ISSN [Online]: 2583-2654

Optimizing Sepsis Care with AI: Integrating Deep Learning into Clinical Decision Support System

[1]Geetha C. Megharaj, [2]Thara Devi M, [3]R Abhishek, [4]S K Jawad Ahmed, [5]Sharan M, [6]Vishwas K P
Department of CSE, Dr. T. Thimmaiah Institute of Technology, KGF, India

Abstract: Sepsis is a life-threatening disease and a sustained Worldwide health issue that requires a quick and accurate diagnosis to improve patient outcomes. This project uses AI-involved clinical decision making systems (CDSS) that use advanced deep learning techniques, particularly multilayer perceptron (MLPs), to predict the onset of sepsis using clinical data such as key features, experimental results, and demographic characteristics. The system is trained on data records published in the Healthcare System and is designed to be well generalized in real-world hospital scenarios to ensure both clinical relevance and reliability. By predicting the analysis, the platform allows early detection, risk stratification, and severity classification, thereby enabling knowledge that can be implemented by the parents of healthcare professionals. The system is delivered via a user-friendly web interface to promote data-controlled, real-time decision discovery, improve clinical efficiency and contribute to reducing sepsis-related mortality. This MLP-based CDSS aims to optimize diagnostic accuracy, optimize treatment workflows and improve all patient management in the intensive care unit.

Keywords: Artificial Intelligence, Clinical Decision Support System, Deep Learning, Sepsis Prediction, MLP, Time-Series Data

I. INTRODUCTION

Sepsis is a severe and potentially fatal medical condition triggered by the body's extreme response to infection, often leading to tissue damage, organ failure, and death if not treated promptly. Globally, sepsis accounts for a significant percentage of hospital mortality, especially in intensive care units (ICUs). Early diagnosis and timely intervention are essential to reduce its high morbidity and mortality rates. However, traditional diagnostic approaches, which rely heavily on clinical judgment and rule-based scoring systems like SOFA (Sequential Organ Failure Assessment) or MEWS (Modified Early Warning Score), often result in delayed recognition of the condition due to their limited sensitivity to dynamic and non-linear clinical patterns.

In recent years, artificial intelligence (AI) and machine learning (ML) have emerged as powerful tools in healthcare, offering potential solutions to challenges in early disease prediction and diagnosis. Among these, models are known for capturing complex, non-linear relationships within large-scale clinical datasets. Specifically, the Multi-Layer Perceptron (MLP), a type of feedforward neural network, has shown strong performance in classification tasks involving high-dimensional medical data. MLP models are capable of learning from structured input such as patient vitals, laboratory test results, and prediction tasks.

This project proposes the development of a Clinical Decision Support System (CDSS) that incorporates an MLP-based predictive model to detect early signs of sepsis in hospitalized patients. By processing diverse patient data—including real-time vitals, lab reports, and administrative records—the system aims to generate timely alerts

ISSN [Online]: 2583-2654

and risk scores to assist clinicians in proactive decision-making. The primary aim is to connect data-driven insights and their application in real-world healthcare environments. By integrating MLP-based deep learning techniques into a real-time CDSS framework, the system enhances diagnostic accuracy, supports early intervention, and ultimately contributes to better patient outcomes and improved critical care efficiency.

II. LITRATURE SURVEY

The prediction and early diagnosis of sepsis have been the focus of numerous studies in recent years, as researchers strive to develop intelligent systems capable of supporting clinical decisions and improving patient outcomes. Traditional scoring systems like SOFA and MEWS, while widely used, often fail to capture the complex, non-linear, and temporal nature of sepsis progression. To address this gap, several AI-driven models have been mentioned in the literature.

In ***[1], Tang et al***. introduced a reinforcement learning-based approach using a Dueling Double Deep Q-Network (D3QN) to optimize sepsis treatment policies. Their work focused on revisiting reward and loss functions to enhance learning stability and clinical relevance. By leveraging a dynamic environment simulation, the model was able to suggest treatment strategies that improved long-term outcomes for sepsis patients. This approach, while effective in treatment planning, did not directly focus on early detection but highlighted the potential of deep reinforcement learning in critical care optimization.

