

Improving Machine Learning Based Sepsis Diagnosis Using Heart Rate Variability

Sai Balaji

The Woodlands College Park High School
The Woodlands, TX, United States

Christopher Sun

Stanford University
Stanford, CA, United States

Anaiy Somalwar

University of California, Berkeley
Berkeley, CA, United States

Abstract—The early and accurate diagnosis of sepsis is critical for enhancing patient outcomes. This study aims to use heart rate variability (HRV) features to develop an effective predictive model for Sepsis-3 detection. Critical HRV features are identified through feature engineering methods, including statistical bootstrapping and the Boruta algorithm, after which XGBoost and Random Forest classifiers are trained with differential hyperparameter settings. In addition, ensemble models are constructed to pool the prediction probabilities of high-recall and high-precision classifiers and improve model performance. Finally, a deep learning model is trained on the HRV features, achieving an F1 score of 0.805, a precision of 0.851, and a recall of 0.763. The best-performing machine learning model is compared to this neural network through an interpretability analysis, where Local Interpretable Model-agnostic Explanations are implemented to determine decision-making criterion based on numerical ranges and thresholds for specific features. This study not only highlights the efficacy of HRV in automated sepsis diagnosis but also increases the transparency of black box outputs, maximizing clinical applicability.

Index Terms—Sepsis-3, Heart Rate Variability, Feature Engineering, Machine Learning, Model Interpretability

I. INTRODUCTION

Sepsis-3 represents the latest consensus definition on sepsis and septic shock, posing a significant global health issue with a high mortality rate and economic burden [1] [2]. Sepsis is characterized by a dysregulated host response to infection, causing fever and an abnormal white blood cell count [3]. Accurate diagnosis of sepsis is critical for reducing mortality rates and substantial healthcare costs [4] [5]. Though advances like the quick Sequential Organ Failure Assessment (qSOFA) may enhance early sepsis detection, the frequent vital sign monitoring and reassessment required adds significant burdens to nursing workload, in return for already-questionable timeliness and accuracy [6]. Heart rate variability (HRV), which reflects autonomic nervous system activity and is extractable from heart-monitoring devices like ECG, emerges as a promising non-invasive tool for early sepsis detection due to its ability to detect dysregulation before clinical symptoms appear [7]–[10].

Bedoya et al. (2020) and van Wijk et al. (2023) employ machine learning and continuous ECG monitoring for sepsis prediction, harnessing both invasive and non-invasive patient data to advance early sepsis detection standards [11] [12]. In contrast, our research zeroes in on exploiting HRV metrics for *non-invasive*, continuous monitoring of sepsis, offering a streamlined, patient-centric approach.

Sendak et al. (2020) successfully implemented “Sepsis Watch,” a deep learning-based sepsis detection tool integrating both invasive and non-invasive clinical data, into everyday clinical practice, showcasing the practical application and benefits of advanced machine learning techniques in enhancing sepsis management [13]. Henry et al. (2015) developed TREWScore, an intricate predictive model for septic shock that utilizes a wide array of physiological and laboratory data and outperforms traditional methods in early detection [14]. Building upon these two works, our study demonstrates that the same types of state-of-the-art models can operate on HRV metrics exclusively, resulting in a simpler feature space and more explainable and effective method for early sepsis detection in a clinical environment.

Additionally, in comparison to Van Wijk et al.’s analysis of 171 emergency organ dysfunction cases, our study examines a more extensive dataset of 4,314 patient cases [12]. Our novel approach spearheads the analysis of 57 HRV-derived metrics, employing statistical analyses, feature selection, and model construction strategies to uncover relevant relationships between HRV and Sepsis-3. We also present case studies demonstrating enriched understanding of model decision-making that may be crucial for physician-mediated environments.

II. METHODS

Fig. 1 encapsulates this study’s methodology, portraying experimental design, data acquisition, statistical analysis, feature selection, model construction, ensembling mechanisms, and model explainability under the umbrella of Sepsis-3 diagnosis.

A. Data Availability Statement

The dataset used in this research is titled “Derivation and Validation of Heart Rate Variability-Based Scoring Models for Continuous Risk Stratification in Patients with Suspected Sepsis in the Emergency Department,” authored by Kuan-Fu Chen from Chang Gung University College of Medicine and dated December 9, 2020 [15].

