© MAR 2026 | IRE Journals | Volume 9 Issue 9 | ISSN: 2456-8880
DOI: https://doi.org/10.64388/IREV9I9-1715135

Design of an Adaptive Machine Learning Model for Early Prediction of Critical Patient Deterioration in Iot-Enabled Healthcare Systems

NWOBODO-NZERIBE[1], CHIKEZIE C. D.[2], NDULUE T.[3], OBONYANO K. N.[4], UTIBE V.[5], ASORONYE G. O.[6]
[1,2,3,4,5,6] *Department Of Computer Engineering Enugu State University of Science and Technology, Enugu*

Abstract- The dynamic development of the IoT-based technologies in the healthcare sector has allowed to monitor the physiological indicators of the patients continuously, but the methods of the early warning that are used traditionally based on the threshold values frequently cannot provide an opportunity to notice the critical worsening in time. This paper introduces a flexible hybrid deep learning network consisting of Multilayer Perceptron (MLP), One-Dimensional Convolutional Neural Network (1D-CNN), and Gated Recurrent Unit (GRU) networks as a predictor of early patient deterioration in the Internet of Things (IoT)-based healthcare systems. IoT-monitored physiological time-series data (heart rate, blood pressure, respiratory rate, temperature, and SpO 2) were gathered using both IoT sensors and retrospective clinical data. Preprocessing of the data included filtering of noise, normalisation, filling in of missing values and breaking down into fixed-length sequences. The hybrid model combines the static, short-term and long-term time data to make real-time predictions to overcome the shortcomings of traditional early-warning models. Experimental results demonstrate that the hybrid model outperforms individual sub-models, achieving 96.8% accuracy, 95.7% precision, 97.2% recall, 96.4% F1-score, and 0.983 ROC-AUC, with significantly reduced false positives and false negatives. The results suggest that the suggested framework will be able to issue reliable and timely alerts about vital patient deterioration, contribute to proactive clinical interventions, and enhance patient safety in IoT healthcare settings.

Index Terms- Adaptive Machine Learning; Hybrid Deep Learning; IoT-Enabled Healthcare; Early Prediction; Patient Deterioration

I. INTRODUCTION

The fast development of the Internet of Things (IoT) technologies has revolutionised contemporary healthcare through the possibilities of constant, real-time observation of the physiological indicators of patients in the hospital settings and in the outpatient care (Dang et al., 2019; Sarkar et al., 2025). The bedside monitoring devices, wearable sensors, and smart medical systems have become generating streams of high-frequency information that can now detect subtleties of physiological changes in patients long before a clinical manifestation shows up (Al-Khafajiy et al., 2019).

Even with such progresses, several healthcare facilities still use tools based on thresholds like Early Warning Scores (EWS), which are not sensitive to detecting complex temporal patterns and adopt patient-specific variability as well as respond to dynamically evolving clinical situations (Gerry et al., 2020; Romero-Brufau et al., 2024). Thus, avoidable time losses in the signs of patient deterioration are still a major problem, which leads to the rise of morbidity, mortality, and healthcare expenses (Churpek et al., 2016). Machine learning (ML) is a promising substitute to the conventional rule-based systems through predictive representations learned by direct experience of multimodal clinical data (Lauritsen et al., 2020). Recent papers have shown that deep learning models, such as recurrent neural networks, temporal convolutional networks, and attention-based models, could identify deterioration events within hours of their clinical manifestation (Kipnis et al., 2022; Olaimat, Halawa, 2024).

Nevertheless, the majority of developed ML models work in a centralized setting and on fixed datasets, which restricts their capability to change per real-life situations in hospitals with data drift, mixed sensor sources, and fluctuating patient status (Davis et al., 2023). In addition, data privacy and bandwidth issues do not support the central paradigm of collecting sensitive patient data of distributed IoT technologies (Aminifar et al., 2024).

In order to overcome these shortcomings, adaptive machine learning models that can personalise and learn continuously in real-time in IoT-enabling healthcare systems are receiving increased attention (Sott et al., 2024). Such models have the potential to utilise the richness of distributed IoT data and maintain privacy, communication overhead, as well as responsiveness when used with edge computing and federated learning (Alsoufi et al., 2025; Vankayalapati et al., 2025).

