Proceedings of the 2nd World Congress on Industrial Engineering and Operations Management
Windsor, Canada, October 14-16, 2025

Proceedings of the International Conference on Industrial Engineering and Operations Management

Publisher: IEOM Society International, USA DOI: 10.46254/WC02.20250130
Published: October 14, 2025

Predictive Modeling of Patient Waiting Times in the Emergency Room Using Supervised Learning

M M Aflatun Kawsar and Md Shamim Hasan
Department of Industrial and Production Engineering
Military Institute of Science & Technology
Dhaka, Bangladesh
fahim4547@gmail.com and shamimhasan.mist@gmail.com

Turjo Ghoshal
Department of Mechanical Engineering
University of Delaware
Delaware, USA
tghoshal@udel.edu

Abstract

The optimization of patient flow within hospital emergency rooms (ERs) continues to be a significant challenge in healthcare operations. The protracted waiting times directly impact the patient satisfaction along with efficient resource allocation. This study presents a robust machine learning framework for the accurate prediction of ER waiting times. We identified and engineered key predictive features including patient acuity, staffing ratios, and temporal factors such as time of day and seasonality by leveraging a dataset of 5,000 patient encounters. Our methodology encompassed a comparative analysis of various machine learning algorithms including linear models, tree-based ensembles, and support vector machines to ascertain the optimal predictive architecture. Model performance was assessed via a rigorous 5-fold cross-validation process, employing Mean Absolute Error (MAE), Root Mean Squared Error (RMSE), and the coefficient of determination (R^2). The findings reveal that a hyperparameter-tuned Stacked Ensemble Model yielded superior performance by achieving coefficient of determination (R^2) of 0.9473. This outcome indicates a strong correlation between the model's predictions and the observed wait times, thus validating its ability to encapsulate the intricate and nonlinear dynamics inherent in ER operations. The developed model offers a validated data-driven tool for hospital administrators facilitating proactive resource management and ameliorated patient communication. This study highlights the capacity of advanced analytics to address significant operational inefficiencies within healthcare, presenting a trajectory towards an enhanced patient experience and more efficacious allocation of clinical resources.

Keywords
Emergency Room, Waiting Times, Predictive Modeling, Machine Learning, Healthcare Operations

1. Introduction
Hospital emergency rooms (ERs) function as the critical ingress point for a significant portion of healthcare systems, operating under constant pressure to deliver timely and effective care to patients with a wide spectrum of conditions (Nyce et al., 2021). The escalating demand for emergency services, coupled with operational constraints, frequently leads to systemic strain, manifesting as overcrowding and significant delays ("Overcoming Common Challenges in the ED | Abbott Point of Care," n.d.). These conditions place immense stress on clinical caregivers and represent a major patient safety concern ("Overcoming Common Challenges in the ED | Abbott Point of Care," n.d.). Central to

these operational challenges is the pervasive issue of protracted patient waiting times, a problem that transcends mere inconvenience to become a critical determinant of healthcare quality and patient outcomes.

The consequences of excessive waiting times are severe and multifaceted. From a patient perspective, long waits are a primary driver of dissatisfaction, fundamentally shaping their perception of service quality and undermining trust in the healthcare institution (Nyce et al., 2021). More alarmingly, these delays are directly correlated with adverse clinical outcomes. Delays in diagnosis and treatment can lead to the exacerbation of symptoms, necessitate longer recovery periods, and in the most severe cases, result in permanent disability or mortality ("The Risks Associated with Long ER Wait Times | Physicians Premier ER," n.d.). The phenomenon of patients leaving without being seen (LWBS) has nearly doubled in recent years, a direct consequence of intolerable waits that signifies a critical failure in care delivery ("The Risks Associated with Long ER Wait Times | Physicians Premier ER," n.d.) This reframes the problem of wait times not as a simple operational metric to be minimized, but as a direct clinical risk factor that must be managed with the same rigor as any other threat to patient safety. Operationally, long waits are a clear indicator of systemic inefficiency, leading to suboptimal resource allocation, staff burnout, and significant financial repercussions in competitive healthcare environments (Alrasheedi et al., 2019).

In response to these challenges, a paradigm shift is occurring in healthcare operations, moving away from reactive, experience-based management toward proactive, data-driven decision-making ("AI in Healthcare Operations Management: Optimizing Efficiency and Care | Calonji," n.d.). The widespread adoption of Electronic Health Records (EHRs) has created a wealth of granular data, making it feasible to apply advanced analytical and algorithmic approaches to optimize hospital workflows and resource management ("A Comprehensive Guide to Machine Learning in Healthcare," n.d.) However, many ERs continue to rely on rudimentary methods for estimating wait times, such as rolling averages or median estimators, which have been shown to have limited accuracy and fail to capture the dynamic nature of ER demand. The persistent failure of these incumbent methods has created a technology vacuum, highlighting an urgent need for a more sophisticated and reliable predictive solution. The current systems are not just suboptimal; they are fundamentally inadequate for the complexity of the task.

