offers a speed advantage in convergence, the LSTM model catches up over prolonged training. Ultimately, both models achieved similar levels of loss, demonstrating comparable training efficiency and effectiveness in minimizing errors in the training set. The choice between the two models may depend on whether faster convergence or capture of complex long-term dependencies is prioritized.

To enhance the interpretability of the deep learning models, we employed SHAP (SHapley Additive exPlanations) to analyze feature contributions to the prediction outcomes of both the LSTM and GRU architectures. Figure 4 (LSTM) and Figure 5 (GRU) display the SHAP summary plots. In both models, Season (year) and Month (sine encoded) demonstrate the highest explanatory power, suggesting that temporal dynamics are key drivers of stadium attendance. Features such as Home Team, Away Team Win Rates, and Real Personal Income also appear consistently among the top contributors, reflecting the role of team performance and regional economic conditions.

Interestingly, Stadium Capacity exhibits higher importance in the GRU model (Figure 5), suggesting that the GRU architecture may be more sensitive to venue-related constraints when predicting attendance. In contrast, in the LSTM model (Figure 4), features such as Week of Season and Over/Under Line carry slightly more weight, indicating a greater emphasis on temporal and betting-related dynamics.

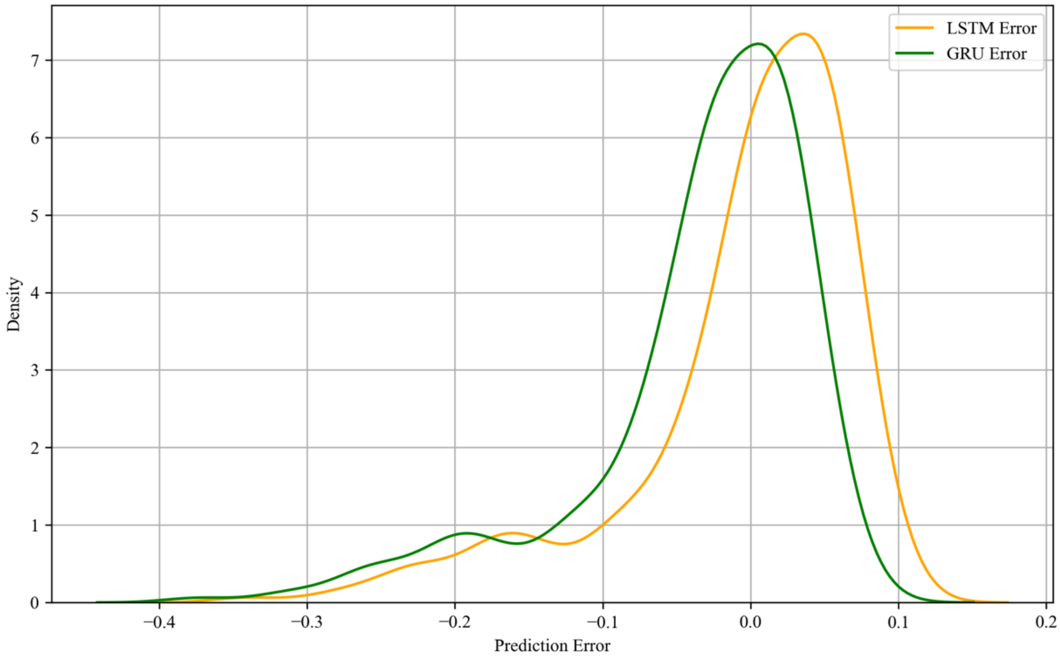

Figure 2. Distribution of prediction errors for both models

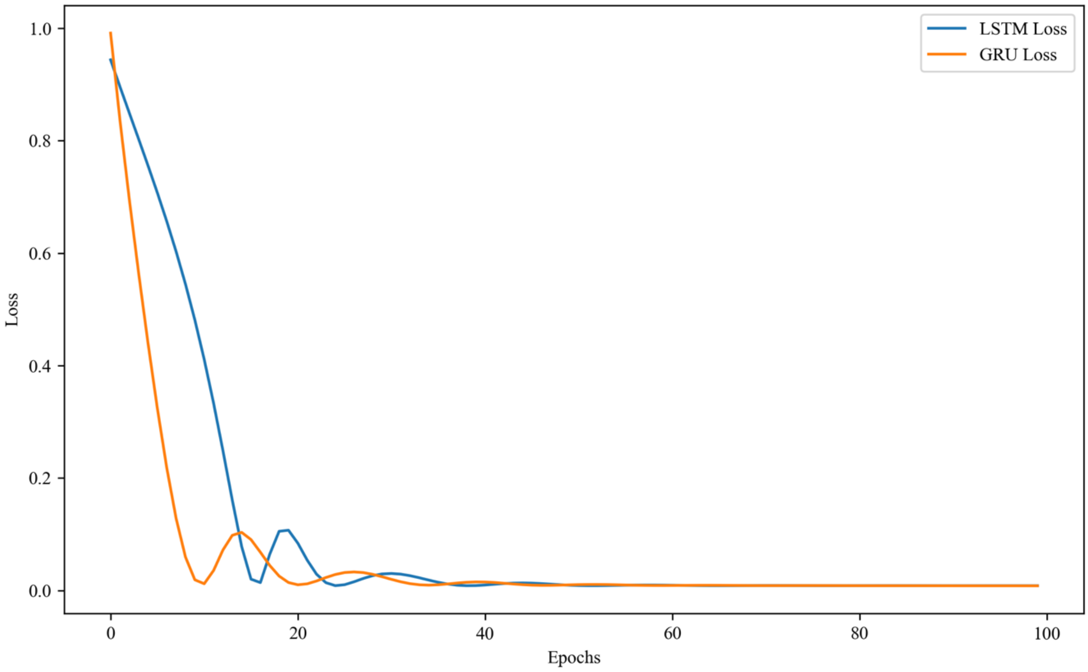

Figure 3. Training loss vs. epochs for both models

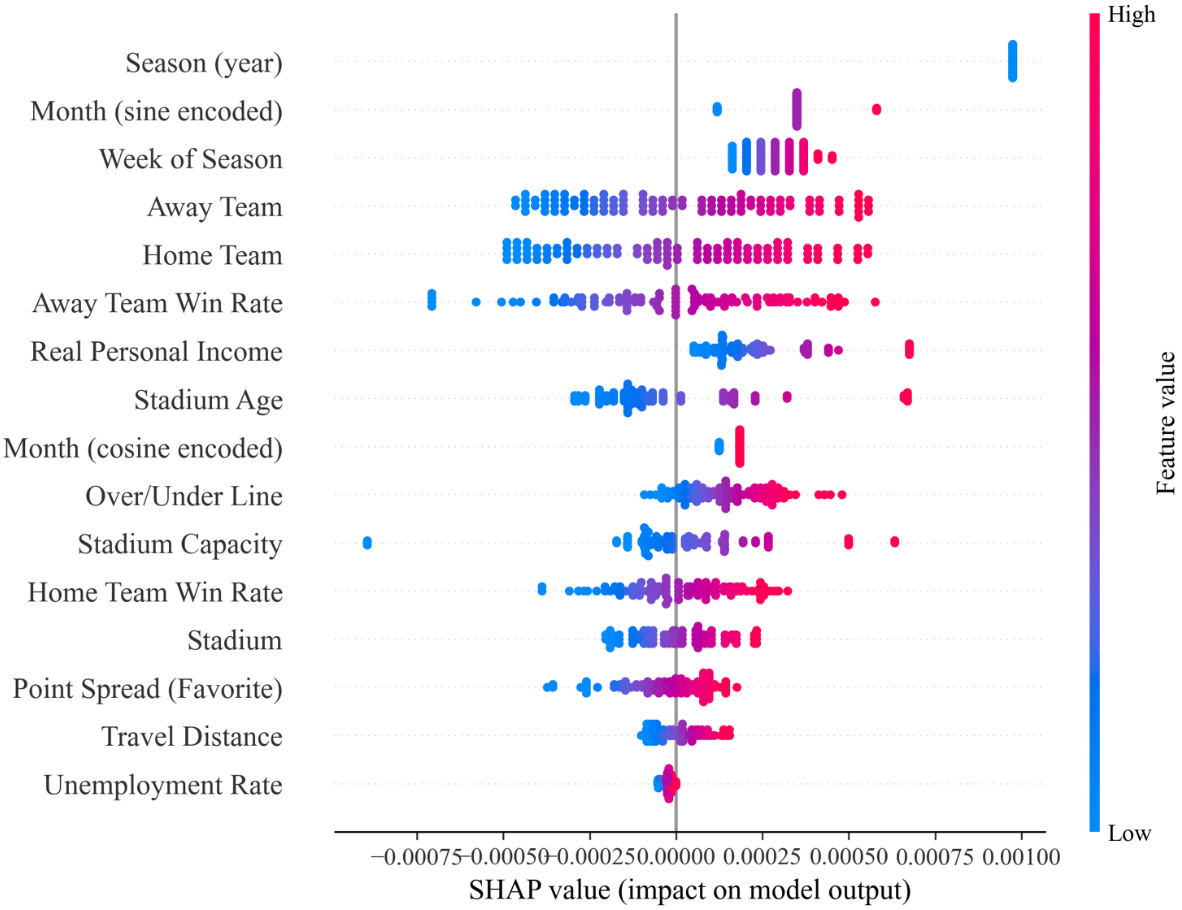

Figure 4. Feature Contributions (SHAP) – LSTM Model

These findings are consistent with results obtained in previous work using tree-based machine learning models, where features such as stadium age, club age, personal income, and unemployment rate were identified as key predictors (Pang & Wang, 2024). The convergence of important features across different modeling paradigms and interpretability techniques demonstrates the robustness of the selected variable set.

Overall, the SHAP analysis confirms that the models are learning meaningful patterns, with importance scores aligning with domain knowledge (e.g., team strength, match timing, and economic context). The consistency of key features across models also supports the robustness of the feature set.