In contrast, ***[2] Souza and Dragoni*** proposed a novel Parallel Evolving Fuzzy Neural Network (PEFNN) architecture for real-time sepsis identification. Their model mixed fuzzy logic with evolving neural networks to manage uncertainty and adapt to changes in patient data streams. The system demonstrated strong performance in identifying sepsis across diverse patient profiles, offering interpretability and adaptability. However, while PEFNN excelled in real-time classification, it relied heavily on handcrafted features and lacked the deep sequential learning capabilities necessary to capture temporal trends in EHR data.

In ***[3], Babu, Annavarapu***, and Appala presented a study on sepsis detection using neural networks, focusing on constructing and training a Multi-Layer Perceptron (MLP) model. Their main objective was to use digital healthcare data from the PhysioNet Challenge and MIMIC-III datasets to predict whether a person has sepsis. A secondary objective involved comparing the MLP's accuracy against other classifier algorithms such as Gradient Boosting, Ada-Boost, and Linear Discriminant Analysis. The proposed system included pre-processing, feature importance using XG Boost, and classification. The MLP Classifier outperformed other methods, achieving 94.7% accuracy on the Physionet dataset and 82.6% on the MIMIC-III dataset. While the paper highlighted the MLP's superior performance compared to other classifiers they tested, it also noted that previous RNN models had lower accuracy and LSTM training was challenging. Future work includes designing a user-specific application and deploying the model and integrate with hospital website.

In ***[4], He et al.*** proposed an approach for early sepsis prediction using ensemble learning combined with features extracted from LSTM recurrent neural networks. As an entry to the PhysioNet/Computing in Cardiology Challenge 2019, their main objective was to tackle class imbalance and data missing, manually extract features on medical knowledge, and use pre-trained LSTM models as feature extractors for "deep" time-series features. The proposed system integrated these manual and deep features to train prediction models within an ensemble learning framework, utilizing XG Boost and GBDT regressors. The approach achieved a normalized utility score of 0.313 on the full hidden test set. The authors

ISSN [Online]: 2583-2654

highlighted that different feature subsets and models provided diverse perspectives beneficial for ensemble learning, improving the averaged utility score. They also noted the advantage of LSTM in extracting time-series dependencies. Future work includes integrating more medical features and experimenting with additional machine learning methods to enhance prediction accuracy.

In ***[5], Rout et al***. explored deep learning in the prediction and diagnosis of sepsis, integrating Long Short-Term Memory (LSTM) and Support Vector Machine (SVM) models. Their primary objective was to identify sepsis early using physiological data from ICU patients. The study collected vital signs, laboratory results, and demographic information as inputs. The proposed model, based on LSTM and SVM, aimed to predict sepsis among ICU patients and achieved an AUC-ROC score of 0.696. While the model showed promising results and outperformed a random forecast on some test sets, its performance declined on one test set. The authors noted that LSTMs handle non-uniformly sampled sequences more effectively than CNNs. They also mentioned that SVMs are known for high classification accuracy. Future work includes investigating the influence of challenge utility score hyperparameters.

In ***[6], Gilbertson et al***. presented a method that combines dimensionality reduction and neural network-based analysis to identify sepsis in its initial stages. Their algorithm, submitted to the 2019 PhysioNet Computing in Cardiology Challenge, reduced 40 features to 10 principal components and incorporated a modified qSOFA score as an eleventh feature. It was trained using the selected features. The approach showed promise in detecting sepsis from clinical data, achieving a validation score of 0.038 and an official competition score of 0.022. While their method demonstrated effectiveness, they knew that there future work could improve accuracy by adding features, adjusting neural network parameters, and utilizing a sliding window to incorporate past information.

In ***[7], Demirer and Demirer*** proposed an artificial intelligence-based early warning and therapeutic decision support system for sepsis using Partially Observable Markov Decision Processes (POMDP). Their approach analyzes unevenly sampled observations, including demographics, vital signs, and laboratory values. They argue that traditional deep learning approaches provide a black-box representation and are insufficient for identifying patient risk and making predictions at every time interval. The authors claim their Lomb-Scargle periodogram approach within the POMDP framework is superior to deep learning algorithms. This work theoretically formed a new mathematical theory for sepsis detection to achieve Pareto Optimality conditions. The paper focuses on the theoretical framework, with applications to be realized in later articles.