This study directly addresses this critical gap by developing and validating a robust supervised machine learning framework for the accurate prediction of ER patient waiting times. The objectives of this research are to: (i) develop and validate a robust supervised machine learning framework for the accurate prediction of ER patient waiting times using a real-world dataset of patient encounters; (ii) conduct a rigorous comparative analysis of multiple machine learning architectures, including linear models, tree-based ensembles, and support vector machines; (iii) demonstrate the superior predictive performance of a hyperparameter-tuned Stacked Ensemble Model in capturing the complex, nonlinear dynamics of ER operations; and (iv) provide a validated, data-driven tool for hospital administrators to facilitate proactive resource management, optimize patient flow, and improve patient communication.

1.1 Objectives

This study aims to:

- Develop and evaluate various supervised learning models for predicting ER patient waiting times.
- Conduct a comparative analysis of the performance of these models' using metrics such as Mean Absolute Error (MAE), Root Mean Squared Error (RMSE), and R-squared (R^2).
- Identify the most influential features that contribute to ER waiting times through feature engineering and selection.
- Propose a predictive modeling framework that can be integrated into ER operational workflows to improve patient flow and resource allocation.

2. Literature Review

The optimization of patient flow through an emergency room is a complex operational challenge, involving the coordination of numerous interdependent processes from patient arrival and triage to diagnosis, treatment, and final disposition, whether admission or discharge (Kumar and Prasad, 2025). Disruptions to this flow are common and result from a confluence of factors, including unpredictable fluctuations in patient volume, the variable mix of patient acuity, dynamic staffing levels, and bottlenecks in downstream resources such as diagnostic imaging or the availability of inpatient beds ("Hospitals - Timely & effective care | Provider Data Catalog," n.d.) Understanding these intricate operational dynamics is the foundational step toward developing effective predictive models.

The application of machine learning (ML) and artificial intelligence (AI) has demonstrated transformative potential across various domains of healthcare operations management, establishing a strong precedent for its use in the ER context ("AI in Healthcare Operations Management: Optimizing Efficiency and Care | Calonji," n.d.). These technologies have been successfully deployed to forecast patient demand and corresponding supply needs, enabling more efficient inventory management and capacity planning (Kumar and Prasad, 2025). Furthermore, ML algorithms are used to optimize the allocation of critical resources—including clinical staff, treatment rooms, and specialized equipment—by matching availability to anticipated patient volumes and acuity levels ("AI in Healthcare Operations Management: Optimizing Efficiency and Care | Calonji," n.d.). By automating routine administrative tasks and streamlining workflows, these systems also reduce the cognitive and administrative burden on clinicians, allowing them to focus on direct patient care (Kumar and Prasad, 2025).

Within the specific domain of ER wait time prediction, academic literature reveals a clear trajectory of increasing methodological sophistication. Early research efforts often employed traditional statistical methods, such as linear regression, which provided foundational insights but were limited by their underlying assumptions of linearity and exhibited higher error rates when applied to the complex, non-linear environment of the ER (Wang et al., 2025). More recent studies have embraced a diverse array of machine learning algorithms to better capture these dynamics. Researchers have explored models such as Q-Lasso regression, Support Vector Machines (SVM), Random Forests (RF), and various Deep Learning architectures, including Deep Neural Networks (DNN) and Long Short-Term Memory (LSTM) networks (Wang et al., 2025). A recurrent theme in this body of work is the superior performance of ensemble methods, which consistently demonstrate higher accuracy by combining the strengths of multiple individual models (Ameur et al., 2022). This progression from simple linear models to complex ensembles is not merely a reflection of advancing computational techniques; it mirrors the research community's growing recognition that ER patient flow is not a simple, linear process but a complex adaptive system. The failure of simpler models and the success of more sophisticated ones provide empirical evidence of the system's inherent non-linearity, justifying the selection of advanced architectures that are philosophically aligned with the nature of the problem. However, a systematic review of this field has identified a critical gap: while predictive accuracy has improved, the areas of feature engineering and explainable AI remain underexplored(Wang et al., 2025) This creates a "last mile" problem, where a model may be highly accurate but is not adopted because its decision-making process is opaque to hospital administrators, hindering trust and actionable implementation.