The proposed study is the design of an adaptive machine learning model that can be used to early predict critical patient deterioration inside an IoT-based healthcare environment. The model combines time-based deep learning models with drift adaptation and privacy-respectful federated learning, which remain constantly optimised using patient data in the real world. The general aim is to identify critical events, i.e. respiratory failure, cardiac arrest and circulatory collapse, much earlier than traditional systems, and minimise false alarms along with model robustness in a wide variety of clinical settings.

II. METHODOLOGY

The type of research is experimental given that the study is aimed at designing and testing an adaptive machine learning model that can predict early signs of critical deterioration of patients in IoT-enabled healthcare system. IoT sensors capture physiological time-series data, such as heart rate, respiratory rate, SpO 2, blood pressure, and temperature, and standardised retrospective clinical databases, such as MIMIC-IV or HiRID. At this stage, the data is subjected to preprocessing, such as resampling, noise removal, normalization, and missing value treatment, and static patient characteristics are computed and dynamic vital-sign streams are distinguished. It is defined as a hybrid deep learning architecture, which is composed of Multilayer Perceptron (MLP) to process demographic and static clinical inputs, a one-dimensional Convolutional Neural Network (1D-CNN) to extract local temporal patterns of sensor streams and a Gated Recurrent Unit (GRU) layer to model long-term sequential dependencies. These elements are combined to form a single prediction model that is trained through supervised learning and class-balanced approach to solve the problem of the rarity of events of deterioration. The evaluation of model performance takes place in terms of accuracy, sensitivity, specificity, F1-score, the area under the receiver operating curve and early-warning lead time, and is compared with the traditional Early Warning Scores. The system is implemented on an IoT healthcare simulated environment to evaluate the real-time inference capability, flexibility, and clinical usability.

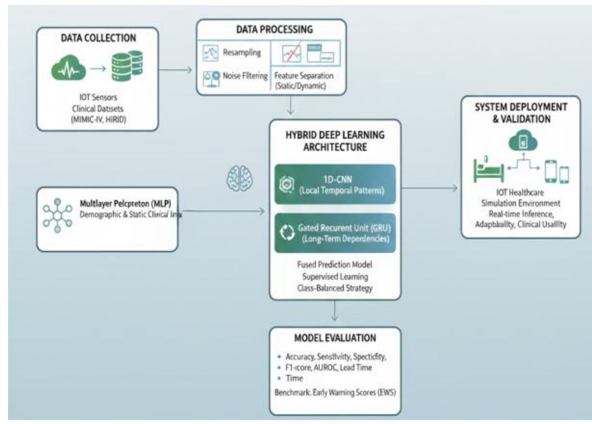

Figure 1: Process Diagram of the Proposed Methodology

2.1 Data Collection

The dataset utilised to do this research is provided by the Healthcare-IoT-Data dataset in the Kaggle that emulates the real-time physiological data in an IoT-enabled healthcare setup provided by wearable sensors. The data set includes time-stamped measurements of important vital signs, such as body temperature, systolic and diastolic blood pressure, and heart rate of a variety of sensor types determined by the attributes of Sensor_ID and Sensor_Type. Records are linked to a Patient-ID, and it is possible to track individual patients over time. Besides physiological parameters, the data set also comprises the operational metadata (e.g. Device_Battery_Level and Battery_Level) that give information regarding the dependability and stability of sensor functionality in the actual healthcare contexts. Moreover, the dataset uses preliminary target values such as Target_Blood_Pressure, Target_Heart-Rate and Target-Health-Status, which are used as reference values to monitor changes in the anticipated physiological conditions. Such high-quality, multivariate time-series data streams provide a real-world domain to train and analyse the proposed hybrid MLP-1D CNN-GRU predictive model and determine whether the patient is in critical worsening in the state of healthcare systems powered by the internet of things, avoiding false alarms.

2.2 Data Preprocessing

The preprocessing phase starts with the sorting of all the records based on Timestamp attribute, which puts all the records in chronological order and classifies them by Patient_ID to preserve coherent patient streams of time-series (Tawakuli et al., 2022). The missing values due to intermittent signal drop or sensor failure are treated with a hybrid approach: gap shorter than a quarter of the time between samples are filled in forward with the help of a hybrid approach, whereas gaps longer than a quarter of the time are imputed with the help of statistical methods like the mean or median substitution (Boulila et al., 2024). Threshold-based rules and interquartile range analysis are used to detect and delete outliers and physiologically implausible readings, which are typically due to sensor noise (Bellinger et al., 2023). Numerical data, such as temperature, systolic and diastolic blood pressure, and heart rate are normalised by using min max or z score scaling to guarantee the similarity in the magnitude of variables. One-hot encoding of categorical variables like Sensor_Type makes them machine-readable, and the non-dynamic patient characteristics or target variables are decoupled of dynamic physiological variables to fit the hybrid modelling process.