B. Statistical Analysis

1) *Relative Histogram Plot and Preliminary Assessment:* We preliminarily assess feature relevance by visually comparing the distribution of HRV features between individuals with and without Sepsis3. We then conduct t-tests and nonparametric bootstrapping for each HRV feature to identify statistically

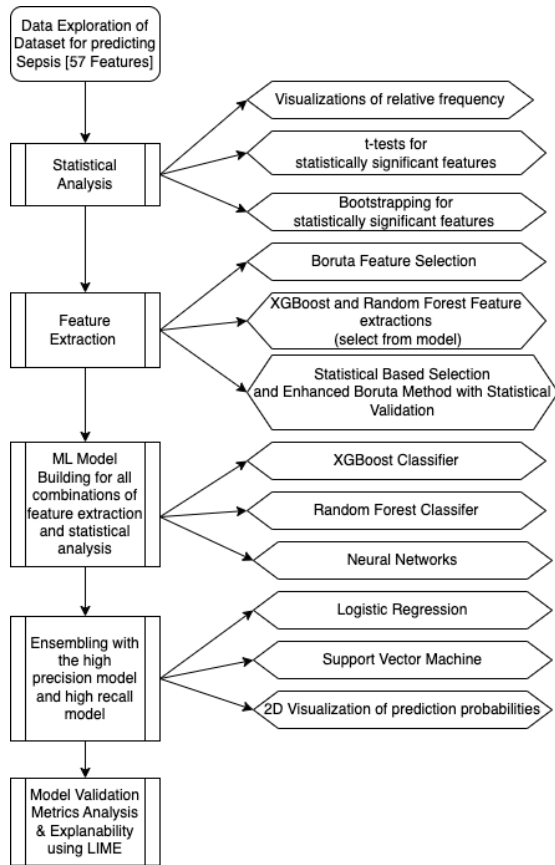


Fig. 1. Experimental Design

significant differences in mean values between patient subgroups. The nonparametric bootstrap enhances the reliability of the t-test by addressing the normality assumption. These statistical analyses establishes a foundation for identifying HRV features most relevant to Sepsis-3, setting the stage for more complex machine learning model development.

C. Feature Selection

We test three separate feature selection mechanisms and discuss a final method that incorporates elements from all.

1) *Model-based Feature Selection*: Model-based feature selection utilizes the intrinsic “importance scores” from ensemble tree-based methods, such as random forest and gradient boosting, to identify and retain HRV features that significantly contribute to predicting Sepsis-3. This approach captures influential predictors and excludes irrelevant noise from the dataset.

2) *Boruta Wrapper Method*: The Boruta algorithm, a wrapper method, systematically identifies relevant HRV features by generating shadow features, comparing their importance with actual features’ importance through a random forest classifier, and only selecting features that consistently show stronger signals than the highest-ranking shadow feature [16].

3) *Statistical Selection*: By employing statistical significance tests through bootstrapping, this method determines features with sheer numerical relevance for Sepsis-3 diagnosis.

We apply the Boruta method in conjunction with additional statistical validation, first using the Boruta algorithm to identify potentially relevant features and then using statistical tests to select only features surpassing a significance threshold. This hybrid method aims to yield an optimized feature set with high predictive ability, thereby improving the accuracy and dependability of the Sepsis-3 predictive model.

D. Model Construction

The model construction phase encompasses data preprocessing to standardize feature scales followed by the training of Extreme Gradient Boosting, Random Forest, and Neural Network models. For machine learning models, we carefully optimize decision thresholds by selecting the β parameter (of the F-beta score) that maximizes F1 score. Moreover, we handle the challenge of imbalanced classes within the dataset by implementing class weights as loss function coefficients. The following are design specifications for our models:

- **Random Forest Classifier**: Employs 1000 trees with a maximum depth of 10.
- **XGBoost Classifier**: Employs maximum of 1,000 estimators and learning rate of 0.01.
- **Neural Network**: Contains four dense layers populated with 64 hidden units each, employing batch normalization and dropout regularization to prevent overfitting.

E. Ensembling Techniques to Couple Model Decision-Making

After the development of individual models, we implement an ensemble method to leverage the strengths of individual predictive models for diagnosing Sepsis-3, aiming to enhance overall predictive performance by combining models with high precision and recall. Prediction probabilities from models excelling in one of precision and recall are stored in a data frame, forming the basis for subsequent ensemble strategies.

1) *ML Ensembling*: This ensemble approach integrates the output probabilities of a high precision model (HPM) and a high recall model (HRM) as input features, then trains a machine learning model on these probabilities with the objective of enhancing overall F1 score. We apply Logistic Regression and Support Vector Classifiers to this task.

2) *Manual Ensembling*: The manual ensembling method introduces a custom, rule-based strategy to combine predictions from models with high precision and recall, delivering a final Sepsis-3 diagnosis based on conditions of agreement or disagreement between the models. By assigning compound prediction statuses through a custom function, this technique provides a nuanced analysis of model consensus, leveraging the comparative advantages of high precision and high recall models.

