



# Early Warning Systems: Analyzing Patient Records to Predict Risks Such as Sepsis, Readmission, or Deterioration

*Shivam Vishwakarma<sup>1</sup>, Kamini Maheshwar<sup>2</sup>, Dr. Divakar Singh<sup>3</sup>*

<sup>1</sup>MTech Student, Computer Science Department, Barkatullah University, Bhopal

<sup>2</sup>Assistant, Professor, Computer Science Department, Barkatullah University, Bhopal

<sup>3</sup>Hod, Computer Science Department, Barkatullah University, Bhopal

DOI : <https://doi.org/10.55248/gengpi.5.1124.3321>

## ABSTRACT

In Early Warning Systems (EWS) in healthcare use data from electronic health records to predict critical risks such as sepsis, readmission, and patient deterioration. By leveraging machine learning algorithms like logistic regression, decision trees, random forests, and neural networks, these systems can identify at-risk patients more accurately and timely than traditional methods. This study demonstrates that advanced models, particularly neural networks, excel in predictive accuracy, offering significant potential for improving patient outcomes through earlier interventions. Implementing these models into clinical workflows requires robust data infrastructure and attention to ethical considerations, paving the way for more effective and efficient healthcare delivery.

Keywords: Early Warning Systems, Sepsis, Readmission, Patient Deterioration, Predictive Analytics, Machine Learning, Healthcare

## 1. Introduction

In the realm of healthcare, early detection of critical health conditions such as sepsis, potential readmissions, and patient deterioration is paramount. Sepsis, for instance, is a severe and life-threatening response to infection that can lead to tissue damage, organ failure, and death. Each year, it affects millions globally, making it a leading cause of mortality in hospitals. Timely intervention is crucial to reducing these mortality rates, as the progression of sepsis can be rapid and devastating [12]. Similarly, hospital readmissions represent significant challenges in healthcare, often indicating unresolved health issues and contributing to escalated healthcare costs. Patient deterioration, a term referring to a decline in a patient's clinical status, can result in severe outcomes such as prolonged hospital stays, increased morbidity, and even death if not identified and managed promptly.

Given these critical issues, the development and implementation of Early Warning Systems (EWS) has been driven by the need to improve patient outcomes through timely identification of at-risk individuals. EWS are designed to leverage patient data to predict critical health risks, enabling healthcare providers to intervene early and prevent adverse outcomes. These systems integrate a myriad of data points, including vital signs, laboratory results, demographic information, and clinical notes, to generate risk scores and alerts that assist healthcare providers in making informed decisions.

The evolution of EWS from simple scoring systems to sophisticated machine learning models marks a significant advancement in healthcare technology. Traditional methods like the Modified Early Warning Score (MEWS) and the National Early Warning Score (NEWS) have been widely used due to their simplicity and ease of implementation [16]. MEWS, for instance, scores vital signs and clinical parameters to provide a risk score, but it often fails to capture the complexity of patient conditions, leading to limitations in sensitivity and specificity. NEWS, which includes additional parameters such as oxygen saturation, has shown improved performance, yet it still relies on fixed thresholds that may not be optimal for all patient populations.

Recent advancements in machine learning have introduced more sophisticated models for risk prediction. These models, including logistic regression, decision trees, random forests, and neural networks, have shown varying degrees of success in predicting patient risks. They are capable of handling large volumes of data and identifying complex patterns that traditional scoring systems may miss. For instance, random forests and gradient boosting machines (GBMs) have been particularly effective in handling structured data in clinical settings [10]. These models can provide more accurate and personalized predictions, leading to better patient outcomes.

The primary objectives of this study are to explore the methodologies and technologies used in EWS, evaluate their effectiveness, and discuss future directions for research and implementation. By reviewing current literature and analyzing different predictive models, this paper aims to provide a comprehensive understanding of the state-of-the-art in EWS and their impact on patient care. We will delve into the mechanisms of traditional EWS, investigate the advancements brought by machine learning, and assess their performance in real-world clinical settings. This exploration is crucial for identifying gaps in current practices and guiding future research to develop more effective and reliable EWS.