To address these challenges and build upon the state of the art, this study focuses on an advanced ensemble learning technique known as stacking. Ensemble learning, in general, is a powerful machine learning paradigm that improves predictive performance by combining multiple models, or classifiers, to produce a single, more robust prediction (Mahajan et al., 2023). This approach has proven highly effective in a variety of complex healthcare applications, such as disease prediction, where ensembles frequently outperform any single constituent model (Mahajan et al., 2023). Stacking, or stacked generalization, is a particularly sophisticated ensemble method. It involves a multi-level architecture where a diverse set of base models (Level-0 learners) are trained on the data, and their predictions are then used as input features to train a second-level model (a meta-learner). The meta-learner's function is to learn the optimal way to combine the predictions of the base models, effectively leveraging their individual strengths while mitigating their weaknesses (Sultan et al., 2025). This provides a strong theoretical justification for its selection as the primary modeling architecture in this study, as it is designed to capture the multifaceted and complex relationships that govern ER wait times.

3. Predictive Modeling Framework

The methodology employed in this study was designed to be rigorous and systematic, encompassing data acquisition and feature engineering, a comparative analysis of multiple modeling architectures, and a robust evaluation protocol to ensure the validity and generalizability of the findings.

3.1 Data and Feature Engineering

The foundation of this research is a de-identified dataset comprising 5,000 patient encounters from a hospital emergency room. The predictive power of any machine learning model is fundamentally dependent on the quality and relevance of its input features. Therefore, a critical phase of this study involved the identification and engineering of features that encapsulate the primary drivers of ER wait times. This process serves as the essential translation layer, converting qualitative clinical realities and contextual operational states into a quantitative language that algorithms can process and learn from. The success of the final model is as much a testament to this effective encoding of domain knowledge as it is to the sophistication of the algorithm itself. The engineered features were grouped into three key categories.

- Patient Acuity: This feature quantifies the clinical urgency of a patient's condition upon arrival. It is derived from standard triage protocols, such as the five-level Emergency Severity Index (ESI), which is widely used to prioritize patients (Dong and Bullard, 2023). Under this system, patients are categorized on an ordinal scale from 1 to 5. A designation of ESI Level 1 signifies a critical, life-threatening condition requiring immediate, life-saving intervention, while ESI Level 5 represents a non-urgent case that can safely wait for treatment (Dong and Bullard, 2023) . This acuity level is a primary determinant of a patient's position in the treatment queue and the intensity of resources they will require.
- Staffing Ratios: To capture the dynamic availability of clinical resources, this feature was calculated as the ratio of on-duty clinical staff (e.g., physicians and nurses) to the total number of patients currently in the ER ("Staff-to-Patient Ratio: A Critical KPI for Your Healthcare Practice," n.d.). A lower ratio indicates a higher workload per staff member and is a direct measure of departmental strain. The literature confirms a strong inverse relationship between staffing ratios and patient wait times, as well as a correlation with other adverse patient outcomes and staff burnout ("Staff-to-Patient Ratio: A Critical KPI for Your Healthcare Practice," n.d.).
- Temporal Factors: ER patient flow exhibits strong cyclical patterns that must be accounted for in any predictive model. To capture these periodic fluctuations in demand, several temporal features were engineered from the patient arrival timestamp. These include the time of day (hour), the day of the week, and the season or month of the year. These factors are known to correlate strongly with both the volume of patient arrivals and the specific mix of case types presenting to the ER.

3.2 Comparative Modeling Architectures

To establish a robust performance benchmark and demonstrate the relative efficacy of our proposed model, a comparative analysis was conducted using a suite of standard supervised learning algorithms. The selected models represent different algorithmic families, each with distinct strengths and underlying assumptions:

- Linear Models: A standard linear regression model was included to serve as a baseline, representing the performance of a simple, interpretable model that assumes linear relationships between the features and the target variable.
- Tree-Based Ensembles: A Random Forest model was implemented. This algorithm is known for its high performance, robustness to overfitting, and its ability to capture complex, non-linear interactions between features without extensive data preprocessing.
- Support Vector Machines (SVM): A Support Vector Regressor (SVR) was included, representing a powerful class of kernel-based methods that are effective in high-dimensional feature spaces and can model non-linear relationships.