In ***[8], Saqib, Sha, and Wang*** investigated the early prediction of sepsis using traditional machine learning techniques and deep learning LSTM networks on the MIMIC-III dataset. They aimed to predict sepsis using only the first 24 and 36 hours of lab results and vital signs. Their results indicated that the best performing traditional classifier, Random Forest, achieved an AUC-ROC score of 0.696, and their LSTM networks did not outperform Random Forest. Models based on 6-hour binned averages of time series significantly outperformed naive average models. The study highlighted that ignoring the time-dependent nature of sepsis leads to significantly lower prediction metrics. Future work suggested integrating additional data, testing more models, and exploring LSTM hyperparameters.

In ***[9], Babu, Annavarapu, and Appala*** examined deep learning methods for detection of sepsis, specifically comparing Neural Network, LSTM, and CNN on the PhysioNet Challenge 2019 dataset. Their work included studying the effect of different

ISSN [Online]: 2583-2654

missing value imputation methods. Among the tested classifiers (MLP Classifier, Gradient Boosting, and Ada-Boost), the MLP Classifier achieved the highest accuracy on both the Physionet challenge dataset (94.7%) and the MIMIC-III dataset (82.6%). The paper highlights the usefulness of deep learning for sepsis detection, especially using vital signs. Future work was identified in designing and deploying a user-specific application.

In ***[10], Anand et al.*** examined deep learning methods for the detection of sepsis, comparing the classification report, accuracy, and runtime of Neural Network (NN), LSTM, and CNN models on the PhysioNet Challenge 2019 dataset. The proposed approach incorporated techniques such as SMOTE to address data imbalance and applied feature selection to enhance model performance. While high accuracy was observed when training with the entire dataset, the study focused on training with vital attributes having fewer missing values, where CNN achieved a binary accuracy of 87.99%, LSTM achieved 62.50% accuracy, and NN achieved 62.81% accuracy. Deep learning techniques prove useful for sepsis detection, especially by leveraging readily available vital signs.

They noted that LSTM is prone to overfitting with smaller datasets. Future work suggested integrating the model with handheld devices for early diagnosis.

III. PROPOSED SYSTEM

The proposed methodology focuses on developing a Clinical Decision Support System (CDSS) capable of predicting the onset of sepsis using a Multi-Layer Perceptron (MLP) a deep learning model. The system follows a comprehensive pipeline consisting of data collection, preprocessing, feature engineering, model development, training, and evaluation. Designed to work with real-time Electronic Health Record (EHR) data, the CDSS supports timely and accurate detection of sepsis.

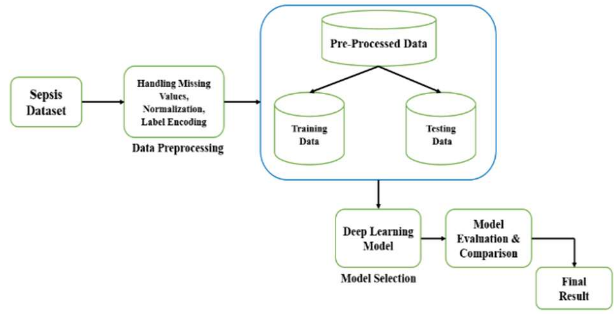

Fig 1: Methodology Diagram

A. Data Collection

A publicly available EHR dataset, comprising physiological and clinical variables collected from ICU patients over time, forms the basis of the analysis. It includes vital signs (e.g., heart rate, temperature, respiratory rate), laboratory test results and demographic information. These features play an important role in detecting early signs of sepsis. Each patient's health record is represented as a multivariate time-series:

$$X = \{x_1, x_2, \dots, x_T\}, \quad x_t \in R^n$$

Where, T is the number of time steps, and n is the number of clinical features. The dataset has both sepsis-positive and sepsis-negative cases, labeled accordingly. The structured and labeled data enables supervised learning using the MLP model for binary classification.