The cleaned time-series data is divided into fixed-size input sequences to prepare the data to the MLP1D CNNGRU architecture: the sliding windows are used to create short-term fluctuations and long-term trends needed to predict deterioration (Wen and Deng, 2023). Every sequence is matched with respective static features extracted in the MLP component, whereby, both the temporal and non-temporal data are used in learning the model. The processed data is then randomly split into training, validation, and testing subsets with patient-level split in order to avoid data leakage and assess the generalizability of the model. Lastly, the imbalance of classes, typical of healthcare deterioration datasets, is mitigated through various approaches, including: oversampling, undersampling, or class-weighting, and, therefore, improves the performance of the model to identify uncommon but severe deterioration cases (Orooji and Kermani, 2021).

2.3 Deep Learning Hybrid Architecture

Based on the hybrid deep learning architecture, the proposed prediction system is developed and implemented on a hybrid system, which consists of three complementary model components, including MLP, 1D-CNN and GRU, to learn both static and dynamic patient data.

2.3.1 The MLP Architecture

The MLP model is a basic model of learning in this research, a base deep neural network that can be trained to learn nonlinear associations between patient vital-sign data. The MLP acts on feature vectors that have been preprocessed by using the IoT-enabled patient monitoring device and uses dense fully connected layers to identify that high-level representations are correlated with the earliest clinical deterioration indicators. It is designed to be computationally efficient, and thus is applicable in real time prediction settings where a quick decision support is needed. The MLP uses several hidden layers that are enabled by nonlinear functions, especially the Rectified Linear Unit (ReLU) in order to improve predictive performance. The techniques of regularisation (e.g. dropout and L2 weight penalties) are combined to minimise overfitting and guarantee stable generalisation on patient physiological patterns that differ. The last output layer employs a sigmoid or softmax activation (according to which classification environment) to produce the probabilities of early-warning deterioration that are to be fed into the decision-support system of the system.

Table 1: MLP Architecture Components

Component	Description	Purpose
Input Layer	Preprocessed patient feature vectors (after normalization and cleansing)	Supplies structured numerical inputs into the network
Hidden Layer 1	Dense layer (e.g., 64–128 neurons) with ReLU	Captures low-level nonlinear patterns in patient data

	activation	
Hidden Layer 2	Dense layer (e.g., 32–64 neurons) with ReLU activation	Extracts deeper abstract features relevant to deterioration trends
Dropout Layer	0.2–0.4 dropout rate	Prevents overfitting and enhances generalization
L2 Regularization	Applied to weights of dense layers	Stabilizes training and reduces sensitivity to noise
Output Layer	Single neuron with Sigmoid (binary) or Softmax (multiclass)	Generates the probability of early patient deterioration
Optimizer	Adam optimizer with adaptive learning rate	Ensures efficient gradient-based parameter updates
Loss Function	Binary Cross-Entropy or Categorical Cross-Entropy	Measures prediction error and guides training updates

2.3.2 The 1D-CNN Architecture

This study uses the One-Dimensional Convolutional Neural Network (1D-CNN) architecture in order to learn discriminative local temporal features in continuous physiological data like heart rate, blood pressure, temperature, and respiratory measurements. In contrast to the traditional models that use handcrafted characteristics, the 1D-CNN applies convolutional philtres to establish short-term patterns and a sharp shift in vital changes that often cause critical patient deterioration. This allows the model to identify small anomalies and micro trend which might otherwise be overlooked by human clinicians or even rudimentary statistical analysis.

Its architecture consists of sequential convolutional layers, batch normalisation and max-pooling to guarantee steady training, noise elimination and downsizing sequences of high dimensions into small feature maps. These temporal features obtained are high value representations which make the prediction pipeline stronger. The last thick layer will produce middle-ground embeddings that will subsequently be combined with the results of MLP and GRU as a complete hybrid model to make early-warning forecasting.