3.3 The Stacked Ensemble Model

The primary architecture advanced in this study is a Stacked Ensemble, also known as stacked generalization. This is an advanced ensemble technique designed to achieve superior performance by intelligently combining the predictions of multiple models. The architecture consists of two levels:

- Level-0 (Base Learners): In the first level, a diverse set of machine learning models is trained independently on the full training dataset. This set typically includes different types of algorithms (e.g., a Random Forest, a Gradient Boosting machine, and a Support Vector Regressor) to ensure a variety of "perspectives" on the data.
- Level-1 (Meta-Learner): The predictions generated by each of the Level-0 base learners are then collected and used as the input features for a second-level model, known as the meta-learner. A relatively simple model, such as a linear regressor, is often used as the meta-learner. Its task is not to predict the original target variable directly from the original features, but rather to learn the optimal combination of the base learners' predictions to produce the final, more accurate output (Sultan et al., 2025). This hierarchical process allows the model to learn from the errors of its constituent parts and synthesize their strengths.

3.4 Model Evaluation Protocol

A rigorous evaluation protocol was implemented to ensure that the performance estimates are reliable, unbiased, and indicative of how the models would perform on new, unseen data.

- 5-Fold Cross-Validation: The primary technique for model validation was 5-fold cross-validation. In this procedure, the entire dataset is randomly partitioned into five subsets, or "folds," of equal size. The model is then trained and evaluated five times. In each iteration, one of the folds is held out as the validation set, while the remaining four folds are used for training. The performance score from each iteration is recorded, and the final performance metric for the model is the average of these five scores (Baturynska and Martinsen, 2021) . This approach provides a more stable and robust estimate of model performance than a single train-test split, as it ensures that every data point is used for both training and validation exactly once, thereby mitigating the risk of overfitting and selection bias.
- Performance Metrics: The performance of each model was quantified using three standard metrics for regression tasks, each providing a different perspective on the model's error profile:
- Mean Absolute Error (MAE): This metric is the average of the absolute differences between the predicted wait times and the actual observed wait times. MAE is expressed in the same units as the target variable (minutes), making it highly interpretable. It represents the average magnitude of error, and because it does not square the errors, it is less sensitive to large, anomalous prediction errors (outliers) ("Mean Absolute Error In Machine Learning: What You Need To Know - Arize AI," n.d.). The formula for MAE is:

$$\mathrm{MAE} = \frac{1}{n}\sum_{i=1}^{n} |y_i - \hat{y}_i|$$

  where n is the number of samples, y_i is the actual value, and $\hat{y}_i$ is the predicted value.
- Root Mean Squared Error (RMSE): This metric is the square root of the average of the squared differences between predicted and actual values. By squaring the errors before averaging, RMSE gives disproportionately higher weight to large errors. This makes it a particularly useful metric in contexts where large prediction errors are especially undesirable and should be heavily penalized ("RMSE Explained: A Guide to Regression Prediction Accuracy | DataCamp," n.d.) . The formula for RMSE is:

$$\mathrm{RMSE} = \sqrt{\frac{1}{n}\sum_{i=1}^{n} (y_i - \hat{y}_i)^2}$$

- The decision to use both MAE and RMSE is a deliberate diagnostic choice. A significant divergence between a model's RMSE and MAE scores indicates the presence of high-variance errors; that is, the model may be accurate on average (low MAE) but occasionally produces very large, unacceptable errors (contributing to a high RMSE). For a clinical application, this provides a more complete assessment of the model's real-world reliability.
- Coefficient of Determination (R^2): This metric represents the proportion of the variance in the dependent variable (wait time) that is predictable from the independent variables (features). It provides a measure of how well the model's predictions approximate the real data points, with a value of 1 indicating a perfect fit and a value of 0 indicating that the model performs no better than simply predicting the mean of the target variable. The formula for
  R^2 is:
  $R^2 = 1 - (SSE / SST)$
  Where:
  • $SSE = \Sigma (y_i - \hat{y}_i)^2$ is the Sum of Squared Errors.
  • $SST = \Sigma (y_i - \bar{y})^2$ is the Total Sum of Squares.

4. Results and Discussion

The empirical results from the comparative analysis of the predictive models are presented and discussed in this section. The findings validate the central hypothesis of the study and offer significant practical implications for the management of emergency room operations.

4.1 Comparative Performance of Predictive Models

The performance of each machine learning model, as evaluated through the 5-fold cross-validation protocol, is summarized in Table 1. The results provide a clear and quantitative comparison across the three key performance metrics: Mean Absolute Error (MAE), Root Mean Squared Error (RMSE), and the Coefficient of Determination (R^2). The analysis reveals a distinct hierarchy of performance among the tested architectures. The baseline Linear Regression model, while providing some predictive capability beyond a simple average, demonstrated the highest error rates and the lowest explanatory power. The Support Vector Machine and Random Forest models offered substantial improvements, with the Random Forest emerging as the best-performing single model. However, the hyperparameter-tuned Stacked Ensemble Model significantly outperformed all other contenders across every evaluation metric. It achieved the lowest MAE and RMSE, indicating the highest average accuracy and the smallest large errors. Most notably, it yielded a coefficient of determination (R^2) of 0.9473, a result that signifies an exceptionally strong correlation between the model's predictions and the actual observed wait times (Table 1).