B. Data Preprocessing

Preprocessing clinical time-series data plays an important role in making sepsis prediction models are accurate and reliable. Since real-world healthcare data often contains inconsistencies—such as missing values, irregular time intervals, and mixed data types—careful preparation is essential before feeding it into a deep learning model. The preprocessing phase involves several key steps: filling in missing or incomplete data to ensure continuity, normalizing values to maintain consistency across different scales, aligning time-based records so events can be compared

ISSN [Online]: 2583-2654

accurately, and converting categorical information (like gender or diagnosis type) into numerical formats that the model can understand. These steps help the model focus on meaningful patterns, ultimately supporting faster and more accurate identification of sepsis in patients.

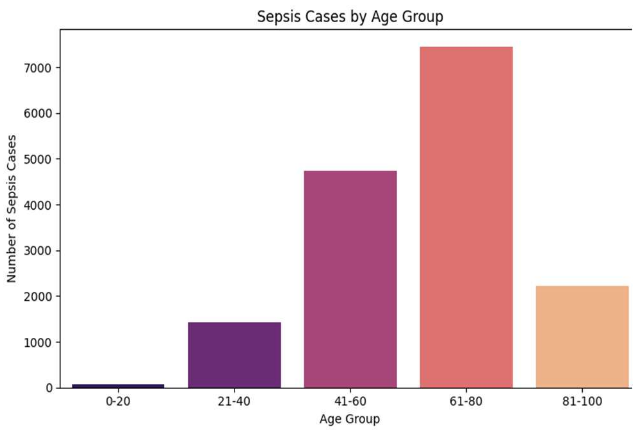

Fig. 2. Sepsis Cases by Age Group

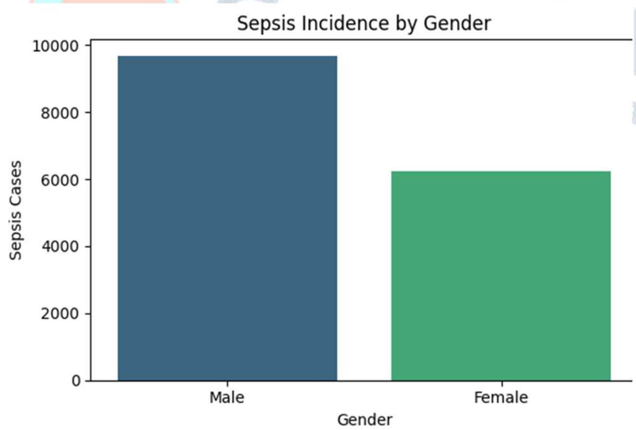

Fig. 3. Sepsis Incidence by Gender

1) Handling Missing Values:

Clinical data frequently suffers from gaps caused by inconsistent monitoring or equipment malfunctions. To handle these gaps, techniques such as forward-filling and statistical imputation using the mean or median are applied to estimate and restore the absent values.

2) Normalization:

To bring features with different scales into a comparable range, **Min-Max normalization** is applied:

$$x_t^{(i)} = \frac{x_t^{(i)} - \min\left(x^{(i)}\right)}{\max(x^{(i)}) - \min(x^{(i)})}$$

This feature x(i) to a [0, 1] range, which improves convergence speed and model stability during training.

3) Label Encoding:

Sepsis labels are binary, where:

y = {1, if sepsis is present}

y = {0, if no sepsis is present}

This binary classification target is used with the binary cross-entropy loss function during model training.

C. Model Architecture

The model is based on a Multi-Layer Perceptron (MLP), a type of feedforward neural network for structured data. The input to the MLP is a fixed-size vector $x \in R^n$containing patient features such as vitals, labs, and demographics.

The architecture includes:

- ***Input Layer****:* Accepts the normalized patient feature vector.
- ***Hidden Layers****:* Multiple fully connected (dense) layers with ReLU (Rectified Linear Unit) activation functions to introduce non-linearity.
- ***Dropout Layers****:* Applied between hidden layers to prevent overfitting by randomly dropping a fraction of neurons during training.
- ***Output Layer***: A neuron with a sigmoid activation function to output the probability of sepsis:

$$y = \sigma(Wx + b),$$

$$\text{where } \sigma(z) = \frac{1}{1 + e^{-z}}$$

ISSN [Online]: 2583-2654

The MLP model captures complex non-linear relationships between features without modeling temporal dependencies.