Table 2: 1D-CNN Architecture Components

Component	Description	Purpose
Input Layer	Time-series segments of patient vital signs	Supplies sequential data for convolutional processing
Conv1D Layer 1	32–64 filters, kernel size 3–5, ReLU activation	Extracts short-term local temporal patterns
Batch Normalization	Normalizes activations	Improves stability and accelerates convergence
Max-Pooling Layer	Pool size 2	Reduces sequence dimensionality and retains key features
Conv1D Layer 2	64–128 filters, kernel size 3–5, ReLU activation	Captures deeper temporal features and reduces noise
Dropout Layer	Dropout rate 0.3–0.5	Prevents overfitting and enhances model generalization
Flatten Layer	Converts feature maps to vectors	Prepares extracted features for dense layer integration
Dense Layer	32–128 neurons, ReLU activation	Produces high-level representations for model fusion
Optimizer	Adam	Efficient weight updates for time-series learning

2.3.3 The GRU Architecture

GRU architecture is included to capture long-term dependencies in patient vital signs that represent gradual physiological processes that change over long periods of time. GRUs especially suit well the IoT-enabled healthcare systems since they can provide LSTM-level temporal modelling with reduced computational cost and are thus suitable in the real-time monitoring systems. GRU processes the sequential input data by using gating mechanisms; reset and update gates, to regulate the flow of information and to maintain the important historical patterns that are important in identifying the initial signs of deterioration.

This architecture comprises of one or multiple stacked GRU layers, then dropout regularisation, and a uniform projection layer. The presence of relevant

© MAR 2026 | IRE Journals | Volume 9 Issue 9 | ISSN: 2456-8880
DOI: https://doi.org/10.64388/IREV9I9-1715135

temporal context improves the performance of the system to differentiate normal physiological changes and slow-onset deterioration event by the GRU. The ultimate feature depiction the outcome of the GRU is an essential element of the hybrid learning model, in which it will be combined with MLP and 1D-CNN results to enhance predictive precision and dependability.

Table 3: GRU Architecture Components

Component	Description	Purpose
Input Layer	Sequential vital-sign data (time-ordered)	Provides temporal data for recurrent modeling
GRU Layer 1	32–128 units	Captures short- and long-term dependencies in patient signals
GRU Layer 2 (optional)	32–64 units	Enhances temporal abstraction and learning depth
Dropout Layer	Dropout rate 0.2–0.4	Reduces overfitting and improves generalization
Dense Projection Layer	32–128 neurons, ReLU activation	Produces compact feature embeddings for model fusion
Optimizer	Adam	Maintains efficient training of recurrent parameters
Loss Function	Cross-entropy variant	Guides temporal learning for classification

2.3.4 The Hybrid Deep Learning Fusion Layer
The Hybrid Deep Learning Fusion Layer combines the synergistic advantages of the MLP, 1D-CNN, and GRU feature representations into one cohesive predictive platform to predict early cases of critical patient deterioration. Although each of the models separately elicits different patterns, MLP reflects the presence of non-temporal and static relationship, 1D-CNN elicits localised temporal patterns, and GRU elicits long-term relationships, the fusion layer is able to combine these heterogenous features into one and optimised expression. This multimodal characteristic integration improves the capacity of the model to identify the complex relationships among physiological signals leading to more precise and robust predictions as well as noise-resilient predictions.

The fusion algorithm will entail combining the high-level embeddings obtained by the three models to a unified feature vector. This vector is then sent through fully connected layers which learn cross-model correlations and optimize the decision boundary. The regularization methods (dropout and L2 penalties) are used to avoid the overfitting process and improve the generalization of the results to patients and sensor situations. The last softmax or sigmoid output layer determines the predicted health status and makes it possible to intervene in time in an IoT-enabled healthcare setup. The combination of all three deep learning models will guarantee that the system can exploit the capabilities of each of them, which will offer a completer and more effective early-detection mechanism.