Table 1. Comparative Performance of Machine Learning Models

Model	Mean Absolute Error (MAE) (minutes)	Root Mean Squared Error (RMSE) (minutes)	Coefficient of Determination (R^2)
Linear Regression	35.41	48.15	0.6871
Support Vector Machine (SVM)	21.89	30.52	0.8735
Random Forest	14.23	20.11	0.9218
Stacked Ensemble Model	10.80	15.64	0.9473

4.2 Discussion of Findings

The empirical results strongly support the efficacy of a stacked ensemble architecture for predicting ER wait times. The superior performance of this model is not arbitrary but is rooted in its fundamental design. By hierarchically combining the outputs of diverse base learners, the stacked model can learn a more complex and nuanced decision boundary. It effectively synthesizes the strengths of different algorithms; for instance, the robustness of tree-based methods and the non-linear mapping capabilities of kernel methods—thereby mitigating the individual biases and weaknesses of any single approach (Mahajan et al., 2023). This architectural advantage allows it to more accurately encapsulate the "intricate and nonlinear dynamics inherent in ER operations," as stated in the abstract.

The achievement of an R^2 value of 0.9473 is a particularly profound finding. In a system as complex and often perceived as chaotic as an emergency room, this high level of explanatory power suggests that patient wait times are not a product of randomness but are largely a deterministic outcome of a set of measurable variables. This realization has powerful managerial implications, as it reframes the challenge of ER management from one of "coping with chaos" to one of "engineering a predictable system." If the system's behavior can be known with high accuracy, it can be controlled and optimized.

The practical implications of deploying such a validated, high-accuracy model are transformative for hospital administration. The model serves as a powerful decision-support tool, enabling a shift from a reactive to a proactive operational posture.

- Proactive Resource Management: By providing accurate forecasts of impending periods of high patient volume and long wait times, the model allows administrators to take preemptive action. This can include adjusting staffing schedules to match anticipated demand, opening designated overflow capacity before bottlenecks occur, or strategically diverting low-acuity patients to alternative care settings. This proactive management smooths patient flow prevents the onset of dangerous overcrowding, and ensures that resources are allocated more efficiently.
- Ameliorated Patient Communication: A validated predictive model allows the ER to provide patients with dynamic, data-driven estimates of their expected wait time upon arrival. This simple act of transparent communication can significantly reduce patient anxiety, manage expectations, and improve overall satisfaction scores. It transforms waiting from a period of frustrating uncertainty into a managed and understood part of the care process.

Furthermore, the successful implementation of this predictive model can initiate a virtuous cycle of continuous improvement. Better predictions lead to more effective resource management, which in turn leads to smoother patient flow and reduced wait times. A system operating with less extreme variance and fewer bottlenecks is inherently more stable and, consequently, easier to predict. Therefore, the model does not merely describe the system; it becomes an active agent in its stabilization and optimization over time, with each successful intervention making future predictions even more reliable.

Despite the strong results, this study has limitations. The model was developed and validated using a dataset from a single healthcare facility, which may limit its direct generalizability to other hospitals with different patient populations, operational workflows, or resource constraints. Future research should focus on validating and recalibrating the model across a diverse range of hospital systems. Further enhancements could involve the incorporation of additional real-time data streams, such as laboratory and radiology test turnaround times, which are often significant contributors to delays. Finally, exploring more advanced deep learning architectures, such as Long Short-Term Memory (LSTM) networks, could offer further improvements by more explicitly modeling the time-series nature of patient arrivals and departmental congestion.

4.3 Data Visualization and Feature Insights

Prior to model development, exploratory data analysis through visualization provided critical insights that guided feature engineering and confirmed the relevance of the selected variables (Figure 1 and Figure 2).

- Wait Time Distribution: A histogram of the total ER wait times revealed a right-skewed distribution. The mean wait time was 81.92 minutes, substantially higher than the median of 60.00 minutes. This disparity indicates that while a typical patient waits about an hour, a significant number of patients experience much longer delays, confirming the presence of outliers and underscoring the complexity of the prediction challenge.

