

SHAP-Enabled Explainable ML for Early Sepsis Risk Prediction Using Tabular EHR Records

Sudipta Bose¹, Srijan Mondal¹, Suman Bhattacharya¹, Mahamuda Sultana¹, Sumit Kumar Banerjee^{1*}, Swarnali Daw², Ritwika Mukherjee², Paramita Sarkar³

¹Guru Nanak Institute of Technology, Kolkata, West Bengal, India.

Email ID:- bose.sudipta.025@gmail.com, ; Email ID:-srijanm.cse.jisu21@gmail.com ,

Email ID:-mahamuda.sultana@gnit.ac.in , Email ID:-suman.bhattacharya@gnit.ac.in,

Email ID:-sumit.banerjee@gnit.ac.in ,

²Narula Institute of Technology, Kolkata, West Bengal, India.

Email ID:-swarnalihiya@gmail.com , Email ID:-swarnali.daw@nit.ac.in , ,

Email ID:-ritwikamukherjee1999@gmail.com , , Email ID:-ritwika.mukherjee@nit.ac.in

³JIS University, Kolkata, West Bengal, India

Email ID:-paramita.sarkar@jisuniversity.ac.in

***Corresponding Author:**

Sumit Kumar Banerjee

Guru Nanak Institute of Technology, Kolkata, West Bengal, India.

Email ID:- sumit.banerjee@gnit.ac.in

ABSTRACT

The sepsis condition is a severe syndrome that necessitates timely treatment. Electronic Health Records (EHRs) contain valuable information that can be used to predict disease early. Structured EHR data exhibit informative missingness, feature heterogeneity, and complex behavior in downstream machine learning (ML) models, thus presenting challenges for application. The current research aims to propose an ML framework for the early prediction of sepsis from structured tabular EHR data. The study will develop an ML framework for data pre-processing, handling missing data, conducting EDA, selecting clinical features, and class balancing. Using the developed framework, we will employ supervised machine learning algorithms to help predict the probability of developing sepsis over a particular period. The framework will be applied to the EHR data, with SHAP used to interpret the results. In addition, the study will examine the associations between input variables and model outcome at multiple levels. The proposed model's performance is analyzed using different metrics, including AUROC, precision, recall, F1 score, sensitivity, and calibration. Additionally, this study will employ SHAP values to assess variable importance and provide patient-level explanations. The research hypothesizes that an interpretable ML framework can be used to create an early-warning system for timely detection of sepsis from structured EHR data.

Keywords :Sepsis prediction, EHR data, explainable AI, machine learning, SHAP, clinical decision support, early warning system, tabular clinical data

INTRODUCTION

Sepsis refers to a serious condition that is triggered by an inappropriate response to infection and is characterized by the presence of dysfunction at the level of organs, which may lead to death. The situation is extremely dangerous because the condition can progress quite rapidly, thus making timely diagnosis and appropriate management critically important. However, early sepsis symptoms can overlap with other acute conditions, making timely and reliable prediction challenging.


Electronic health records typically contain vast amounts of records, including crucial signs, diagnostic results, medication records, & demographics, that can be used to identify risks early. EHR, together with ML, was used to identify patterns in structured data to predict future sepsis occurrences. While ML algorithms can sometimes be remarkably accurate, only a small percentage can explain why they reach particular conclusions, which may hinder clinicians' efforts to provide treatment.

Explainable artificial intelligence (XAI) could enhance the interpretability of ML-based clinical systems. SHapley Additive exPlanations (SHAP) are model-agnostic, enabling calculation of the importance of each feature in the final prediction at both global and patient-specific levels, thereby identifying key clinical predictors and offering explanations for a particular patient's high- or low-risk classification. The study offers an explainable ML framework for early sepsis prediction from tabular EHR data and SHAP analysis. First of all, the article concentrates on data pre-processing, missing data, feature engineering, and class imbalance. Second, the study trains ML models, applies SHAP analysis, and evaluates them using AUROC, AUPRC, sensitivity, specificity, precision, and F1 score, thereby providing an interpretable ML model for early sepsis prediction from routine EHR data.