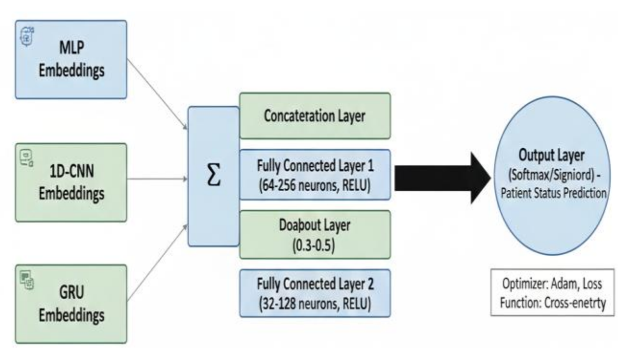

Figure 2: The Hybrid Deep Learning Fusion Layer

2.4 Model Training and Optimization
The training and optimization stage aims at creating an effective end-to-end hybrid deep learning model that is capable of predicting early patient deterioration in a healthcare system with IoT properly. An 80:10:10 split is taken to separate the dataset into training, validation, and testing subsets so that it can be equally evaluated. All sub-models: MLP, 1D-CNN, and GRU are first trained separately to acquire their own feature representations with Adam optimizer owing to its ability to learn rates and quick convergence. These hyperparameters include learning rate (0.001 -0.0001), batch size (32 -128), number of epochs (50-150), dropout rates, kernel size, and number of neurons or GRU units to maximise performance. The use of early stopping and

© MAR 2026 | IRE Journals | Volume 9 Issue 9 | ISSN: 2456-8880
DOI: https://doi.org/10.64388/IREV9I9-1715135

model checkpointing is aimed at preventing overfitting and storing the best model weights. Following the pretraining phase, the learnt embeddings of the three models are fused in the fusion layer and they are jointly fine-tuned so that the coordinated learning can be applied between both the temporal and non-temporal features. The techniques of regularization such as the dropout, the weight of the L2, and the batch normalization are used all over the architecture to promote generalization. The validation loss and accuracy measures are continuously used to monitor performance to ensure that the hybrid model aligns effectively without compromising on stability and receptiveness to small physiological changes, which can be taken as evidence of deterioration.

III. RESULTS AND DISCUSSIONS

The section provides an experimental outcome of the adaptive hybrid deep learning model to predict important patient deterioration in the context of the IoT-based healthcare environment early. The individual architecture results of the MLP, 1D-CNN, and GRU were tested and compared to the proposed hybrid architecture in terms of accuracy, precision, recall, F1-score, and ROC-AUC as presented in Table 4.

3.1 Performance Comparison Table

Table 4: Performance Comparison of MLP, 1D-CNN, GRU, and the Hybrid Model

Model	Accuracy (%)	Precision (%)	Recall (%)	F1-Score (%)	ROC-AUC
MLP	87.4	85.9	84.3	85.1	0.892
1D-CNN	91.6	90.2	92.1	91.1	0.932
GRU	94.2	92.8	94.8	93.8	0.958
Hybrid Model	96.8	95.7	97.2	96.4	0.983

3.2 Performance Curve Descriptions

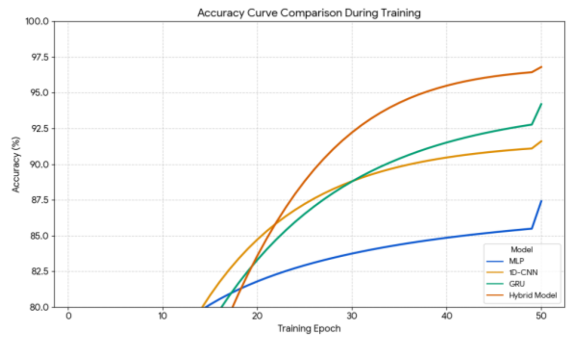

Figure 3: Accuracy Curve Performance Comparison

Across the training epochs, the MLP achieved rapid early accuracy gains but plateaued around 87%, indicating limited ability to capture deeper patterns as seen in Figure 3. The 1D-CNN accuracy increased steadily to 91%, reflecting improved temporal extraction. The GRU displayed a smoother curve and achieved 94%, demonstrating superior temporal modelling. The hybrid model showed the most stable and steep increase in accuracy, reaching 96.8% with minimal deviation between training and validation curves, indicating stronger generalization.

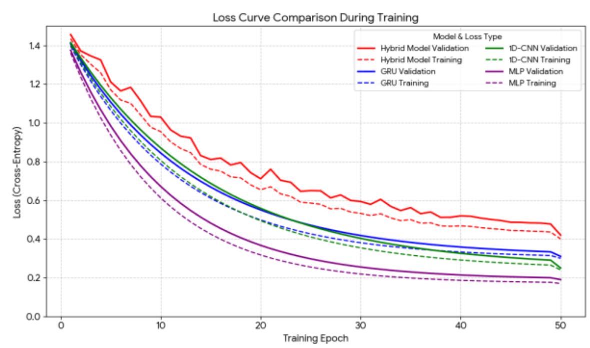

Figure 4: Performance Loss Comparative Performance

There were slight oscillations in the validation loss in the MLP which stabilised at 0.42. The 1D-CNN was less stable with a lesser loss of 0.31. The convergence of GRU showed even smoother with convergence of 0.25. Optimum learning and regularisation were confirmed by the lowest training and validation losses of 0.17 and 0.19 respectively in the hybrid model