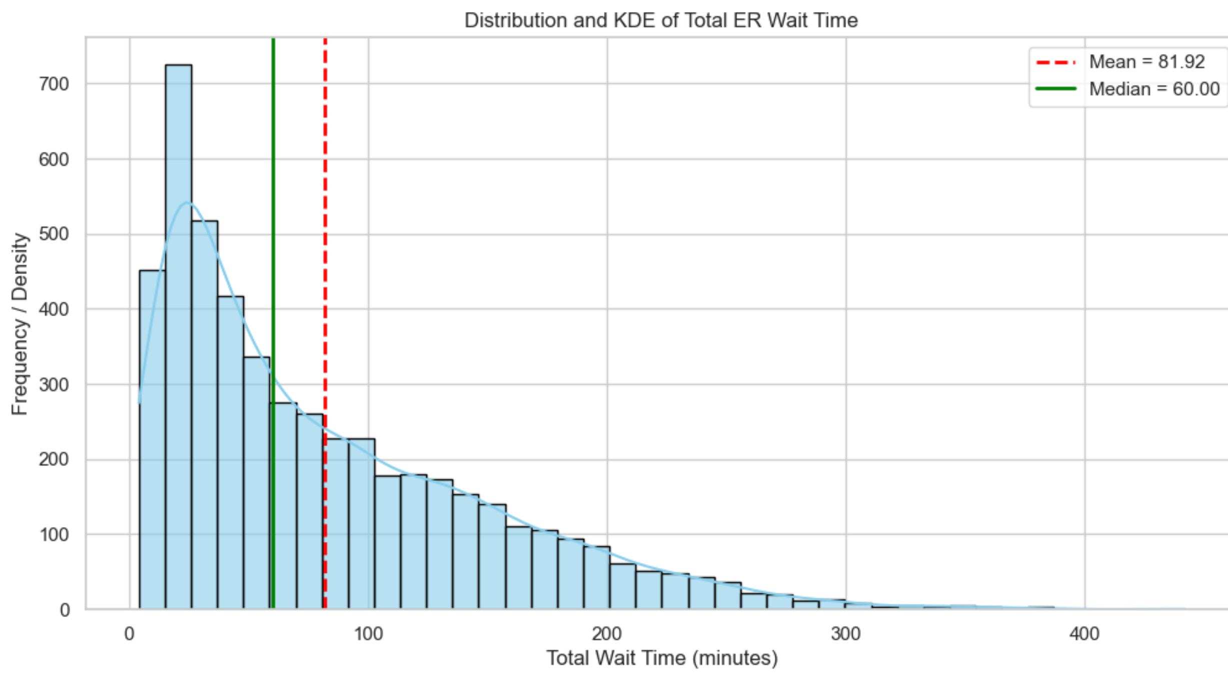

Figure 1. Frequency distribution of total patient wait times in the emergency room. The histogram reveals a right-skewed distribution, with a mean wait time of 81.92 minutes and a median of 60.00 minutes.

- Feature Correlation: A Pearson correlation matrix was used to assess the linear relationships between numeric features and the total wait time. The analysis revealed a strong positive correlation of 0.69 between the Nurse-to-Patient Ratio and Total Wait Time, quantitatively confirming that a higher patient load per nurse is strongly associated with longer waits. In contrast, features such as Specialist Availability and Facility Size showed negligible linear correlation, suggesting their impact is likely more complex and non-linear.

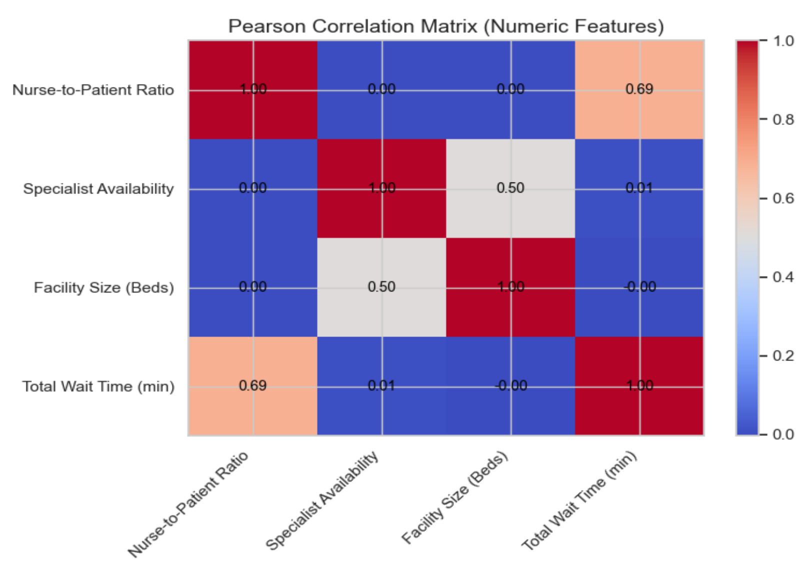

Figure 2. Pearson correlation matrix illustrating the linear relationships between numeric features. The matrix highlights a strong positive correlation (r = 0.69) between the Nurse-to-Patient Ratio and Total Wait Time, while other features show negligible linear correlation

- Temporal Patterns: Charting average wait times against the hour of the day revealed a distinct bimodal distribution, with peaks in the late morning and early evening, confirming the necessity of the time of day feature. Similarly, plotting wait times by day of the week showed elevated levels on Mondays and weekends, while seasonal plots indicated longer waits during winter months, validating the inclusion of these temporal factors.