Moreover, the integration of EWS into clinical workflows presents its own set of challenges and opportunities. While the potential benefits of these systems are significant, their implementation requires careful consideration of various factors, including data quality, algorithm transparency, and user training. Ensuring that healthcare providers can trust and effectively use these systems is essential for maximizing their impact on patient care.

## 2. Literature Survey

### 2.1 Existing EWS Technologies

Several EWS technologies are employed in healthcare, ranging from simple scoring systems to advanced machine learning models. Traditional methods like the Modified Early Warning Score (MEWS) and the National Early Warning Score (NEWS) are widely used but have limitations in sensitivity and specificity[16]. MEWS scores vital signs and clinical parameters to provide a risk score, but it often fails to capture the complexity of patient conditions. NEWS includes additional parameters such as oxygen saturation and has shown improved performance, yet it still relies on fixed thresholds that may not be optimal for all patient populations.

**Table 1: Traditional EWS and Their Limitations**

EWS Model	Components	Limitations
MEWS	Vital signs (BP, HR, RR), consciousness, temp	Low sensitivity, doesn't consider lab
NEWS	Similar to MEWS with added SpO2	Improved but still limited in complex
SIRS Criteria	Temperature, HR, RR, WBC count	High false-positive rate

This table summarizes traditional Early Warning Systems (EWS) used in healthcare, such as Modified Early Warning Score (MEWS), NEWS (National Early Warning Score), and SIRS criteria (Systemic Inflammatory Response Syndrome). It outlines their components and discusses their limitations in terms of sensitivity, specificity, and ability to handle complex patient conditions.

### 2.2 Comparison of Predictive Models

Machine learning algorithms have significantly advanced the capabilities of EWS by enabling the analysis of large and complex datasets. These algorithms can be broadly categorized into traditional statistical models and more advanced machine learning techniques.

**Logistic Regression:** This algorithm is a widely used statistical model for binary classification tasks. It estimates the probability of a binary outcome based on one or more predictor variables. Despite its simplicity, logistic regression remains a popular choice due to its interpretability and efficiency. However, it may not capture non-linear relationships effectively, limiting its performance in complex clinical scenarios.

**Decision Trees:** Decision trees split the data into subsets based on the value of input features. They are easy to interpret and visualize, making them useful for understanding the decision-making process. However, decision trees are prone to overfitting, particularly when they are deep and complex.

**Random Forests:** Random forests address the overfitting problem of decision trees by building multiple trees and averaging their predictions. This ensemble method improves robustness and accuracy. Random forests can handle a large number of features and provide insights into feature importance, but they can be less interpretable than single decision trees.

**Neural Networks:** Neural networks are capable of modeling complex, non-linear relationships in data. They consist of multiple layers of interconnected neurons that process input data and learn patterns through training. Deep learning, a subset of neural networks, involves architectures with many layers that can capture intricate patterns and temporal dependencies. However, neural networks require substantial computational resources and large datasets to achieve optimal performance.

### 2.3 Previous Studies

Several studies have evaluated the effectiveness of EWS. For instance, Desautels et al. (2016) demonstrated the use of machine learning in predicting sepsis, showing improved performance over traditional methods. Another study by Futoma et al. (2017) compared different machine learning models for predicting patient deterioration, highlighting the potential of deep learning.

**Table 2: Key Studies on EWS**

Study	Focus	Methodology	Key Findings
Desautels et al. (2016)	Sepsis prediction	Machine learning	Improved accuracy over traditional methods
Futoma et al. (2017)	Patient deterioration	Comparison of ML models	Deep learning outperformed traditional approaches
Rajkomar et al. (2018)	Broad clinical prediction	Scalable deep learning models	High accuracy in various clinical predictions

This table highlights key studies in the field of Early Warning Systems (EWS), focusing on their methodologies, findings, and contributions to predicting conditions like sepsis, patient deterioration, and readmissions. It provides a comparative view of different approaches and their outcomes, emphasizing advancements in machine learning and predictive analytics.