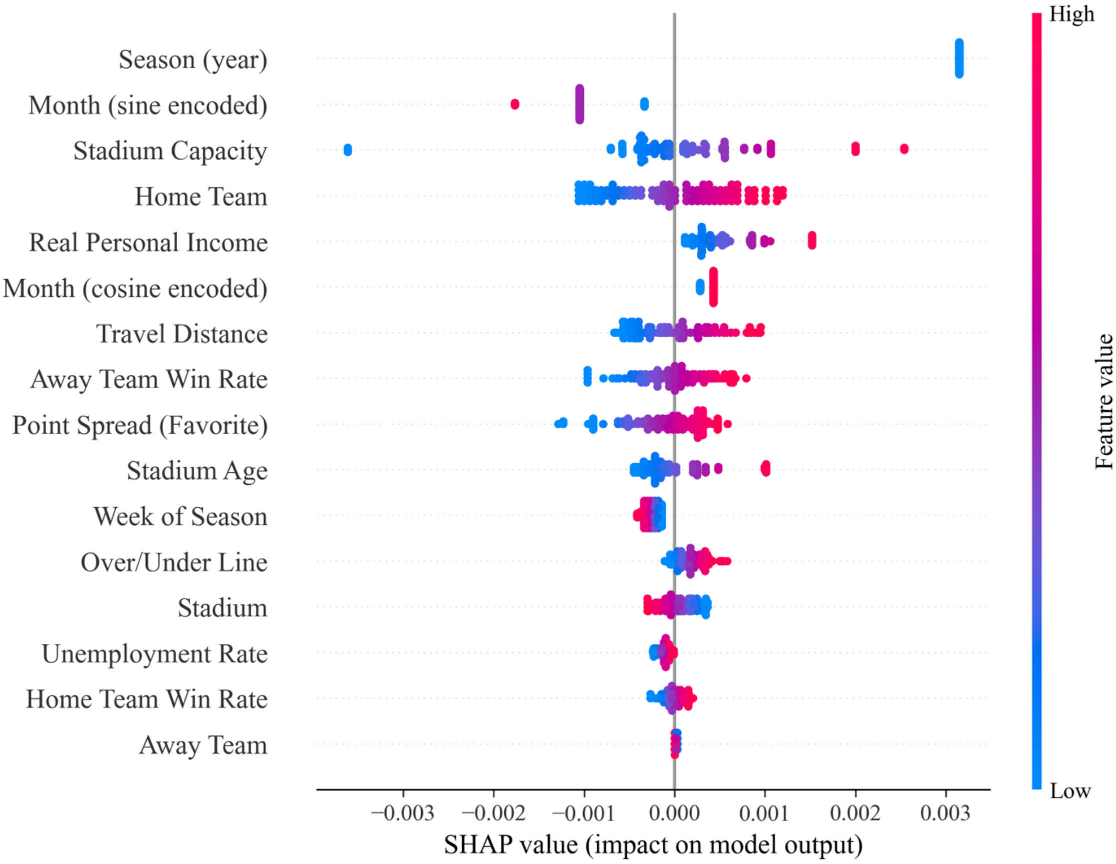

Figure 5. Feature Contributions (SHAP) – GRU Model

Conclusion

This study investigated the use of Long Short-Term Memory (LSTM) and Gated Recurrent Unit (GRU) models to predict stadium attendance in NFL games, offering valuable insights for financial decision-making in sports economics. By integrating a range of game-specific, economic, and temporal features, the analysis demonstrated that both models effectively captured the underlying patterns in attendance data, with the LSTM model showing a slight improvement in predictive accuracy. A comparative evaluation of different hidden sizes and layer counts revealed that moderate configurations strike an optimal balance between the model complexity and performance. Both models achieved their lowest Mean Squared Error (MSE) on the test set with a hidden size of 50 and three layers. The results indicated a strong predictive performance with a Root Mean Squared Error (RMSE) of 7.81% and a Mean Absolute Error (MAE) of 5.62%, which is suitable for practical applications. These findings represent a significant improvement in accuracy compared with the Artificial Neural Network (ANN) results reported in prior studies (Şahin & Erol, 2018; Şahin & Uçar, 2022). While the predictive accuracy was slightly lower than in prior ensemble-based models (Pang & Wang, 2024), the deep learning models provided substantial advantages in capturing temporal structures and dynamic attendance patterns.

Further analyses of the error density and residuals highlighted the strengths and limitations of each model: the LSTM model exhibited superior adaptability to sudden variations in attendance, whereas the GRU model delivered consistent predictions with faster convergence. In addition to performance metrics, this study incorporated SHapley Additive exPlanations (SHAP) to interpret the models' decisions. The SHAP analysis revealed that seasonality, team identity, and economic indicators were the most influential features, offering greater transparency into the model's behavior and aligning with established domain knowledge. These insights highlight the trade-off between raw predictive power and interpretability, and reinforce the value of deep learning models in contexts where temporal reasoning is critical. This research contributes to the growing field of AI-driven demand forecasting in sports, showing how advanced machine-learning techniques can enhance prediction accuracy and support financial planning in high-stake environments.