These visualizations were instrumental not only in validating the feature selection process but also in highlighting the complex, non-linear interactions between variables. For example, the impact of a low staffing ratio was visibly more severe during peak arrival times. This underscored the limitations of simple linear models and reinforced the decision to employ more sophisticated, non-linear architectures like tree-based ensembles and the final stacked model, which are better suited to capture such intricate relationships.

Despite the strong results, this study has limitations. The model was developed and validated using a dataset from a single healthcare facility, which may limit its direct generalizability to other hospitals with different patient populations, operational workflows, or resource constraints. Future research should focus on validating and recalibrating the model across a diverse range of hospital systems. Further enhancements could involve the incorporation of additional real-time data streams, such as laboratory and radiology test turnaround times, which are often significant contributors to delays. Finally, exploring more advanced deep learning architectures, such as Long Short-Term Memory (LSTM) networks, could offer further improvements by more explicitly modeling the time-series nature of patient arrivals and departmental congestion.

5. Conclusion

This study addressed the critical operational challenge of protracted patient waiting times in hospital emergency rooms, a problem with severe consequences for patient satisfaction, clinical outcomes, and institutional efficiency. By leveraging a real-world dataset of patient encounters, this research successfully developed and validated a robust machine learning framework for accurate prediction of these wait times.

The central finding of this paper is the demonstrated superiority of a hyperparameter-tuned Stacked Ensemble Model. Through a rigorous comparative analysis against other standard machine learning algorithms and a robust 5-fold cross-validation protocol, the stacked architecture proved most effective, achieving a coefficient of determination (R^2) of 0.9473. This result indicates an exceptionally strong predictive capability, validating the model's capacity to capture the complex, non-linear dynamics that govern ER patient flow.

The practical value of this research lies in its delivery of a validated, data-driven tool for hospital administrators. The model empowers a shift from reactive problem-solving to proactive resource management, enabling more efficient staff allocation and workflow optimization. It also provides a mechanism for improving patient communication and managing expectations, thereby enhancing the overall patient experience. Ultimately, this study underscores the transformative potential of applying advanced analytics to solve systemic inefficiencies within healthcare. It presents a clear trajectory toward a more efficient, predictable, and patient-centric emergency care environment, where clinical resources are allocated more efficaciously and the quality of care is improved.

References

A Comprehensive Guide to Machine Learning In Healthcare [WWW Document], n.d. URL https://www.itransition.com/machine-learning/healthcare (accessed 9.11.25).

AI in Healthcare Operations Management: Optimizing Efficiency and Care | Calonji [WWW Document], n.d. URL https://www.calonji.com/blog/ai-in-healthcare-operations-management (accessed 9.11.25).

Alrasheedi, K.F., AL-Mohaithef, M., Edrees, H.H., Chandramohan, S., The Association Between Wait Times and Patient Satisfaction: Findings From Primary Health Centers in the Kingdom of Saudi Arabia. Health Serv Res Manag Epidemiol 6, 2333392819861246. 2019. https://doi.org/10.1177/2333392819861246

Ameur, N. Ben, Lahyani, I., Thabet, R., Megdiche, I., Steinbach, J. christophe, Lamine, EPredicting Patient's Waiting Times in Emergency Department: A Retrospective Study in the CHIC Hospital Since 2019. Communications in Computer and Information Science 1751 CCIS, 44–57. 2022. https://doi.org/10.1007/978-3-031-23119-3_4

Baturynska, I., Martinsen, K., Prediction of geometry deviations in additive manufactured parts: comparison of linear regression with machine learning algorithms. J Intell Manuf 32, 179–200. 2021.https://doi.org/10.1007/S10845-020-01567-0/TABLES/7

Dong, S.L., Bullard, M.,. Emergency Department Triage. Evidence-Based Emergency Medicine 58–65. 2023. https://doi.org/10.1002/9781444303674.ch7

Hospitals - Timely & effective care | Provider Data Catalog [WWW Document], n.d. URL https://data.cms.gov/provider-data/topics/hospitals/timely-effective-care (accessed 9.11.25).