RELATED WORK

ML is currently being used to predict sepsis from electronic health records (EHRs), as traditional clinical scores often fail to capture the complex underlying associations between physiological and laboratory variables. Islam et al. conducted a scoping review of the literature on machine-learning-based prediction of adult sepsis. The authors identified 42 studies for inclusion in the review, highlighting substantial heterogeneity in data sources, definitions, features, and outcome assessment across the included articles, which prompted the development of an initiative to standardize and harmonize the reporting of such studies (Islam et al., 2023). Yang et al. also conducted a similar review, identifying 23 studies within their scope using data from over 4.3 million patients. The authors found that the most common algorithms used to predict sepsis were ML models, with RF and XGBoost performing the best (Yang et al., 2023). Various researchers have investigated the potential of using structured clinical variables to predict sepsis early. Lauritsen et al. used the EHR event logs to develop an ML model to predict sepsis. The authors' results demonstrated that the model's accuracy generally declined as the horizon length increased and was only moderately accurate several hours before the diagnosis (Lauritsen et al., 2019). Srmedha et al. designed an ML pipeline to predict sepsis up to 6 hours before diagnosis using a combination of demographic data, vital signs, and EHR data. The study demonstrated that an RF pipeline was the best-performing approach (Srmedha et al., 2022). Kijpaisalratana et al. evaluated the performance of LR, RF, gradient boosting, and neural networks, reporting that, compared with qSOFA, MEWS, and SIRS, ML models performed well in predicting sepsis using emergency department EHR data (Kijpaisalratana et al., 2022). Su et al. compared the performance of LR, RF, and XGBoost for predicting sepsis-related outcomes using the first 6 hours of ICU data and found that RF was the best-performing algorithm overall (Su et al., 2021). ML-powered early sepsis detection is moving outside the ICU walls. Valik et al. developed SepsisFinder, which used sparse, routine EHR data. They showed that it could predict sepsis onset up to 48 hours before ICU admission, outperforming NEWS2 and alerting clinicians earlier than conventional scores (Valik et al., 2023). Several researchers have suggested addressing data imbalance when developing clinical machine learning models. Zhou et al. proposed using multiple imputation methods with SMOTE. The authors concluded that the RF method could be used to predict sepsis with a reasonably good AUROC on an independent test set (Zhou et al., 2024). This study underlines the importance of data pre-processing and balancing strategies for clinical ML performance. Explainability has become a major concern in clinical ML because many successful algorithms are generally considered black boxes. Therefore, several researchers have sought to describe the nature of the associations between clinical variables and the outcomes predicted by ML models. Nesaragi and Patidar described an explainable ML framework based on LightGBM and physiological and time-lag features; they used SHAP values to identify relevant variables associated with sepsis (Nesaragi & Patidar, 2021). Hu et al. designed an understandable and transparent ML model for sepsis prognosis using MIMIC-IV and XGBoost; their results showed that SHAP values could be used to explain feature associations at the patient level (Hu et al., 2022). Aygun et al. evaluated XGBoost, LightGBM, and AdaBoost algorithms for emergency department sepsis prediction and showed that SHAP values could explain the associations between clinical features and sepsis (Aygun et al., 2024). Similarly, Zhang et al. showed that SHAP values could be used to interpret a CBC-based ML model for the early detection of sepsis (Zhang et al., 2024). Recent studies have begun embedding methods that explain predictions made by high-complexity models into the design of advanced machine learning methods. This allows for the investigation of the role of clinical factors and the provision of predictive performance that is comparable to that of more complex models. Tang et al. proposed CNN- and BERT-based models, referred to as CNN-Transformer and LSTM-Transformer, respectively, for early sepsis prediction. The authors highlighted the use of SHAP values for interpreting the results (Tang et al., 2024). Shi et al. developed a prediction model by employing both structured electronic health records (EHRs) and natural language processing-extracted features from clinicians' notes using XGBoost and GRU algorithms. These authors externally validated the results in a pediatric population and used SHAP values to interpret the model (Shi et al., 2025). Liu et al. engineered structured emergency information to develop an interpretable machine learning model for early sepsis prediction in pre-hospital settings (Liu et al., 2025). The authors have shown that an interpretable gradient-boosting model can perform the prediction task proficiently by including demographics, history of present illness, history, and vital signs. They validated their findings in an external cohort and used SHAP values to demonstrate that the model could identify associations between clinical variables and clinician-relevant outcomes (Liu et al., 2025). Ma et al. showed that an XGBoost-based model achieved the best performance for predicting sepsis in elderly patients. Still, they also noted the need to evaluate their findings in an independent sample (Ma et al., 2025). However, most of these studies have several limitations, including being retrospective, using different definitions of sepsis, or having limited external and temporal validation (Islam et al., 2023), (Yang et al., 2023). Some studies focused on predicting sepsis mortality or outcomes rather than on early sepsis onset (Hu et al., 2022), (Gao et al., 2024). Earlier studies showed that sepsis could be predicted using minimal sets of electronic health record (EHR) variables (Desautels et al., 2016), whereas combining engineered clinical and deep features with ensemble approaches yielded even better performance (He et al., 2020). Subsequent works focused on other settings such as physiological signals (Sadasivuni et al., 2021)