F. Model Explainability

In an effort to bridge the gap between complex machine learning models, simple diagnostic outputs, and practical clinical applicability, this study integrates Local Interpretable Model-agnostic Explanations (LIME) to illuminate the decision-making processes of the highest-performing

model in predicting Sepsis-3 [17]. LIME provides granular insights into individual predictions, highlighting the relevance of specific HRV features in the model’s conclusions, thereby demystifying the model’s operational logic for clinicians and healthcare practitioners.

III. RESULTS

A. Statistical Analysis

Fig. 2 shows the distributions of two selected features for healthy and infected patients. We can see that the difference in mean values is visually noticeable; such observations guided our statistical feature selection framework. We further quantify the degree of phenotype difference using the t -statistic and p -value for each HRV feature. Statistical analysis reveals that mean heart rate, though elementary, is a strong predictor for Sepsis-3. Alongside mean heart rate are features like Shannon entropy and forbidden words, derived from a symbolic dynamics representation of time-series heart data. Fig. 3 displays the log of the absolute value of the difference in means for the fifteen most statistically significant features in the dataset.




Fig. 2. Feature Distributions by Sepsis-3 Status

B. Feature Selection and Importance

Utilizing the XGBoost and Random Forest model-based feature selection methods result in the identification of 52 and 25 key features, respectively, with a significant overlap, underscoring crucial HRV metrics detailed in Table III. The Boruta algorithm identifies 42 features, and a subsequent statistical-based bootstrapping analysis pinpoints 41 of these 42 features as statistically significant. These 41 features are collectively noted as Boruta-bootstrap features as in Table III. Our multi-faceted feature selection process not only underscores the complexity of Sepsis-3 pathophysiology but also highlights the diverse HRV metrics essential for accurate prediction models.

C. Model Performance Metrics

As a result of our exhaustive methodology applying different hyperparameters, feature sets, and machine learning architectures to the sepsis diagnosis task, we create three tables

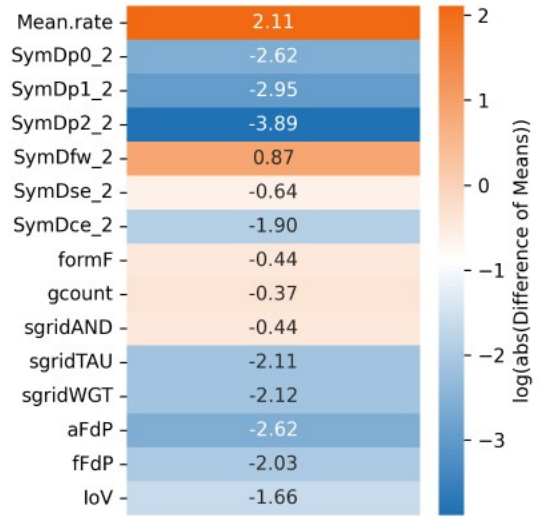


Fig. 3. Top 15 Statistically Significant Sepsis Features Based on Bootstrapping Difference Of Means

summarizing the performance metrics for all architectures tested. Tables IV, V and VI capture the nuances of the relationship between said settings and model performance. Namely, XGBoost is the better overall performer compared to Random Forest, and the Boruta-bootstrap feature selection mechanism tends to result in the best model when optimizing for precision and recall – individually and in tandem. On the other hand, using random forest’s intrinsic feature importance scores seem to lead to weak performance, which may not necessarily mean that these importance scores are irrelevant, but may very well be a testament to the added value of fine-tuning. Overall, we are unable to configure a model yielding both a precision and recall higher than 0.8; the best ML model is an XGBoost with bootstrap-determined features, achieving a F1 score of 0.745.

Notably, the Neural Network outperforms all machine learning models, emerging as a promising HRV-based classifier of Sepsis-3. Utilizing Boruta-bootstrap features, our model achieves a F1-score of 0.805, with a precision of 0.851 and recall of 0.763. We display learning curves in Fig. 4, which illustrate convergence after roughly 500 training epochs and mitigation of overfitting. Table I displays a confusion matrix for the deep learning model’s correct and incorrect predictions on the testing data, and Table II displays performance metrics.