Future studies could explore the integration of additional features, such as weather conditions or ticket pricing, and examine the applicability of these models across other sports leagues and event types. Continued refinement of these predictive models has the potential to provide sports organizations with powerful tools for making informed financial and operational decisions. In addition, recent advances in deep learning—particularly Transformer-based architectures that rely on attention mechanisms—have demonstrated strong performance in time-series forecasting due to their ability to model long-range dependencies and complex feature interactions. While this study focused on standard recurrent models (LSTM and GRU), future research may consider incorporating Transformer-based models or hybrid architectures to further enhance prediction accuracy and interpretability, especially when multimodal data sources become available.

Ethics approval and informed consent

This study does not involve human subjects, human data, or animal experiments

Competing interests

No potential conflict of interest exists among the authors.

Funding

There was no external funding in this study.

Acknowledgments

The author would like to acknowledge the support of the China Scholarship Council (CSC).

References

Abumohsen, M., Owda, A. Y., & Owda, M. (2023). Electrical load forecasting using LSTM, GRU, and RNN algorithms. *Energies*, *16*(5), 2283.

Baimbridge, M., Cameron, S., & Dawson, P. (1996). Satellite television and the demand for football: A whole new ball game? *Scottish Journal of Political Economy*, *43*(3), 317–333.

Borland, J., & MacDonald, R. (2003). Demand for sport. *Oxford Review of Economic Policy*, *19*(4), 478–502.

Busari, G. A., & Lim, D. H. (2021). Crude oil price prediction: A comparison between AdaBoost-LSTM and AdaBoost-GRU for improving forecasting performance. *Computers & Chemical Engineering*, *155*, article 107513.

Cho, K., Van Merriënboer, B., Gulcehre, C., Bahdanau, D., Bougares, F., Schwenk, H., & Bengio, Y. (2014). *Learning phrase representations using RNN encoder-decoder for statistical machine translation*. ArXiv Preprint arXiv:1406.1078.

Cho, M., Kim, C., Jung, K., & Jung, H. (2022). Water level prediction model applying a long short-term memory (LSTM)–gated recurrent unit (GRU) method for flood prediction. *Water*, *14*(14), 2221.