or federated learning (Alam & Rahmani, 2023). Feature engineering, in particular, proved critical to the task: Bomrah et al. found that dynamic vital signs, laboratory variables, and specific strategies for extracting features from them were relevant across 29 studies (Bomrah et al., 2024). In turn, Prithula et al. evaluated traditional and stacked machine learning (ML) methods, with enhanced data processing using SMOTE and COPULA, on EHR data (Prithula et al., 2024). Recently, there has been growing interest in clinical validation: thus, SepsisAI tackled the issues of false alarms and alarm fatigue (Gupta et al., 2024), whereas TREWS's prospective multicenter study demonstrated the clinical value of ML-based predictive early-warning systems when appropriate alerts were integrated into the workflow (Adams et al., 2022). NAVOY® Sepsis was also subjected to a prospective study to validate its clinical performance (Persson et al., 2024). Explainability is another task of growing interest: Zhou et al., for example, combined external validation with SHAP explanations for their septic shock prediction model (Zhou et al., 2025). Ogunwale et al. have recently proposed a multi-horizon framework that applies structured data pre-processing, evaluates multiple classical ML models, and calibrates and interprets overall model performance and patient-specific predictions using SHAP values (Ogunwale et al., 2026). This proposed work is unified by the clinical pipeline, in which the authors analyze leakage over time, perform clinical feature engineering and pre-processing, address class imbalance, comparatively explore table modeling, and interpret global and patient-specific models using SHAP.



PROPOSED METHODOLOGY

This approach consists of several subsequent steps: EHR preparation and featurization, data pre-processing, prediction using a machine learning model, and post hoc explanation using SHapley Additive exPlanations (SHAP). The prediction part takes into account only the information known before the beginning of the prediction window.

EHR Data Collection: The first step in developing the proposed framework is to identify the structured EHR records needed to estimate early sepsis. All of this required data can be obtained from patients' hospital or intensive care unit (ICU) records, including demographics, clinical parameters, lab tests, & other relevant clinical variables. In these cases, patients' vital signs can provide essential information on the state of their circulatory and respiratory systems. In contrast, laboratory tests can help detect abnormal physiological functions, including potential infections.

The framework uses only structured clinical variables that were available before the specified prediction time. Such an approach helps ensure that the model is based on information that can actually be used to predict sepsis onset. The collected records are then organized at the patient- or clinical-encounter level to facilitate cohort selection, outcome labeling, and feature engineering.

Cohort Selection and Labeling: The second step is to select a relevant patient cohort and apply the chosen labels. Inclusion/exclusion criteria are used to remove irrelevant, duplicate, and incomplete patient records and ensure consistency within the selected patient cohort. The patient cohort is split into patients with and without sepsis groups based on the presence/absence of sepsis, as defined by the chosen criteria and the outcome information available in the selected EHR database.

For the cases, the clinically recorded or dataset-defined occurrence of sepsis is selected as the reference event. The controls are selected from the non-cases in which the outcome of interest was not observed during the follow-up. After labeling, the target variable can be used to train the supervised ML model. It is particularly important to ensure that the labeling process does not include any clinical information unavailable at the time of prediction.

Prediction Window Definition: The third stage determines the period during which the model is expected to predict sepsis before its onset. In other words, a specific prediction time is set, and observations of EHR before that timeline point are used as input data for the machine learning models. The following period is considered the prediction window, during which sepsis occurrence is evaluated.

Therefore, the temporal arrangement of the task guarantees that an early prediction model will not be trained to recognize instances of sepsis already present in the data. It also minimizes the risk of data leakage, in which information from EHRs near the onset of sepsis could undeservedly improve the model's performance. Thus, it can be concluded that the prediction window is the period during which early-warning systems can detect sepsis and help clinicians with further examinations.