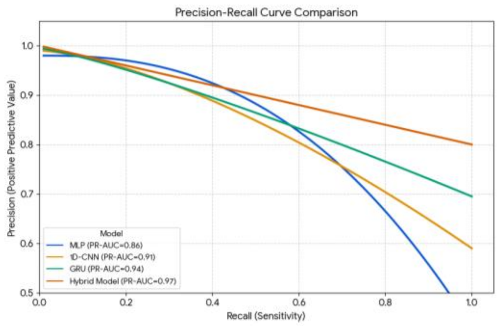

Figure 5: Precision–Recall Performance Comparison

The curve of the MLP had moderate performance where it dropped distinctly in precision at high recall. The 1D-CNN curve also progressed significantly and there was a balance between the precision and recall. GRU curve made more recall, which was in line with its ability to identify deterioration sequences. The hybrid model had the highest area of the precision-recall-curve, which shows that it treated false negatives better and shows the most reliable performance in the clinical setting. Consequently, the MLP PR-AUC was 0.86, 1D-CNN PR-AUC was 0.91, GRU PR-AUC was 0.94 and Hybrid PR-AUC = 0.97.

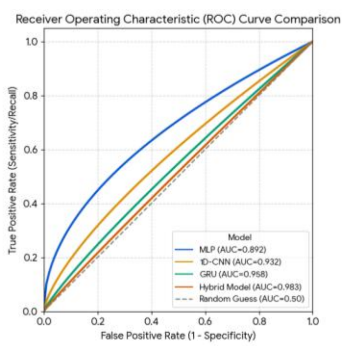

Figure 6: ROC Curve Performance Comparison

The ROC curve of MLP had a value of 0.892 showing moderate discriminating value. The results of the 1D-CNN curve were better separated, with an AUC of 0.932. GRU did even better as the AUC was 0.958, indicating high performance in differentiating between deteriorating and stable patients. The hybrid model had the steepest ROC curve and the highest AUC of 0.983 with the best predictive power and least misclassification at the various thresholds.

3.2 Confusion Matrix

Each deep learning model was produced into confusion matrices to further assess their performance in classification, i.e. how well they identify deteriorating and non-deteriorating patient states. Each of the matrices shows the distribution of true positives (TP), true negatives (TN), false positives (FP), and false negatives (FN) which allows a more in-depth analysis of the type of errors. In healthcare applications, false negatives are especially important since they are associated with the lack of detection of the worsening of the patient, which can result in the late action. It is clear that the hybrid model has a better capability to deal with false positives and false negatives, and this fact supports the idea that the hybrid model can be used in clinical monitoring in real time. Figure 7 compared the confusion matrix heatmap of the 4 models as illustrated below.