TABLE I
CONFUSION MATRIX FROM NEURAL NETWORK MODEL

	Predicted: No Sepsis	Predicted: Sepsis	
Actual: No Sepsis	TN = 926	FP = 18	944
Actual: Sepsis	FN = 32	TP = 103	135
	958	121	1079

D. Ensembling Techniques

In theory, ensembling classifiers together for coupled model decision-making should improve performance [18]. Our results

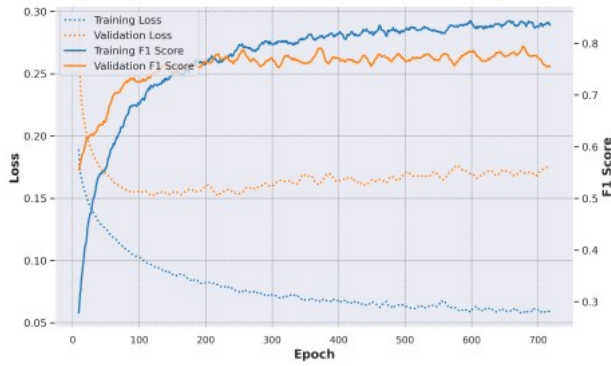


Fig. 4. Neural Network Learning Curves: Loss and F1 Score

weakly support this claim, as we find that the best ensemble does not perform as well as the best fine-tuned individual ML model, but outperforms an off-the-shelf XGBoost baseline. The logistic regression ensemble, merging a high-precision XGBoost model with a high-recall Random Forest model, achieves an F1-score of 0.724, with precision and recall metrics also presented. Interestingly, the support vector machine ensemble demonstrates better performance by 0.4% but has a 11.2% higher recall and 10.2% lower precision. Overall, both ensembling approaches yield favorable outcomes summarized in Table II. We can think of the input data frame of probabilities as an information-dense, specially-selected set of features; this allows for greater coupling accuracy even if the ensemble is a simple model like logistic regression.

We further visualize the results of ensembling by plotting two sets of probabilities against each other: Fig. 5 contains prediction probabilities of the high precision model (HPM) on the x-axis and prediction probabilities of the high recall model (HRM) on the y-axis. Ideally, this plot yields two “regions” separable by a straight line (or any other decision boundary that can be learned by an ML model). In our case, there seems to be only a faint curve of separation between purple and yellow points, representing an inability to perfectly separate healthy and affected patients.

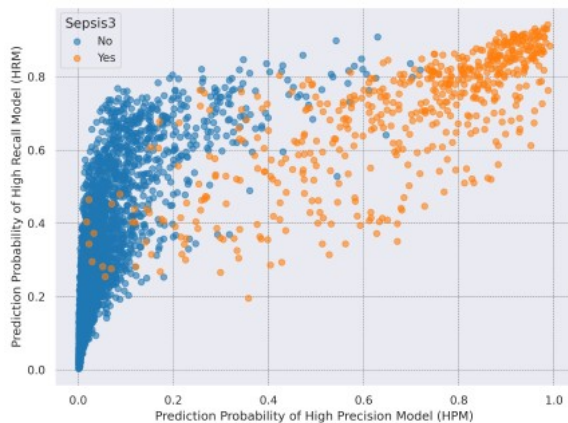


Fig. 5. 2-D Feature Space: HPM and HRM Prediction Probabilities

The manual ensemble method also reveals a notable recall of 0.825 alongside a model accuracy of 72.07%. Regarding clinical implementation, we encounter a 26.76% classification disagreement rate between the high precision and high recall models. In practice, these patients should be referred for physician diagnosis.

TABLE II
PREDICTIVE PERFORMANCE ACROSS MACHINE LEARNING MODELS

Model	F1 Score	Precision	Recall
Baseline XGBoost	0.682	0.689	0.676
Ensemble: Logistic Regression	0.724	0.791	0.667
Ensemble: SVM	0.728	0.689	0.778
Best ML Model (XGBoost)	0.745	0.760	0.731
Neural Network	0.805	0.851	0.763

E. Explainability Insights

Local Interpretable Model-agnostic Explanations (LIME) for the XGBoost model and Neural Network provide insightful perspectives on how HRV features are differentially weighed in a model prediction. From this, we may be able to untangle the complexity of HRV feature and even infer physiological implications.

1) Key Findings from XGBoost (Fig. 6):

- **Frequency Domain Measures:** The feature fFdp surfaces with significant importance across various ranges: when $fFdp > 0.24$, the XGBoost model interprets this as an elevated Sepsis-3 risk by almost 8% on average. When $-0.31 < fFdp \leq -0.24$, the severity of increased risk drops to around 3%. When $-0.67 < fFdp \leq -0.31$, the model reverses its decision-making, adding around 3% to the healthy diagnosis. Finally, when $fFdp \leq -0.67$, the model is confident in a decreased Sepsis-3 risk by almost 8%. This case study highlights the nuanced impact of HRV frequency domain measures on the model’s Sepsis-3 diagnosis. As one can see, the XGBoost model “discretizes” the fFdp feature, transforming a continuous feature scale into relevant percentages toward or against a Sepsis-3 diagnosis.
- **Amplitude and Non-linear Characteristics:** Features such as aFdp and KLPE appear in multiple ranges, indicating the critical role of amplitude variations and non-linear dynamics in HRV analysis, potentially reflecting the autonomic nervous system’s regulation.
- **Variability Indicators:** IoV (index of variability) and mean heart rate are identified as vital for understanding HRV patterns, as they may signify physiological adaptability and resilience.