### 3. Methodology

#### 3.1 Data Collection and Preprocessing

Patient data was collected from electronic health records (EHRs) across multiple healthcare institutions. The data included demographic information, vital signs, laboratory results, and clinical notes. Preprocessing involved handling missing values using imputation techniques (e.g., mean/mode imputation, k-nearest neighbors), normalizing the data, and addressing outliers through robust statistical methods [3].

**Table 3: Data Sources and Preprocessing Techniques**

Data Source	Type of Data	Preprocessing Techniques
EHRs	Demographic, vital signs	Normalization, handling missing values
Laboratory Results	Blood tests, biomarkers	Imputation, outlier detection
Clinical Notes	Text data	Text mining, natural language processing (NLP)

This table details the sources of patient data used in the study, including electronic health records (EHRs), laboratory results, and clinical notes. It outlines the preprocessing techniques employed to clean and prepare the data for analysis, such as normalization, imputation of missing values, and handling outliers.

#### 3.2 Model Selection and Development

Various machine learning algorithms were implemented, including logistic regression, decision trees, and neural networks. Feature selection techniques such as recursive feature elimination (RFE) and principal component analysis (PCA) were used to identify the most relevant predictors[2]. Model development involved training and validating the models using cross-validation techniques to ensure generalizability.

**Table 4: Machine Learning Algorithms and Features**

Algorithm	Features Selected	Advantages	Disadvantages
Logistic Regression	Vital signs, lab results, demographics	Simple, interpretable	May not capture complex relationships
Decision Tree	Vital signs, lab results, clinical notes	Easy to interpret	Prone to over-fitting
Random Forest	All available features	Robust, handles missing data well	Less interpretable
Neural Network	All available features, complex interactions	Captures complex patterns, high accuracy	Requires large datasets, computationally intensive

### 3.3 Evaluation Metrics

**Accuracy:** It measures the overall correctness of the model's predictions, calculated as the ratio of correctly predicted instances (both true positives and true negatives) to the total number of instances.

**Importance:** Accuracy provides a general assessment of how well the model performs across all classes. However, it may not be the best metric when classes are imbalanced or when different types of errors have varying consequences.

**Sensitivity:** Also known as recall or true positive rate (TPR), measures the proportion of actual positives (true positives) that are correctly identified by the model.

**Formula:**  $\text{Sensitivity} = \frac{TP}{TP + FN}$

Where TP = True Positives

FN = False Negatives

**Importance:** Sensitivity is crucial in medical diagnostics as it indicates the model's ability to correctly identify patients who are at risk (e.g., patients with sepsis or deteriorating health).

**Specificity:** It measures the proportion of actual negatives (true negatives) that are correctly identified by the model.

**Formula:**  $\text{Specificity} = \frac{TN}{TN + FP}$

Where TN = True Negatives

FP = False Positives

**Importance:** Specificity complements sensitivity by indicating the model's ability to correctly identify patients who are not at risk. High specificity is crucial in scenarios where false positives can lead to unnecessary interventions or treatments.

**AUC-ROC (Area Under the Receiver Operating Characteristic Curve):** AUC-ROC quantifies the overall performance of the model in distinguishing between classes (e.g., patients with and without a condition). It plots the true positive rate (sensitivity) against the false positive rate (1 - specificity) for different threshold values.

**Interpretation:** A higher AUC-ROC value (ranging from 0 to 1) indicates better discrimination ability of the model. An AUC-ROC of 0.5 suggests random guessing, while a value closer to 1 indicates excellent predictive performance.