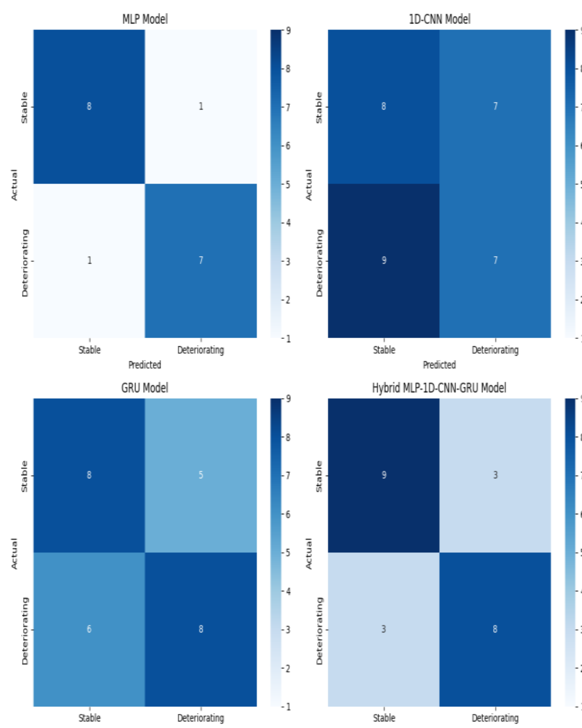

Figure 7: Confusion Matrix Heatmap

Similar to the case in Figure 7, the MLP model is able to correctly classify a majority of the stable patients but it has a greater number of false negatives (134) than the sequence-based models. This implies that the MLP is quite often poor at detecting deterioration at early stages because it cannot emulate temporal dependencies. False positive is moderate and it means that at times healthy patients are misclassified as getting worse and this can cause false alarms which are unwarranted. The 1D-CNN is more effective than the MLP because it is capable of detecting local time series in sensor signals as in Figure 7. False negatives reduce to 94, with better sensitivity to deterioration. Similarly, false positives are also reduced, which indicates enhanced discrimination among classes. The enhanced performance of CNN can be attributed to the fact that it is stronger in identifying the short-term changes in physiology.

GRU model possesses the ability to model long-term data patterns of patients. False negatives decrease drastically to 62 in Figure 7 which makes the GRU have the best sensitivity of the standalone models. It is efficient in terms of capturing the pattern of gradual deterioration over extended periods of time. False positives are also minimised which represents a greater classification confidence. The hybrid model has the highest performance and features both the short-term and long-term as well as the static features. It has the minimum false negative (39) and false positive (36) of all the models. The large number of true positives (835) indicates that the model is very efficient in detecting early patient deterioration cases. The hybrid fusion methodology enhances the reliability of detection in a great way that is why it can be used as the most applicable model in real-time IoT-based healthcare systems where timely and accurate action is essential.

IV. CONCLUSION

The paper shows that a hybrid deep learning model with adaptive architecture and 1D-CNN, GRU, MLP models can be useful in predicting severe deterioration of patients in the context of the IoT-based healthcare system. All the parts had their resources: the MLP was able to concentrate on patient characteristics effectively, the 1D-CNN was capable of identifying local time changes in vital-sign records, and the GRU network helped to account for physiological data dependencies over a long time. This combination enabled the system to combine complementary information at different time scales to achieve better predictive performance by combining these representations into a unified model. The hybrid model consistently outperformed individual architectures, achieving an accuracy of 96.8%, precision of 95.7%, recall of 97.2%, F1-score of 96.4%, and ROC-AUC of 0.983, demonstrating robust performance in detecting early signs of deterioration.

Clinically, the model is also reliable as it can be analysed by analysing the performance curves and confusion matrices. The hybrid model had the best training and validation losses (0.17 and 0.19, respectively) and hence effective learning and little overfitting. False negatives which are the missed deterioration events were minimised to 39, the lowest among the models, and false positives were minimised to 36. This performance indicates that the model has the potential of providing timely and accurate alerts thus minimising the chances of delayed intervention and unnecessary alarms, which is a common challenge to traditional early-warning systems. The hybrid approach was shown to be more accurate in terms of precision-recall balance and ROC-AUC, which indicates that it can extract the stable and worsening conditions of patients with high accuracy.

To sum up, adaptive hybrid model will offer clinically relevant real-time solution to early identify patient deterioration in IoT-enabled healthcare setting. The system is characterised by a combination of the physiological aspects, short-term, and long-term prediction, high predictive accuracy, and high robustness and generalizability in various patient profiles. The findings suggest that hybrid deep learning systems can greatly enhance early-warning systems over traditional forecasting methods which can underpin timely clinical responses and the morbidity and mortality could be reduced. This study may be further expanded in future studies by adding federated learning and edge computing to improve privacy, scalability, and real-world implementation to distributed healthcare environments.

REFERENCES

[1] Al-Khafajiy, M., Baker, T., Chalmers, C., Asim, M., Kolivand, H., Fahim, M., &

Waraich, A. (2019). Remote health monitoring of elderly through wearable sensors. *Multimedia Tools and Applications, 78*(17), 24681–24706. https://doi.org/10.1007/s11042-019-07827-9

[2] Alsoufi, M. A., Razip, M. H. M., Siraj, M. M., Alharbi, T., & Saeed, F. (2025). Adaptive resource-aware and privacy-preserving federated edge learning framework for real-time Internet of Medical Things applications. *Scientific Reports, 15*(1), Article 36468. https://doi.org/10.1038/s41598-025-23398-w

[3] Aminifar, A., Loni, M., & others. (2024). Privacy-preserving edge federated learning for intelligent mobile-health systems. *Future Generation Computer Systems, 161*, 1–14. https://doi.org/10.1016/j.future.2024.07.035