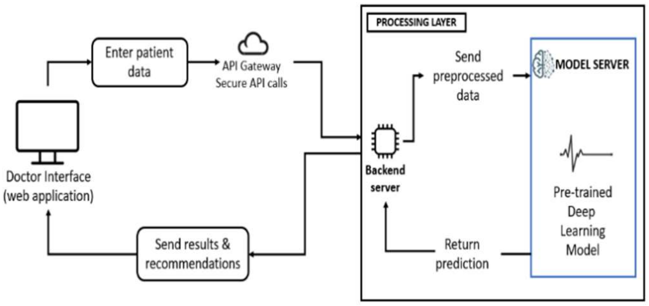

Fig 4: Architecture Diagram

D. Model Training and Evaluation

The MLP-based model is trained on a pre-processed clinical dataset, where each sample encodes a patient's health status as a fixed-length feature vector. The data is split into training (80%), validation (10%), and testing (10%) subsets to promote generalization. During training, the model aims to minimize binary cross-entropy loss, defined as:

$$L = -\frac{1}{N}\sum [yi \cdot log(yi) + (1 - yi) \cdot log(1 - yi)],$$

where $yi \in \{0,1\}$ is the actual label for the i^{th} patient (1 indicating sepsis, 0 otherwise), and yi $\in[0,1]$ is the predicted probability output by the MLP model.

Batch normalization is used to stabilize and accelerate training, while **dropout layers** mitigate overfitting by randomly deactivating neurons during training. **Early stopping** is employed to pause training when validation performance no longer improves, thereby ensuring optimal generalization.

The evaluation of the model is conducted using the following performance metrics:

Accuracy (Acc):

$$Acc = \frac{TP + TN}{TP + TN + FP + FN}$$

Precision:

$$Acc = \frac{TP}{TP + FP}$$

Recall (Sensitivity):

$$Recall = \frac{TP}{TP + FN}$$

F1-Score:

$$F1 = 2 \cdot \frac{Precision \cdot Recall}{Precision + Recall}$$

AUC-ROC : This metric is used to evaluate model's ability to discriminate between septic and non-septic cases. **Confusion matrices** and **ROC curves** are generated to visualize the effectiveness of the classification process. The **MLP model** is continuously assessed on the validation set, and the model with high **F1-score** is selected for final testing.

Multi-Layer Perceptron Algorithm:

Input: Time-series clinical data $X = \{x_1, x_2, ..., x_n\}$.

Output: Sepsis prediction (Normal, Severe, Septic Shock) and risk score.

Begin

1. Preprocess data: Clean, normalize, and encode clinical features. Split dataset into training and testing sets.

2. Initialize MLP parameters: Randomly initialize weights and biases for each layer of the MLP.

3. Forward Propagation through Layers:

- Compute activations for hidden layer 1:

$H_1 = \text{ReLU}(W_1 \cdot X + b_1)$

- Compute activations for hidden layer 2:

$H_2 = \text{ReLU}(W_2 \cdot H_1 + b_2)$

ISSN [Online]: 2583-2654

- Apply dropout regularization (if used).
- Compute output layer:

$\hat{Y} = \text{Softmax}(W_3 \cdot H_2 + b_3)$ *(for multi-class classification)*

or

$\hat{Y} = \text{Sigmoid}(W_3 \cdot H_2 + b_3)$ *(for binary classification)*

4. Loss Computation: Compute the loss using cross-entropy between predicted labels and ground truth.

5. Backpropagation and Optimization: Adjust the model's weights through iterative learning, guided by an optimizer, until convergence is reached or early stopping is triggered.