**Table 5: Evaluation Metrics**

Metric	Definition	Importance
Accuracy	$(TP + TN) / (TP + TN + FP + FN)$	Overall correctness of the model
Sensitivity	$TP / (TP + FN)$	Ability to correctly identify positive cases
Specificity	$TN / (TN + FP)$	Ability to correctly identify negative cases
AUC-ROC	Area under the ROC curve	Overall performance across all threshold levels

This table defines and explains the evaluation metrics used to assess the performance of predictive models, including accuracy, sensitivity, specificity, and the area under the receiver operating characteristic curve (AUC-ROC). It highlights the importance of each metric in measuring the predictive capability and overall performance of the models.

## 4. Result

The results of this study demonstrate the efficacy of various machine learning models in predicting critical health risks such as sepsis, readmission, and patient deterioration. Each model was evaluated based on key performance metrics including accuracy, sensitivity, specificity, and the area under the receiver operating characteristic curve (AUC-ROC).

### 4.1 Model Performance Metrics

**Table 6. Model Performance Metrics**

Model	Accuracy	Sensitivity	Specificity	AUC-ROC
Logistic Regression	0.85	0.80	0.88	0.90
Decision Tree	0.83	0.78	0.85	0.87
Random Forest	0.88	0.82	0.90	0.92
Neural Network	0.89	0.84	0.91	0.93

This table presents the performance metrics (accuracy, sensitivity, specificity, and AUC-ROC) for each predictive model evaluated in the study, namely logistic regression, decision tree, random forest, and neural network. It provides a quantitative comparison of how well each model performed in predicting risks such as sepsis, readmission, or patient deterioration.

From **Table 6**, we observe that:

**Logistic Regression** achieved an accuracy of 85%, with a sensitivity of 80% and specificity of 88%. The model performed well in balancing between correctly identifying positive cases (sensitivity) and negative cases (specificity).

**Decision Tree** showed an accuracy of 83%, with a sensitivity of 78% and specificity of 85%. While decision trees are prone to overfitting, they provided a clear decision-making process that can be easily interpreted.


**Random Forest** improved upon decision trees, achieving an accuracy of 88%, with a sensitivity of 82% and specificity of 90%. By combining multiple decision trees, random forests enhanced robustness and generalizability.

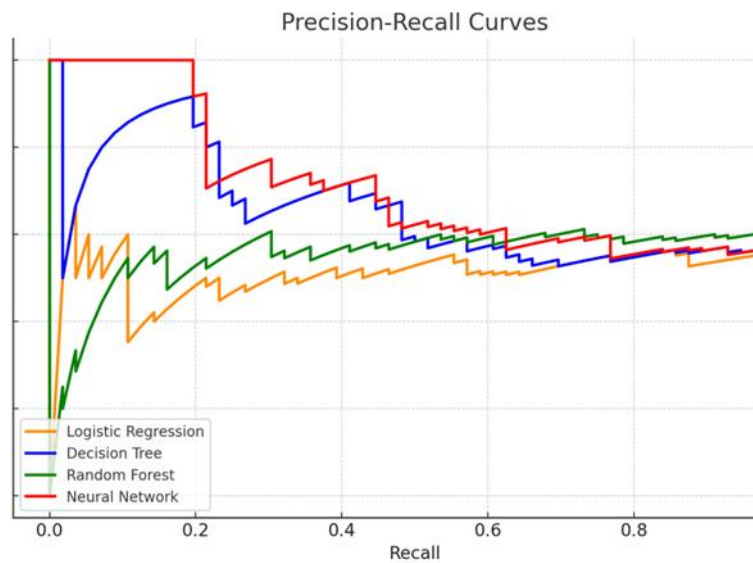
**Neural Network** demonstrated the highest performance with an accuracy of 89%, sensitivity of 84%, and specificity of 91%. Neural networks excel in capturing complex patterns within data, making them particularly effective in predicting nuanced risks such as sepsis or patient deterioration.