DeSchriver, T. D., Rascher, D. A., & Shapiro, S. L. (2016). If we build it, will they come? Examining the effect of expansion teams and soccer-specific stadiums on Major League Soccer attendance. *Sport, Business and Management: An International Journal*, *6*(2), 205–227. https://doi.org/10.1108/SBM-05-2014-0025

Du, P., Wang, Y., Liao, C., & Xian, T. (2022). *Sports games attendance forecast using machine learning*. 2022 IEEE 2nd International Conference on Data Science and Computer Application (ICDSCA). https://doi.org/10.1109/ICDSCA56264.2022.9987748

Hall, J., O'Mahony, B., & Vieceli, J. (2010). An empirical model of attendance factors at major sporting events. *Special Issue on Event Studies*, *29*(2), 328–334. https://doi.org/10.1016/j.ijhm.2009.10.011

Hochreiter, S., & Schmidhuber, J. (1997). Long short-term memory. *Neural Computation*, *9*(8), 1735–1780.

Karg, A., Nguyen, J., & McDonald, H. (2021). Understanding season ticket holder attendance decisions. *Journal of Sport Management*, *35*(3), 239–253.

Mueller, S. Q. (2020). Pre- and within-season attendance forecasting in major league baseball: A random forest approach. *Applied Economics*, *52*(41), 4512–4528. https://doi.org/10.1080/00036846.2020.1736502

Pang, Y., & Wang, F. (2024). Forecasting stadium attendance using machine learning models: A case of the national football league. *Studia Sportiva*, *18*(2), 147–164.

Park, J., & Park, S. (2017). A study on prediction of attendance in Korean Baseball League using artificial neural network. *KIPS Transactions on Software and Data Engineering*, *6*(12), 565–572.

Paul, R. J., Ehrlich, J. A., & Losak, J. (2021). Expanding upon the weather: Cloud cover and barometric pressure as determinants of attendance for NFL games. *Managerial Finance*, *47*(6), 749–759. https://doi.org/10.1108/MF-06-2020-0295

Paul, R. J., & Weinbach, A. P. (2011). NFL bettor biases and price setting: Further tests of the Levitt hypothesis of sportsbook behaviour. *Applied Economics Letters*, *18*(2), 193–197. https://doi.org/10.1080/13504850903508242

Şahin, M., & Erol, R. (2018). Prediction of attendance demand in European football games: Comparison of ANFIS, fuzzy logic, and ANN. *Computational Intelligence and Neuroscience*, *1*, article 714872. https://doi.org/10.1155/2018/5714872

Şahin, M., & Uçar, M. (2022). Prediction of sports attendance: A comparative analysis. *Proceedings of the Institution of Mechanical Engineers, Part P: Journal of Sports Engineering and Technology*, *236*(2), 106–123. https://doi.org/10.1177/1754337120983135

Schreyer, D., & Ansari, P. (2021). Stadium attendance demand research: A scoping review. *Journal of Sports Economics*, *23*, 749–788. https://doi.org/10.1177/15270025211000404

Spenner, E. L., Fenn, A., & Crooker, J. (2004). The demand for NFL attendance: A rational addiction model. *Colorado College Economics and Business Working Paper No. 2004–01*. https://doi.org/10.2139/ssrn.611661

Sun, X., Wang, Y., & Khan, J. (2023). Hybrid LSTM and GAN model for action recognition and prediction of lawn tennis sport activities. *Soft Computing*, *27*(23), 18093–18112. https://doi.org/10.1007/s00500-023-09215-4

Wakefield, K. L., & Sloan, H. J. (1995). The effects of team loyalty and selected stadium factors on spectator attendance. *Journal of Sport Management*, *9*(2), 153–172. https://doi.org/10.1123/jsm.9.2.153

Yeung, C., Sit, T., & Fujii, K. (2025). Transformer-based neural marked spatio temporal point process model for analyzing football match events. *Applied Intelligence*, *55*(5), 335. https://doi.org/10.1007/s10489-024-05996-9

Zhang, Q., Zhang, X., Hu, H., Li, C., Lin, Y., & Ma, R. (2022). Sports match prediction model for training and exercise using attention-based LSTM network. *Digital Communications and Networks*, *8*(4), 508–515.

Zhou, Z.-H. (2021). *Machine learning*. Springer Nature.