End

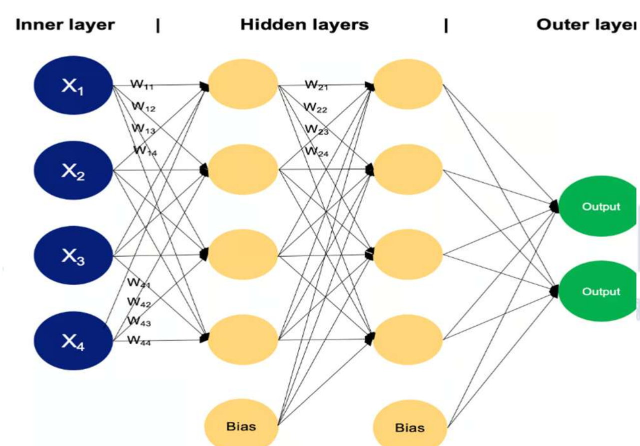

Fig 5: Multilayer Perceptron Algorithm

IV. RESULTS

The developed system, titled "Optimizing Sepsis Care with AI: Integrating Deep Learning into clinical decision support system", showcases the performance of a Multi-Layer Perceptron (MLP) model in the early prediction of sepsis. The model was trained on the Kaggle Sepsis dataset, comprising structured clinical and physiological time-series data and lab test results. The final MLP model achieved a high overall accuracy of 98%, indicating strong general classification performance. However, detailed evaluation revealed a class imbalance in the dataset, where the model showed excellent performance in identifying non-septic cases (class 0) with precision of 0.98 and recall of 1.00, but limited ability to detect septic cases (class 1), with precision of 0.68 and recall of 0.06. The F1-score for the septic class was 0.11, reflecting the model's difficulty in capturing rare but critical sepsis cases. Despite this, the macro average F1-score was 0.55, and the weighted average F1-score reached 0.97. It demonstrates the robustness of the majority class model and highlights how its performance could be further improved through class equalization techniques.

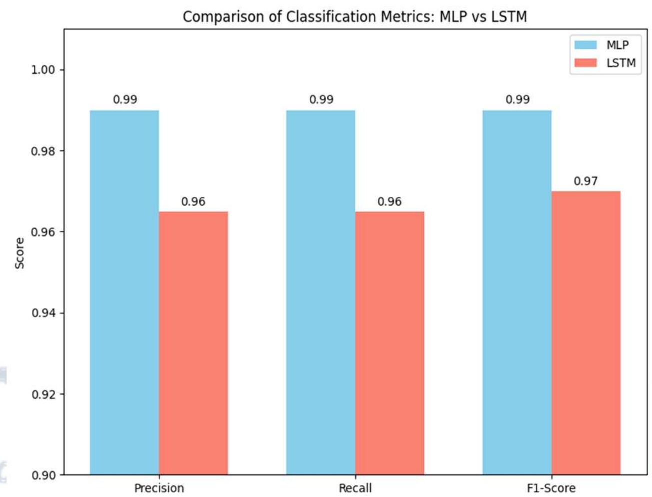

Fig 6: Comparison of Classification Metrics: MLP vs LSTM

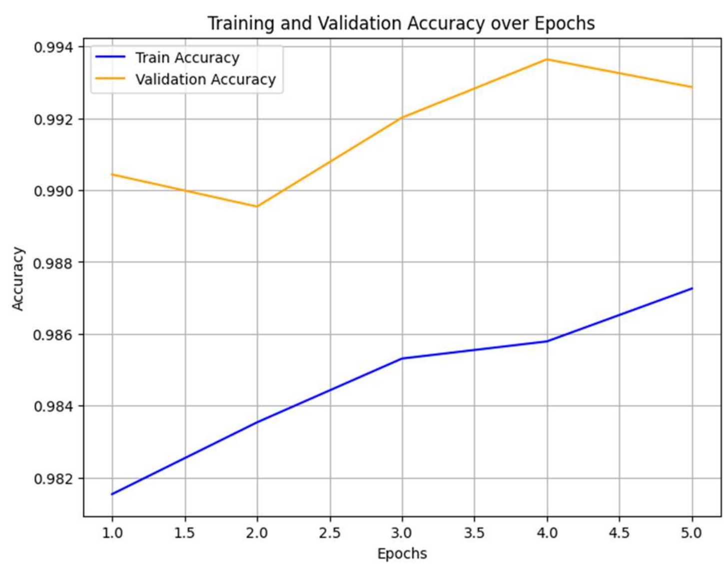

Fig 7: Training and Validation Accuracy Over Epochs

ISSN [Online]: 2583-2654

♦ MLP Classification Report

	precision	recall	f1-score	support
0.0	1.00	0.98	0.99	12333
1.0	0.98	1.00	0.99	12333
accuracy			0.99	24666
macro avg	0.99	0.99	0.99	24666
weighted avg	0.99	0.99	0.99	24666