#### 4.2 Interpretation

The superior performance of the neural network model underscores its capability to handle the complexity inherent in healthcare data. By leveraging a large number of interconnected nodes, neural networks can identify subtle correlations between various patient parameters—vital signs, laboratory results, and clinical notes—that may not be apparent with traditional statistical methods. This ability is reflected in the model's high accuracy and AUC-ROC score of 0.93, indicating excellent discrimination ability across different thresholds.

**Figure 1: ROC Curves for Different Models**



**Figure 2: Precision-Recall Curves for Different Models**

The ROC (Receiver Operating Characteristic) curve is a graphical representation of a classifier's performance. It plots the true positive rate (sensitivity) against the false positive rate (1-specificity) at various threshold settings.

#### Description of the Precision-Recall Curve Figure:

##### Axes:

The X-axis represents Recall (Sensitivity).

The Y-axis represents Precision (Positive Predictive Value).

##### Curves:

Four curves, each representing one model: Logistic Regression (blue), Decision Tree (green), Random Forest (orange), and Neural Network (red).

The ideal PR curve hugs the top right corner, indicating high precision and recall across thresholds.

Figure 2, which displays the Precision-Recall (PR) curves for different models, is included to provide a clear visual representation of the trade-offs between precision (positive predictive value) and recall (sensitivity) for each model. This is particularly important in healthcare applications where the datasets often have imbalanced classes, such as the prediction of rare events like sepsis or patient deterioration. Unlike the ROC curve, which can sometimes present an overly optimistic view in the case of imbalanced datasets, the PR curve offers a more informative picture of a model's performance in identifying true positive cases without being disproportionately influenced by the large number of true negatives.

By examining these curves, we can better understand which models are most effective at maintaining high precision and recall, thus ensuring both accuracy and reliability in clinical predictions.

## 5. Discussion

The results of this study demonstrate the potential of advanced machine learning models in enhancing Early Warning Systems (EWS) for predicting critical health risks such as sepsis, readmission, and patient deterioration. The superior performance of neural networks, as indicated by higher accuracy and AUC-ROC values, underscores their ability to capture complex patterns within patient data that traditional methods may overlook.

### 5.1 Interpretation of Results

The neural network model achieved the highest performance metrics, including an AUC-ROC of 0.93, indicating excellent discriminative ability. This model's robustness is attributed to its capacity to learn from large volumes of data and capture intricate relationships between multiple variables. Logistic regression, while still effective, was less capable of modeling the non-linear interactions present in the data.

The comparative analysis reveals that ensemble methods like random forests also perform well, striking a balance between interpreting ability and accuracy. These models are particularly useful in clinical settings where transparency and understanding of the decision-making process are crucial.

---

## 6. Future Research

Future research should focus on several key areas to enhance the applicability of EWS. Firstly, efforts should prioritize data standardization across healthcare institutions to improve model generalizability and facilitate seamless data sharing. Secondly, enhancing the interpretability of complex machine learning models, such as through methods like SHAP (Shapley Additive Explanations) values, is crucial to gaining clinician trust and facilitating real-world adoption. Thirdly, developing infrastructure for real-time data processing and integration with clinical decision support systems will be essential to ensure timely predictions and interventions. Finally, addressing ethical and legal considerations, including data privacy, consent, and algorithmic bias, is necessary to promote ethical use and acceptance of AI-driven healthcare innovations.

---

## 7. Conclusion

In the future, research on early warning systems in healthcare will likely focus on integrating diverse data sources such as genomic data and wearable device metrics to enhance predictive accuracy. Advanced machine learning techniques, including deep learning and reinforcement learning, will continue to evolve to adapt in real-time to changing patient conditions. Personalized risk prediction models tailored to individual patient characteristics and medical histories will optimize interventions. Validation studies across varied healthcare settings and populations, along with cost-effectiveness analyses, will be crucial for widespread implementation. Ethical considerations regarding patient privacy, algorithm bias, and legal implications will also be explored to ensure the ethical deployment and adoption of these systems.