2) Neural Network Model Insights (Fig. 7):

- **Highlighted Features:** Unlike the XGBoost model, the neural network does not prioritize fFdp, and rather seems to choose a whole new set of features altogether to give prominence to in the decision-making process. For example, the features vImax, dlmax, shannEn, and Poincar.SD2 have a pronounced association with HRV outcomes.

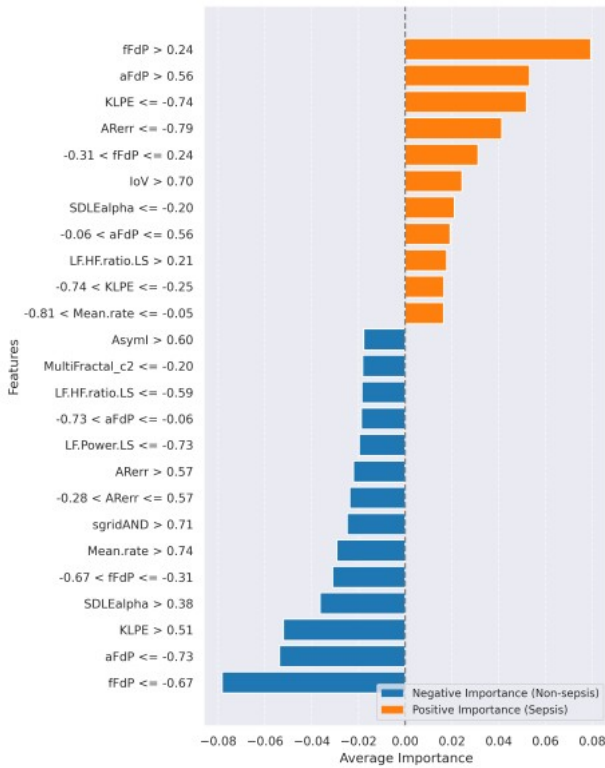


Fig. 6. LIME Average Feature Importance on Test Dataset using XGBoost Model

3) *Comparison*: Analysis for both models, while showcasing some overlap in feature significance (e.g. frequency domain measures, non-linear dynamics), also revealed distinct differences in the importance range of specific features. This discrepancy may arise from inherent differences in how each model approaches the task, with the neural network modeling non-linear relationships for example.

IV. DISCUSSION

Advancements in HRV-Based Early Sepsis Detection

Our comprehensive framework not only highlights the potential for artificial intelligence to significantly advance Sepsis-3 diagnostics but also emphasizes the promise of non-invasive monitoring methodologies through heart rate variability. Our initial trial of machine learning models with different hyperparameters and feature selection mechanisms reveals a maximum F1 score of 0.745. Introducing deep learning methods significantly increases this metric, as our neural network achieves a 0.805 F1 score. This suggests nonlinearity between HRV features and Sepsis-3.

We try to bring clinical utility to this work in two ways. We first offer interpretable explanations that allow physicians and users to understand a model’s decision-making criteria. We also prescribe a manual ensembling mechanism where two models, each specializing in one of precision and recall, are locked in a system of “checks and balances” to provide a final Sepsis-3 diagnosis. This manual ensemble is a feasible clinical




Fig. 7. LIME Average Feature Importance on Test Dataset using Neural Network

tool, as instances of model agreement can be trusted (0.825 recall), while instances of model disagreement can simply call for human intervention. This represents a pivotal progression toward integrating machine learning in clinical settings for predictive diagnostics; overall, our study corroborates the claims of van Wijk et al. (2023) regarding the indispensable role of HRV analysis in healthcare outcomes [12].

A. XGBoost LIME Explanations

Frequency Domain Measures: Van Wijk et al. found that the average of the NN-interval, ultra-low frequency (ULF), very low frequency (VLF), low frequency (LF), and total power differed significantly between groups with varying degrees of organ dysfunction in sepsis patients [12]. In our case, LIME explanations for the XGBoost model identify frequency domain-related features such as LF.Power.LS and LF.HF.ratio.LS as having considerable importance. This alignment suggests that changes in the frequency domain measures of HRV are indeed indicative of physiological alterations, potentially related to organ dysfunction or sepsis.

B. Neural Network LIME Explanations

Continuous Monitoring Capability: The utility of continuous ECG monitoring highlighted by van Wijk et al. aligns with the importance of dynamic features such as Poincar.SD2 and shannEn identified by the Neural Network’s LIME explanations [12]. Such features, derived from continuous data, could provide a granular view of HRV changes over time, potentially

enhancing the predictive capability for clinical deterioration in sepsis patients.