Feature Construction: The fourth stage involves transforming the identified EHR measurements into machine-learning features. Features can be broadly grouped into demographic variables, clinical parameters, lab test results, and clinically relevant derived variables. All of the features include the vital signs that need to be measured regularly. These include heart rate, breathing rate, blood pressure, temperature, and oxygen saturation. Those reports come from the laboratory and include white blood cell count, lactate, creatinine, platelets, and glucose. Where available, additional features can be engineered from these time-dependent variables that capture changes or trends in the values of monitored parameters.

Trends in physiological measurements often contain important predictive information that individual measurements may not capture. Similarly, other clinically relevant derived variables can be constructed if these can be meaningfully predicted from the set of available variables. Finally, note that all features should ideally be derived from data available before the prediction time t to respect the early prediction paradigm.

Data Pre-processing: The fifth step involves preparing the derived clinical features for the machine learning task. EHR data may have missing values that require imputation, inconsistent or incorrect measurements that need to be removed, and variables on different scales. The clinical features with missing data are then imputed according to an established procedure.

Additionally, any invalid measurements are flagged and excluded from the model training process. The categorical variables, if any, were converted into their numerical equivalent. The variables were then prepared as required by the algorithm to which the data would be submitted for training. Of particular note, encoding and scaling parameters are determined from the training set and then applied in the same way to the validation and test data. The technique is important because the information from the test sets should not be used to shape the model; rather, the model should be used to process the test data.

Class-Imbalance Handling: The sixth step aims to address class imbalance in sepsis data, where the number of non-sepsis cases may outnumber sepsis cases. This could lead to an ML model being biased towards one class over the other and, consequently, failing to predict the less frequent class (sepsis) at all. To mitigate this situation, a proper class-weighting or resampling procedure is applied only to the training data, thereby increasing the representation of the underrepresented class. Class-imbalance handling procedures must be applied only to the training set, leaving the test set untouched. This way, the model is tested on the test set using metrics such as sensitivity, precision, F1-score, AUROC, & AUPRC, rather than accuracy, given the class imbalance.

Machine Learning Models: The seventh step uses the selected supervised machine learning models to predict sepsis risk based on the pre-processed features. Various algorithms are used to compare the performance of different models in learning patterns in tabular data extracted from clinical records. For instance, the Logistic Regression algorithm was initially applied due to its simplicity and interpretability. The second method, Random Forest, was implemented to identify nonlinear dependencies and interactions between the input variables by using a tree-based approach and an ensemble learning framework. Finally, another gradient boosting method, XGBoost, was used for its high performance in approximating complex relationships in structured data compared to other algorithms.

All three methods were trained on the prepared training set and further tuned using either the training or validation subset. As a result, each model predicted the probability of sepsis occurrence for a given patient and was later used to calculate SHAP values required to explain the results. The best-performing model was eventually selected for further interpretation.

Model Evaluation: The proposed models were assessed for predictive performance and evaluated on the test set. Since the data were imbalanced, the performance cannot be measured using a single metric. In binary classification for early sepsis prediction, both true positives and false positives are important. For that reason, the area under the ROC curve and the precision-recall curve were calculated.

Additionally, these predictive performances are analyzed across 4 categories of performance metrics. These categories are sensitivity, specificity, precision, and F1 Score. Precision is used to check how well the model

minimizes false positives. Sensitivity and specificity also help measure the model's performance; they check TP and TN rates from the confusion matrix. The comparison was performed on the same test set used in this study. Accordingly, the best model was chosen for further interpretation using the SHAP framework.

SHAP-Based Explainability: The ninth step is dedicated to interpreting the predictions with the selected model using SHAP. This approach allows determining the importance of each input variable in the model predictions. In other words, the explanation of the predictions is based on the definition of each clinical feature's contribution to the outcome. Thus, it becomes possible to solve the unexplainable problem in ML and increase the transparency of the selected model's predictions regarding the risk of sepsis development. First, there are two main ways to interpret SHAP: global and local. The former refers to the importance of input variables for the overall model, whereas the latter focuses on explaining the prediction for a particular instance.