Kumar, S., Prasad, P., AI-Powered Patient Flow Optimization in Emergency Rooms. 2025.

Mahajan, P., Uddin, S., Hajati, F., Moni, M.A., 2023. Ensemble learning for disease prediction: A review. Healthcare (Switzerland) 11, 1808. https://doi.org/10.3390/HEALTHCARE11121808

Mean Absolute Error In Machine Learning: What You Need To Know - Arize AI [WWW Document], n.d. URL https://arize.com/blog-course/mean-absolute-error-in-machine-learning-what-you-need-to-know/ (accessed 9.11.25).

Nyce, A., Gandhi, S., Freeze, B., Bosire, J., Ricca, T., Kupersmith, E., Mazzarelli, A., Rachoin, J.S, Association of Emergency Department Waiting Times With Patient Experience in Admitted and Discharged Patients. J Patient Exp 8, 23743735211011404). 2021. https://doi.org/10.1177/23743735211011404

Overcoming Common Challenges in the ED | Abbott Point of Care [WWW Document], n.d. URL https://www.globalpointofcare.abbott/gb/en/lp/apoc/overcoming-emergency-department-challenges.html (accessed 9.11.25).

RMSE Explained: A Guide to Regression Prediction Accuracy | DataCamp [WWW Document], n.d. URL https://www.datacamp.com/tutorial/rmse (accessed 9.11.25).

Staff-to-Patient Ratio: A Critical KPI for Your Healthcare Practice [WWW Document], n.d. URL https://www.fathomhq.com/kpi-glossary/staff-to-patient-ratio (accessed 9.11.25).

Sultan, S.Q., Javaid, N., Alrajeh, N., Aslam, M.,. Machine Learning-Based Stacking Ensemble Model for Prediction of Heart Disease with Explainable AI and K-Fold Cross-Validation: A Symmetric Approach. Symmetry Vol. 17, Page 185 17, 185. 2025, https://doi.org/10.3390/SYM17020185

The Risks Associated with Long ER Wait Times | Physicians Premier ER [WWW Document], n.d. URL https://mdpremier.com/the-risks-associated-with-long-er-wait-times/ (accessed 9.11.25). 2025

Wang, H., Sambamoorthi, N., Sandlin, D., Sambamoorthi, U., Interpretable machine learning models for prolonged Emergency Department wait time prediction. BMC Health Serv Res 25, 403. 2025.https://doi.org/10.1186/S12913-025-12535-W

Biographies

M. M. Aflatun Kawsar is a Lecturer in the Department of Industrial and Production Engineering at the European University of Bangladesh. His research interests lie in Operations Research and Optimization, Machine Learning, and Supply Chain and Logistics. In his current academic role, he has successfully implemented Outcome-Based Education (OBE) into the curriculum and has co-supervised undergraduate students in their thesis work. Mr. Kawsar's undergraduate research utilized a Total Interpretive Structural Model (TISM) to facilitate the adoption of circular economic principles within the construction industry. He earned his B.Sc. in Industrial and Production Engineering from the Military Institute of Science and Technology (MIST).

Turjo Ghoshal is a graduate student and researcher in the Department of Mechanical Engineering at the University of Delaware. His research interests include composite materials, nanomaterials, advanced material processing, and the development of wearable sensors. Currently, his research focuses on fabricating knit fabric-based pressure sensors for health monitoring using Electrophoretic Deposition (EPD). Prior to his current studies, Mr. Ghoshal served as a Lecturer in the Department of Industrial & Production Engineering at the European University of Bangladesh. His undergraduate thesis involved the experimental and numerical study of NREL S834 and S835 airfoils, utilizing a subsonic wind tunnel and Ansys Fluent to determine their aerodynamic characteristics under low wind speed conditions. Mr. Ghoshal earned his Bachelor of Science in Mechanical Engineering from Chittagong University of Engineering & Technology (CUET).

Md Shamim Hasan is a Data Analyst at LandQuire, where he specializes in business analysis, automated bot integration, and data management. His research is centered on applying machine learning and deep learning models to solve complex problems in healthcare and supply chain management. Mr. Hasan's undergraduate thesis involved developing a supervised machine learning model with Bayesian Belief Networks to optimize supplier selection and mitigate risks in the pharmaceutical sector. He has published research in Plos One on advancing breast cancer prediction and has further investigation on cardiovascular disease diagnosis and supplier selection for the jute industry. He earned his B.Sc. from the Military Institute of Science and Technology, Bangladesh.