Previous research, in conjunction with our LIME explanations, underscore the potential of HRV analysis in monitoring clinical deterioration in sepsis patients. We advocate for the adoption of HRV features in automated sepsis detection; we argue that the equipment and capabilities required are already in place, as HRV features can be extracted from devices such as ECG and passed to machine learning pipelines for patient diagnosis. As a non-invasive, real-time indicator of physiological change, HRV is a prime feature for risk stratification and monitoring tools in clinical settings.

V. CONCLUSION

This research effectively leverages heart rate variability to enhance Sepsis-3 detection. We pinpoint a critical set of HRV features by implementing a hybrid Boruta-bootstraping feature engineering mechanism. We assemble a suite of advanced machine learning models with differential β values, then ensemble high-precision and high-recall models to visualize a 2-D feature space of probabilities. Though ensembling does not boost model performance significantly, we find that deep learning methods result in the highest F1 score, with our Dense Neural Network achieving an F1 Score of 0.805, a precision of 0.851, and a recall of 0.763. Furthermore, we include Local Interpretable Model-agnostic Explanations to elevate this model from a mere predictive tool to a transparent system with clinical applicability. Consequently, this research marks a significant advancement in sepsis risk stratification, offering healthcare providers a tool that identifies at-risk patients and provides insights into physiological indicators, merging early detection with actionable intelligence.

VI. ADDITIONAL INFORMATION

The authors have made the code used in this research available on GitHub (click).

REFERENCES

- [1] C. W. Seymour, J. N. Kennedy, S. Wang, *et al.*, “Derivation, validation, and potential treatment implications of novel clinical phenotypes for sepsis,” *JAMA*, vol. 321, no. 20, pp. 2003–2017, 2020.
- [2] T. Iwashyna, “Sepsis: Understanding a common cause of life-threatening infection,” *Michigan Health Lab.Iwashyna*, 2016.
- [3] A. Rhodes, L. E. Evans, W. Alhazzani, *et al.*, “Surviving sepsis campaign: International guidelines for management of sepsis and septic shock 2016,” *Intensive care medicine*, vol. 46, pp. 854–887, 2020.
- [4] K. E. Rudd, S. C. Johnson, K. M. Agesa, *et al.*, “Global, regional, and national sepsis incidence and mortality, 1990–2017: Analysis for the global burden of disease study,” *Lancet*, vol. 395, no. 10219, pp. 200–211, 2020.
- [5] R. Kleinpell, L. Aitken, and C. Schorr, “Impending sepsis: Public awareness and education,” *AACN advanced critical care*, vol. 24, no. 1, pp. 60–69, 2013.
- [6] K. R. Islam, J. Prithula, J. Kumar, *et al.*, “Machine learning-based early prediction of sepsis using electronic health records: A systematic review,” *Journal of clinical medicine*, vol. 12, no. 17, p. 5658, 2023.
- [7] S. P. Shashikumar, M. D. Stanley, I. Sadiq, *et al.*, “Early sepsis detection in critical care patients using multiscale blood pressure and heart rate dynamics,” *Journal of electrocardiology*, vol. 60, pp. 113–120, 2020.
- [8] T. Ahmad, P. Fiadjoe, G. Wilson, and D. Friedman, “Impending onset of sepsis encoded in heart rate variability,” *Pediatric Critical Care Medicine*, vol. 19, no. 9, e437–e444, 2018.
- [9] F. M. de Castilho, A. L. P. Ribeiro, V. Nobre, G. Barros, and M. R. de Sousa, “Heart rate variability as predictor of mortality in sepsis: A systematic review,” *PloS one*, vol. 13, no. 9, e0203487, 2018.
- [10] C. León, G. Carrault, P. Pladys, and A. Beuchée, “Early detection of late onset sepsis in premature infants using visibility graph analysis of heart rate variability,” *IEEE journal of biomedical and health informatics*, vol. 25, no. 4, pp. 1006–1017, 2021. DOI: 10.1109/JBHI.2020.3021662.
- [11] A. D. Bedoya, J. Futoma, M. E. Clement, *et al.*, “Machine learning for early detection of sepsis: An internal and temporal validation study,” *JAMIA Open*, vol. 3, no. 2, pp. 252–260, 2020.
- [12] R. J. van Wijk, V. M. Quinten, M. C. van Rossum, H. R. Bouma, and J. C. Ter Maaten, “Predicting deterioration of patients with early sepsis at the emergency department using continuous heart rate variability analysis: A model-based approach,” *Scandinavian Journal of Trauma, Resuscitation and Emergency Medicine*, vol. 31, no. 1, p. 15, 2023.
- [13] M. P. Sendak, W. Ratliff, D. Sarro, *et al.*, “Real-world integration of a sepsis deep learning technology into routine clinical care: Implementation study,” *JMIR medical informatics*, vol. 8, no. 7, e15182, 2020.
- [14] K. E. Henry, D. N. Hager, P. J. Pronovost, and S. Saria, “A targeted real-time early warning score (trewscore) for septic shock,” *Science Translational Medicine*, vol. 7, no. 299, 299ra122, 2015.
- [15] K.-F. Chen, “Hrv sepsis dataset,” *Mendeley Data*, vol. V1, 2020. DOI: 10.17632/5tgk53r4v5.1.
- [16] M. B. Kursu, A. Jankowski, and W. R. Rudnicki, “Boruta—a system for feature selection,” *Fundamenta Informaticae*, vol. 101, no. 4, pp. 271–285, 2010.
- [17] M. T. Ribeiro, S. Singh, and C. Guestrin, “‘‘ why should i trust you?’’ explaining the predictions of any classifier,” in *Proceedings of the 22nd ACM SIGKDD international conference on knowledge discovery and data mining*, 2016, pp. 1135–1144.
- [18] M. K. Hasan, M. A. Alam, D. Das, E. Hossain, and M. Hasan, “Diabetes prediction using ensembling of different machine learning classifiers,” *IEEE Access*, vol. 8, pp. 76 516–76 531, 2020.