In the current research, two types of SHAP explanations can be used for analysis. Global explanations can be represented by a summary plot that highlights the most relevant features for the model. At the same time, local explanations can help clarify exactly how the prediction result for each patient was derived and which features influenced the outcome. Finally, SHAP values can be visualized in various ways, including feature-importance plots, summary plots, and local explanations. It is worth noting that these visualizations do not indicate a causal relationship between the input variables and the final prediction. Instead, they provide evidence of predictive associations and help analyze the model behavior from a clinical perspective.

Clinical Interpretation and Insights: The final step is to discuss the interpretation of the model's predictions and the obtained SHAP explanation from the clinician's perspective. First, it should be discussed which features possess the highest values in the global explanation and how these features are connected to the development of physiological decline and an increased risk of sepsis.

Secondly, the clinician should review the local SHAP explanation to determine which vital signs contribute to the high- or low-risk prediction for the patient. The results can then be used to provide a rationale for the model's behavior and reveal the hidden patterns the algorithm discovered. The results of the interpretation are meant to support clinicians in their practice rather than replace them. The model's SHAP values should be considered input to clinical decisions, not definitive evidence. The link between the features and the outcome is associative and not causal.

RESULT AND DISCUSSION

Expected Prediction Performance: The proposed framework is anticipated to facilitate the demonstration of machine learning algorithms' capability to predict patient presentations at high risk of developing sepsis before the specified prediction window. Logistic Regression, Random Forest, and XGBoost models will be compared using the same set of pre-processed EHR features.

The author expects Logistic Regression to have the best interpretability and, therefore, the best baseline performance. At the same time, the ensemble method is expected to utilize the best possible nonlinearities and interactions between features the fastest. The comparison of the models will be based on indicators such as AUROC, AUPRC, sensitivity, specificity, precision, and F1-score.

The author's main focus was on increasing sensitivity and AUPRC, since a high-risk population is the target and outcome imbalances need to be reduced. The model that demonstrated the best results for discriminating and identifying the outcome of interest will be used for the interpretation.

Expected Explainability Analysis: The final chosen model is expected to be evaluated using the SHAP approach to enhance interpretability & the findings. In general, global SHAP explanations are meant to provide the main insights into the results obtained. In particular, features associated with physiological decline and laboratory test deviations are expected to have the greatest impact on prediction outcomes. In addition, with a local explanation that highlights the features driving specific patients' records, the model can be evaluated for the proper inclusion of clinical logic. However, SHAP values should not be used to infer causal relationships.

Expected Clinical Interpretation: It is anticipated that the proposed framework would offer clinically relevant insights by evaluating the model's features that influence the algorithm's predictions. The features with the greatest impact on the model will be identified, and their implications for sepsis development, deterioration of physiological functions, and the most frequent symptoms of the disease will be determined. The patient-specific explanations will help clinicians understand the reasons for the patient's low or high likelihood of developing sepsis. However, it should be noted that the information provided by the framework will only illustrate the data patterns discovered by the algorithm, while no causality will be implied. As a result, the proposed method must be used to facilitate time-critical clinical decisions about the early prediction of sepsis. Still, they should not be used as a substitute for professional medical advice or as a stand-alone solution.

Expected Overall Outcome: The proposed study will aim to design a framework that enables the selected sepsis prediction model to flag patients at risk of developing sepsis before the prediction window. Moreover, applying the proposed methodology will enable the model to explain the rationale behind its predictions, thereby assisting clinicians in selecting the best interventions.

As such, it will be essential to test the framework on an appropriate EHR-based ML dataset, which will allow the model to demonstrate its ability to predict patient outcomes accurately and enable timely and adequate interventions.

Discussion: Several prior works show that it is possible to predict sepsis from events captured in EHRs using machine learning. Lauritsen et al. (Lauritsen et al., 2019) studied the use of deep learning models trained on EHR representations as sequences of events. Srmedha et al. (Srmedha et al., 2022) & Kijpaisalratana et al. (Kijpaisalratana et al., 2022) describe classical ML methods for early sepsis prediction. Nesariagi and Patidar (Nesaragi & Patidar, 2021) then apply the XAI method SHAP to train an interpretable early-warning model for sepsis. Finally, Liu et al. (Liu et al., 2025) use machine learning to extract features from structured emergency information relevant to triage, and visualize them using the SHAP and LIME explainer methods. Multiple manuscripts address the issue of imbalanced data (Zhou et al., 2024), perform external validation of the model, or use the SHAP explanation method (He et al., 2020).