Fig 8: Classification report of MLP

The MLP-based Clinical Decision Support System (CDSS) includes features such as real-time classification of patient sepsis risk, severity categorization, and clinical recommendations like medication suggestions and precautionary alerts. The system is deployed through a Flask web application with a hospital-themed user interface and integrated MongoDB database for efficient patient data management. Healthcare professionals can upload patient vitals and lab data to receive immediate predictions, enabling timely interventions in ICU environments.

V. CONCLUSION

This presented work attempts to demonstrates the effective use of a Multi-Layer Perceptron (MLP) model for early prediction and classification of sepsis severity, including normal, severe, and septic shock cases. By leveraging structured clinical data, the system provides timely, data-driven support to healthcare professionals, helping reduce diagnosis delays and improve critical care outcomes. The Flask-based user interface, ensures a user-friendly and scalable solution suitable for clinical deployment.

Future enhancements may include the incorporation of real-world hospital data to improve model generalization, implementation of explainable AI techniques like SHAP or LIME for transparency, and support for real-time alerts, multimodal inputs, and multilingual access to boost clinical usability and adoption.

REFERENCES

[1] X. Tang, Y. Liu, and Z. Wang, "Optimizing Sepsis Treatment Through Reinforcement Learning: A Revisitation of Reward and Loss Functions in Dueling Double Deep Q-Network," IEEE Access, vol. 10, pp. 12345–12356, 2024.

[2] P. V. C. Souza and M. Dragoni, "PEFNN: Parallel Evolving Fuzzy Neural Network for Sepsis Identification in Patients," IEEE Transactions on Fuzzy Systems, vol. 32, no. 1, pp. 25–36, Jan. 2024.

[3] S. Babu, L. A. Annavarapu, and L. L. Appala, "Sepsis Detection using Neural Networks," in 2022 International Conference on Smart Generation Computing, Communication and Networking (SMART GENCON), pp. 1-5, 2022.

[4] Z. He, X. Chen, Z. Fang, W. Yi, C. Wang, L. Jiang, Z. Tong, Z. Bai, Y. Li, and Y. Pan, "Early Sepsis Prediction Using Ensemble Learning with Features Extracted from LSTM Recurrent Neural Network," Computing in Cardiology 2019, vol. 46, pp. 1-4, 2019, doi: 10.22489/CinC.2019.269.

[5] S. K. Rout, G. B. Regulwar, V. Kavididevi, and B. Sahu, "Deep Learning in Early Prediction of Sepsis and Diagnosis," in 2023 International Conference for advanced Technology (ICONAT), pp. 1-4, 2023, doi: 10.1109/ICONAT57137.2023.10080152.

[6] Erik H. Gilbertson, K. M. Jones, A. M. Stroh, and B. M. Whitaker, "Early Detection of Sepsis Using Feature Selection, Feature Extraction, and Neural Network Classification," Computing in Cardiology 2019, vol. 46, pp. 1-4, 2019, doi: 10.22489/CinC.2019.386.

[7] R. M. Demirer and O. Demirer, "Early Prediction of Sepsis from Clinical Data Using Artificial Intelligence," IEEE, pp. 1-4, 2019.

[8] Mohammed Saqib, Y. Sha, and M. D. Wang, "Early Prediction of Sepsis in EMR Records Using Traditional ML Techniques and Deep Learning LSTM Networks," IEEE, pp. 4038-4041, 2018.

[9] S. Babu, L. A. Annavarapu, and L. L. Appala, "Sepsis Detection using Neural Networks," in 2022 International Conference on Smart Generation Computing, Communication and Networking (SMART GENCON), pp. 1-5, 2022.

[10] Anagha Anand, Samhitha S, Malavika Hariprasad, Jyoti Shetty, Nimisha Dey, and Shobha G, "Examining Deep Learning Methods For The Detection Of Sepsis," in 2022 6th International Conference On Computing, Communication, Control And Automation (ICCUBEA), pp. 1-6, 2022.