TABLE III
SEPSIS FEATURE SELECTION AND IMPORTANCE

Feature Name	Description	XGB	RF	B	Bo	B_Bo
Mean.rate	heart rate in BPM	Y	Y	Y	Y	Y
Coefficient.of.variation	ratio of the standard deviation to the mean	Y	N	Y	Y	Y
Poincar..SD1	Poincaré plot standard deviation perpendicular the line of identity	Y	Y	Y	Y	Y
Poincar..SD2	Poincaré plot standard deviation along the line of identity	Y	Y	Y	Y	Y
LF.HF.ratio.LS	Ratio of LF-to-HF power	Y	Y	Y	Y	Y
LF.Power.LS	power of the low-frequency band	Y	Y	Y	Y	Y
HF.Power.LS	power of the high-frequency band	Y	Y	Y	Y	Y
DFA.Alpha.1	Detrended fluctuation analysis, which describes short-term fluctuations	Y	N	Y	Y	Y
DFA.Alpha.2	Detrended fluctuation analysis, which describes long-term fluctuations	N	N	Y	Y	Y
Largest.Lyapunov.exponent	measures a non-linear system's sensitive dependence on starting conditions	N	N	N	Y	N
Correlation.dimension	Estimate of required model variables	Y	N	Y	Y	Y
Power.Law.Slope.LS	Power Law (based on frequency) slope x2	Y	Y	Y	Y	Y
Power.Law.Y.Intercept.LS	Power Law (based on frequency) y-intercept x2	Y	Y	Y	Y	Y
DFA.AUC	Detrended fluctuation analysis: area under the curve	Y	Y	Y	Y	Y
Multiscale.Entropy	Multiscale entropy - measures the regularity and complexity of a time series	Y	N	Y	N	Y
VLF.Power.LS	Absolute power of the very-low-frequency band	Y	Y	Y	Y	Y
Complexity	Hjorth parameter complexity	Y	Y	Y	Y	Y
eScaleE	Embedding scaling exponent	N	N	N	Y	N
pR	Recurrence quantification analysis: percentage of recurrences	Y	N	N	Y	N
pD	Recurrence quantification analysis: percentage of determinism	Y	N	N	N	N
dlmax	Recurrence quantification analysis: maximum diagonal line	Y	N	N	Y	N
sedl	Recurrence quantification analysis: Shannon entropy of the diagonals	Y	N	N	Y	N
pDpR	Recurrence quantification analysis: determinism/recurrences	Y	N	N	Y	N
pL	Recurrence quantification analysis: percentage of laminarity	N	N	N	N	N
vLmax	Recurrence quantification analysis: maximum vertical line	Y	N	Y	Y	Y
sevl	Recurrence quantification analysis: Shannon entropy of the vertical lines	N	N	N	N	N
shannEn	Shannon entropy, measures uncertainty in a random variable	Y	Y	Y	Y	Y
PSeo	Plotkin and Swamy energy operator: average energy	Y	Y	Y	Y	Y
Teo	Teager energy operator: average energy	Y	Y	Y	Y	Y
SymDp0_2	Symbolic dynamics: percentage of 0 variations sequences, non-uniform case	Y	N	N	Y	N
SymDp1_2	Symbolic dynamics: percentage of 1 variations sequences, non-uniform case	Y	N	N	Y	N
SymDp2_2	Symbolic dynamics: percentage of 2 variations sequences, non-uniform case	Y	N	Y	Y	Y
SymDfw_2	Symbolic dynamics: forbidden words, non-uniform case	Y	N	N	Y	N
SymDse_2	Symbolic dynamics: Shannon entropy, non-uniform case	Y	N	N	Y	N
SymDce_2	Symbolic dynamics: modified conditional entropy, non-uniform case	Y	N	N	Y	N
formF	Form factor	Y	N	Y	Y	Y
gcount	Grid transformation feature: grid count	Y	Y	Y	Y	Y
sgridAND	Grid transformation feature: AND similarity index	Y	Y	Y	Y	Y
sgridTAU	Grid transformation feature: time delay similarity index	Y	N	Y	Y	Y
sgridWGT	Grid transformation feature: weighted similarity index	Y	N	Y	Y	Y
aFdP	Allan factor distance from a Poisson distribution	Y	Y	Y	Y	Y
fFdP	Fano factor distance from a Poisson distribution	Y	Y	Y	Y	Y
IoV	Index of variability distance from a Poisson distribution	Y	Y	Y	Y	Y
KLPE	Kullback-Leibler permutation entropy	Y	Y	Y	Y	Y
AsymI	Multiscale time irreversibility asymmetry index	Y	Y	Y	Y	Y
CSI	Poincaré plot cardiac sympathetic index	Y	N	Y	Y	Y
CVI	Poincaré plot cardiac vagal index	Y	Y	Y	Y	Y
ARerr	Predictive feature: error from an autoregressive model	Y	Y	Y	Y	Y
histSI	Similarity index of the distributions	Y	N	N	Y	N
MultiFractal_c1	Multifractal spectrum cumulant of the first order	Y	N	Y	Y	Y
MultiFractal_c2	Multifractal spectrum cumulant of the second order	Y	N	Y	Y	Y
SDLEalpha	Scale-dependent Lyapunov exponent slope	Y	N	Y	Y	Y
SDLEmean	Scale-dependent Lyapunov exponent mean value	Y	N	Y	Y	Y
QSE	Quadratic sample entropy	Y	Y	Y	Y	Y
Hurst.exponent	Rate at which autocorrelations decrease as the lag between pairs of values increases	Y	Y	Y	Y	Y
mean	Mean value	Y	N	Y	Y	Y
median	Median value	Y	N	Y	Y	Y