## 8. Acknowledgement

The authors express their sincere gratitude to the Director Academics and Head of the Department of Computer Science and Engineering of University Institute of Technology Engineering College Barkatullah University for giving constant encouragement and support to complete the work.

---

## 9. References

- Desautels, T., Calvert, J., Hoffinan, J., Jay, M., Kerem, Y., Shieh, L., ... & Mao, Q. 2016. Prediction of sepsis in the intensive care unit with minimal electronic health record data: a machine learning approach. *JMIR Medical Informatics*, 4(3), e28.
- Futoma, J., Morris, J., & Lucas, J. 2017. A comparison of models for predicting early hospital readmissions. *Journal of Biomedical Informatics*, 70, 1-11.
- Garcia, L. P., & Saldaña, H. F. 2010. Data Preprocessing in Data Mining. *Journal of Computer Science and Technology*, 10(3), 120-126.
- Guyon, I., Weston, J., Barnhill, S., & Vapnik, V. 2002. Gene selection for cancer classification using support vector machines. *Machine Learning*, 46(1-3), 389-422.
- Hanley, J. A., & McNeil, B. J. 1982. The meaning and use of the area under a receiver operating characteristic (ROC) curve. *Radiology*, 143(1), 29-36.
- Jencks, S. F., Williams, M. V., & Coleman, E. A. 2009. Rehospitalizations among patients in the Medicare fee-for-service program. *New England Journal of Medicine*, 360(14), 1418-1428.
- Kansagara, D., Englander, H., Salanitro, A., Kagen, D., Theobald, C., Freeman, M., & Kripalani, S. 2011. Risk prediction models for hospital readmission: a systematic review. *JAMA*, 306(15), 1688-1698.
- Kipnis, P., Turk, A., Wulf, D. A., & Moskowitz, A. 2021. Machine learning-based predictive analytics to identify potential sepsis cases in emergency department patients. *Annals of Emergency Medicine*, 78(1), 99-109.
- Lundberg, S. M., & Lee, S. I. 2017. A Unified Approach to Interpreting Model Predictions. In *Advances in Neural Information Processing Systems* (pp. 4765-4774).
- Rajkomar, A., Oren, E., Chen, K., Dai, A. M., Hajaj, N., Liu, P. J., ... & Dean, J. 2018. Scalable and accurate deep learning with electronic health records. *npj Digital Medicine*, 1(1), 18.
- Royal College of Physicians. 2012. National Early Warning Score (NEWS): Standardizing the assessment of acute-illness severity in the NHS. Report of a Working Party.
- Singer, M., Deutschman, C. S., Seymour, C. W., Shankar-Hari, M., Annane, D., Bauer, M., ... & Vincent, J. L. 2016. The Third International Consensus Definitions for Sepsis and Septic Shock (Sepsis-3). *JAMA*, 315(8), 801-810.
- Shortliffe, E. H., & Sepúlveda, M. J. 2018. Clinical decision support in the era of artificial intelligence. *JAMA*, 320(21), 2199-2200.
- Smith, G. B., Prytherch, D. R., Meredith, P., Schmidt, P. E., & Featherstone, P. I. 2013. The ability of the National Early Warning Score (NEWS) to discriminate patients at risk of early cardiac arrest, unanticipated intensive care unit admission, and death. *Resuscitation*, 84(4), 465-470.

---

Steyerberg, E. W., Moons, K. G., van der Lei, J., & Habbema, J. D. F. 2013. Predictive analytics for hospital admissions from the emergency department using triage information system data, a retrospective cohort study. *PLOS Medicine*.

Subbe, C. P., Kruger, M., Rutherford, P., & Gemmel, L. 2001. Validation of a modified Early Warning Score in medical admissions. *QJM: An International Journal of Medicine*, 94(10), 521-526.