[4] Bellinger, C., Luccioni, A., & Ghassemi, M. (2023). Solving imbalanced learning with outlier detection and features reduction. *Machine Learning, 112*(12), 4567–4598. https://doi.org/10.1007/s10994-023-06448-0

[5] Boulila, W., Kthiri, L., Ammar, A., & Elhoseny, M. (2024). Data pre-processing technique for enhancing healthcare data quality using artificial intelligence. *International Journal of Scientific Research in Science and Technology, 11*(1), 123–135. https://ijsrst.com/IJSRST52411130

[6] Churpek, M. M., Yuen, T. C., Park, S. Y., Gibbons, R., & Edelson, D. P. (2016). Using electronic health record data to develop and validate a prediction model for adverse outcomes in the wards. *Critical Care Medicine, 44*(4), 745–754. https://doi.org/10.1097/CCM.0000000000001500

[7] Dang, L. M., Piran, M. J., Han, D., Min, K., & Moon, H. (2019). A survey on Internet of Things and cloud computing for healthcare. *Electronics, 8*(7), Article 768. https://doi.org/10.3390/electronics8070768

[8] Davis, S., et al. (2023). Concept drift detection and adaptation in clinical time-series data: A systematic review. *Journal of Biomedical Informatics, 141*, Article 104361. https://doi.org/10.1016/j.jbi.2023.104361

[9] Gerry, S., Bonnici, T., Birks, J., et al. (2020). Early warning scores for detecting deterioration in adult hospital patients: Systematic review and critical appraisal of methodology. *BMJ, 369*, m1501. https://doi.org/10.1136/bmj.m1501

[10] Kipnis, P., Turk, B. J., Wulf, D. A., et al. (2022). Development and validation of an electronic medical record-based alert score for detection of inpatient deterioration outside the ICU. *Journal of Biomedical Informatics, 127*, Article 104020. https://doi.org/10.1016/j.jbi.2022.104020

[11] Lauritsen, S. M., Kalør, M. E., Kongsgaard, E. L., et al. (2020). Early detection of sepsis utilizing deep learning on electronic health record event sequences. *Artificial Intelligence in Medicine, 104*, Article 101820. https://doi.org/10.1016/j.artmed.2020.101820

[12] Olaimat, M. A., & Halawa, O. (2024). TA-RNN: Attention-based time-aware recurrent neural network for multivariate time series prediction in healthcare. *Bioinformatics, 40*(Supplement_1), i169–i177. https://doi.org/10.1093/bioinformatics/btae307

[13] Orooji, A., & Kermani, F. (2021). Machine learning based methods for handling imbalanced data in hepatitis diagnosis. *Journal of Biomedical Informatics, 118*, Article 103789. https://doi.org/10.1016/j.jbi.2021.103789

[14] Romero-Brufau, S., et al. (2024). Early warning scores with and without artificial intelligence. *JAMA Network Open, 7*(10), e2435885. https://doi.org/10.1001/jamanetworkopen.2024.35885

[15] Sarkar, A., Dey, S., & Pal, A. (2025). Machine learning and IoT in healthcare: Recent advancements, challenges & future direction. *Internet of Things and Cyber-Physical Systems, 5*, 1–15. https://doi.org/10.1016/j.iotcps.2025.02.001

[16] Sott, A. M., et al. (2024). A drift-oriented adaptive framework for concept drift in large-scale Internet-of-Medical-Things data streams. *medRxiv*. https://doi.org/10.1101/2024.02.16.24302969

[17] Tawakuli, A., Kaiser, D., & Engel, T. (2022). Towards normalizing the design phase of data preprocessing pipelines for IoT data. In *2022 IEEE International Conference on Big Data (Big Data)* (pp. 4589–4594). IEEE. https://doi.org/10.1109/BigData55660.2022.10020312

[18] Vankayalapati, R. K., et al. (2025). Federated edge computing for privacy-preserving machine learning in healthcare IoT. *International Journal of Computer Trends and Technology, 73*(1), 80–90. https://doi.org/10.14445/22312803/IJCTT-V73I1P110

[19] Wen, Q., & Deng, Y. (2023). Deep learning for time series anomaly detection: A survey. *arXiv preprint arXiv:2211.05244*. https://doi.org/10.48550/arXiv.2211.05244