XGB refers to XGBoost, RF refers to Random Forest, B refers to Boruta, Bo refers to Bootstrap, and B_Bo refers to Boruta-bootstrap, each representing different feature selection methodologies applied in this study.

TABLE IV
MODEL PERFORMANCE METRICS, BETA = 1.0

Model	Features selected	F1 Score	Precision	Recall
XGBoost	bootstrap	0.745	0.760	0.731
	boruta_bootstrap	0.742	0.726	0.759
	boruta	0.732	0.773	0.694
	xgboost	0.731	0.809	0.667
	randomforest	0.726	0.695	0.759
Random Forest	boruta_bootstrap	0.633	0.636	0.630
	boruta	0.631	0.571	0.704
	bootstrap	0.628	0.567	0.704
	xgboost	0.628	0.567	0.704
	randomforest	0.619	0.600	0.639

TABLE V
MODEL PERFORMANCE METRICS, BETA = 0.5

Model	Features selected	F1 Score	Precision	Recall
XGBoost	bootstrap	0.597	1.000	0.426
	boruta_bootstrap	0.629	0.980	0.463
	boruta	0.675	0.934	0.528
	xgboost	0.651	0.931	0.500
	randomforest	0.674	0.836	0.565
Random Forest	xgboost	0.551	0.780	0.426
	bootstrap	0.556	0.770	0.435
	boruta	0.583	0.761	0.472
	boruta_bootstrap	0.580	0.750	0.472
	randomforest	0.573	0.729	0.472

TABLE VI
MODEL PERFORMANCE METRICS, BETA = 2.0

Model	Features selected	F1 Score	Precision	Recall
Random Forest	boruta_bootstrap	0.505	0.350	0.907
	boruta	0.511	0.357	0.898
	bootstrap	0.501	0.348	0.899
	randomforest	0.500	0.348	0.889
XGBoost	boruta_bootstrap	0.689	0.574	0.861
	bootstrap	0.679	0.560	0.861
	boruta	0.709	0.622	0.824
	randomforest	0.693	0.597	0.824
Random Forest	xgboost	0.725	0.659	0.806
Random Forest	xgboost	0.587	0.474	0.769