The main differences between the works under consideration and the presented research are described in Table 1. While in the works under consideration, the main focus was on improving a certain aspect of the prediction model (a certain ML algorithm, prognosis, or interpretability), the current study provides a more comprehensive approach that includes prediction window definition, clinical feature engineering, pre-processing, class imbalance analysis, comparison of tabular ML algorithms, and SHAP-based analysis of both global and patient-specific importance.

Unlike the works under consideration, the present study offers an integrated approach that includes defining the prediction window, clinical feature engineering, pre-processing, analysis of class imbalance, comparison of different tabular ML algorithms, and SHAP-based analysis of both global and patient-specific importance.

Therefore, rather than implementing a specific ML algorithm or using particular SHAP explanation methods, the current study proposes an integrated approach. Moreover, the introduced framework has not been experimentally verified; thus, further research is needed to confirm its performance.

Comparison With Existing Approaches

Study	Early Sepsis Prediction	Tabular EHR	Imbalance Handling	SHAP/XAI	Patient-Level Explanation
Lauritsen <i>et al.</i> (Lauritsen et al., 2019)	YES	YES	NO	NO	NO
Srmedha <i>et al.</i> (Srmedha et al., 2022)	YES	YES	NO	NO	NO
Kijpaisalratana <i>et al.</i> (Kijpaisalratana et al., 2022)	YES	YES	NO	NO	NO
Zhou <i>et al.</i> (Zhou et al., 2024)	YES	YES	YES	NO	NO
Nesaragi and Patidar (Nesaragi & Patidar, 2021)	YES	YES	NO	YES	YES
Hu <i>et al.</i> (Hu et al., 2022)	NO	YES	NO	YES	YES
Aygun <i>et al.</i> (Aygun et al., 2024)	YES	YES	NO	YES	YES
Liu <i>et al.</i> (Liu et al., 2025)	YES	YES	NO	YES	YES
Proposed Framework	YES	YES	YES	YES	YES

CONCLUSION

The work presented in this paper describes a SHAP-enabled explainable ML approach for early sepsis risk prediction using tabular EHR records. It was created by performing EHR preparation, study cohort identification and labeling, prediction-window specification, clinically relevant feature engineering, pre-processing, class imbalance handling, machine-learning-based prediction, and SHAP-based interpretation. The authors argue that their solution enables early sepsis risk prediction, demonstrates model behavior, and identifies factors associated with this outcome. Since this research provides a framework for predicting sepsis, it is too early to tell how effective such a model would be. Further studies are needed, and additional machine learning algorithms should be explored to find the best approach for this work. It would also be informative to identify external validation and temporal validation results for the framework, conduct calibration studies, and assess subgroups' performance. Such an algorithm might really help with the early detection of sepsis and be incorporated into clinicians' practice. However, it is essential to test the proposed model on actual data before applying the results of the research. In addition, it might be interesting to see whether any of the other proposed ideas, such as integrating the algorithm into EHR systems or having clinicians use SHAP explanations for early sepsis prediction, could work.

DECLARATIONS

Funding: None

Conflict of interest/Competing interests: None

Ethics approval and consent to participate: Not applicable

Consent for publication: Not applicable

Data availability: None

Materials availability: None

Code availability: None

Author contribution: All the authors contributed equally..

REFERENCES:

1. Islam, K. R., Prithula, J., Kumar, J., Tan, T., Reaz, M., Sumon, M. S. I., & Chowdhury, M. E. (2023). Machine learning-based early prediction of sepsis using electronic health records: A systematic review. *Journal of Clinical Medicine*, 12.
2. Yang, Z., Cui, X., & Song, Z. (2023). Predicting sepsis onset in ICU using machine learning models: A systematic review and meta-analysis. *BMC Infectious Diseases*, 23.
3. Lauritsen, S., et al. (2019). Early detection of sepsis utilizing deep learning on electronic health record event sequences. *Artificial Intelligence in Medicine*, 104, 101820.
4. Srimedha, B. C., Naveen Raj, R., & Mayya, V. (2022). A comprehensive machine learning based pipeline for an accurate early prediction of sepsis in ICU. *IEEE Access*, 10, 105120–105132.
5. Kijpaisalratana, N., et al. (2022). Machine learning algorithms for early sepsis detection in the emergency department: A retrospective study. *International Journal of Medical Informatics*, 160, 104689.
6. Su, L., et al. (2021). Early prediction of mortality, severity, and length of stay in the intensive care unit of sepsis patients based on Sepsis 3.0 by machine learning models. *Frontiers in Medicine*, 8.
7. Valik, J., et al. (2023). Predicting sepsis onset using a machine learned causal probabilistic network algorithm based on electronic health records data. *Scientific Reports*, 13.
8. Zhou, L., et al. (2024). An early sepsis prediction model utilizing machine learning and unbalanced data processing in a clinical context. *Preventive Medicine Reports*, 45.
9. Nesaragi, N., & Patidar, S. (2021). An explainable machine learning model for early prediction of sepsis using ICU data. *Infectious Diseases and Sepsis*.
10. Hu, C., et al. (2022). Interpretable machine learning for early prediction of prognosis in sepsis: A discovery and validation study. *Infectious Diseases and Therapy*, 11, 1117–1132.
11. Aygun, U., et al. (2024). Assessment of sepsis risk at admission to the emergency department: Clinical interpretable prediction model. *Diagnostics*, 14.
12. Zhang, T., et al. (2024). Development and validation of a machine learning-based interpretable model for predicting sepsis by complete blood cell parameters. *Heliyon*, 10.
13. Tang, Y., Zhang, Y., & Li, J. (2024). A time series driven model for early sepsis prediction based on transformer module. *BMC Medical Research Methodology*, 24.
14. Shi, X., et al. (2025). Accurate prediction of sepsis from pediatric emergency department to PICU using a machine-learning model. *Frontiers in Pediatrics*, 13.
15. Liu, Z., et al. (2025). Interpretable machine learning for predicting sepsis risk in emergency triage patients. *Scientific Reports*, 15.

16. Ma, X., et al. (2025). Constructing an early warning model for elderly sepsis patients based on machine learning. *Scientific Reports*, 15.
17. Gao, J., et al. (2024). Prediction of sepsis mortality in ICU patients using machine learning methods. *BMC Medical Informatics and Decision Making*, 24.
18. Desautels, T., et al. (2016). Prediction of sepsis in the intensive care unit with minimal electronic health record data: A machine learning approach. *JMIR Medical Informatics*, 4.
19. He, Z., et al. (2020). Early sepsis prediction using ensemble learning with deep features and artificial features extracted from clinical electronic health records. *Critical Care Medicine*.
20. Sadasivuni, S., et al. (2021). Fusion of fully integrated analog machine learning classifier with electronic medical records for real-time prediction of sepsis onset. *Scientific Reports*, 12.
21. Alam, M. U., & Rahmani, R. (2023). FedSepsis: A federated multi-modal deep learning-based Internet of Medical Things application for early detection of sepsis from electronic health records using Raspberry Pi and Jetson Nano devices. *Sensors*, 23.
22. Bomrah, S., et al. (2024). A scoping review of machine learning for sepsis prediction—feature engineering strategies and model performance: A step towards explainability. *Critical Care*, 28.
23. Prithula, J., et al. (2024). A novel classical machine learning framework for early sepsis prediction using electronic health record data from ICU patients. *Computers in Biology and Medicine*, 184, 109284.
24. Gupta, A., et al. (2024). Improving sepsis prediction in intensive care with SepsisAI: A clinical decision support system with a focus on minimizing false alarms. *PLOS Digital Health*, 3, e0000569.
25. Adams, R., et al. (2022). Prospective, multi-site study of patient outcomes after implementation of the TREWS machine learning-based early warning system for sepsis. *Nature Medicine*, 28, 1455–1460.
26. Persson, I., et al. (2024). Early prediction of sepsis in intensive care patients using the machine learning algorithm NAVOY® Sepsis, a prospective randomized clinical validation study. *Journal of Critical Care*, 80, 154400.
27. Zhou, H., Li, F., & Liu, X. (2025). Early prediction of septic shock in ICU patients using machine learning: Development, external validation, and explainability with SHAP. *International Journal of Medical Informatics*, 206, 106169.
28. Ogunwale, Y., et al. (2026). Explainable machine learning for multi-horizon early sepsis prediction in intensive care units. *Cureus Journal of Computer Science*.