iJRASET
International Journal For Research in Applied Science and Engineering Technology



INTERNATIONAL JOURNAL FOR RESEARCH

IN APPLIED SCIENCE & ENGINEERING TECHNOLOGY

Volume: 14 **Issue:** III **Month of publication:** March 2026

DOI: https://doi.org/10.22214/ijraset.2026.77694

www.ijraset.com

Call: 08813907089 | E-mail ID: ijraset@gmail.com

International Journal for Research in Applied Science & Engineering Technology (IJRASET)
ISSN: 2321-9653; IC Value: 45.98; SJ Impact Factor: 7.538
Volume 14 Issue III Mar 2026- Available at www.ijraset.com

Foresight-ICU: AI-Driven Real-Time Multi-Modal ICU Patient Monitoring System

Dr.K.Jagadeesh[1], Karthik A[2], Mukilesh Kumar S[3], Nivash S[4], Nithyasri R[5]
[1]*Associate Professor,* [2, 3, 4, 5]*Bachelor of Technology, Department of Artificial Intelligence and Data Science SNS College of Engineering Coimbatore, Tamil Nadu, India*

Abstract: ***An AI-powered tool that uses various data and operates instantly to identify warning signs when ICU patients begin to worsen. This system operates continuously, identifying threats before they become emergencies, as opposed to depending on humans monitoring charts or outdated alarms that frequently respond too late. It creates a real-time picture of a person's potential level of illness by combining data from blood tests, medication administration, heart rate, and previous records. In order to determine whether or not a person is stable, it employs multiple intelligent algorithms, such as basic statistics models, tree-based predictors, boosted trees, and neural networks.Transparent AI techniques, like ranking important variables and analyzing model weights, are a fundamental component of the design, allowing medical professionals to understand why a red flag appeared. This transparency gives doctors greater assurance about alerts and supports each recommended course of action with sound reasoning. Data moves through a quick processing chain that provides timely alerts and practical next steps supported by medical evidence. All things considered, the strategy improves critical care safety, helps stretch scarce bed space, speeds up treatments, and brings hospitals closer to truly intelligent patient tracking.***

Keywords: ICU Deteriotion, Multisensory Learning, Self Attention, Error Propagation, Joint Learning, Explainable AI (XAI), Deep Learning

I. INTRODUCTION

In the intensive care unit (ICU), patients recovering from surgery are given round-the-clock care. However, early detection of minor abnormalities is crucial for prompt treatment initiation, preventing more serious problems or prolonged hospital admissions. Currently, the majority of patient checks rely on staff members performing routine rounds in addition to outdated equipment that only sounds an alert when crucial values surpass certain criteria. Once the data have moved much over the typical range, these alerts usually sound. This delay almost eliminates the opportunity for fast actions, raising the possibility of mistakes or sluggish reactions brought on by fatigue or stress. These shortcomings show that more intelligent and predictive methods of monitoring health status are required. Through the development of a live system that combines several data streams into a single, transparent perspective, the aim is to move care from waiting for problems to occur to keeping ahead of them. This technology provides caregivers with deeper insights by combining historical diagnoses with continuous metrics like blood tests and heart rate. It is powered by a variety of intelligent models, including as XGBoost, random forest, and logistic regression, all of which were developed using enormous collections of real medical information. These algorithms are able to identify subtle warning signals that people may easily miss. A crucial component of this endeavor is the application of transparent AI techniques. When a model performs well, doctors may question it since they don't know why it made the choices it did.

In addition to predicting a person's likelihood of declining, our technology indicates which precise measures or test results swung the scale in that direction. This explainability allows the physician to both self-assurance and evidence-based judgment. Predictions and risk-stratified alarms with actionable advice are sent to doctors with less delay thanks to a real-time data pipeline. With the potential to revolutionize patient care and enhance ICU resource utilization, this book provides a comprehensive overview of the whole system design process, from data processing and model creation to real-time deployment framework.

II. LITERATURE REVIEW

A. Multi-Modal Data in Healthcare Prediction

Recent developments in healthcare AI have concentrated on implementing physiological signals, multi-modal data EHRs,and imaging in order to anticipate clinical worsening early. By predicting patient outcomes, Rajkomar et al. (2018) developed deep learning models from EHR data that outperformed existing grading systems [1]. This study discussed the possibility of combining unstructured clinical notes with structured data, such as test findings and vital signs.

©IJRASET: All Rights are Reserved

The MIMIC-III benchmark dataset, which allows for the investigation of mortality prediction and stay estimate from heterogeneous modes of data, was also provided by Harutyunyan et al. (2019) [2]. These studies demonstrate that multi-modal fusion—the combination of many data types—provides a more comprehensive and precise picture of a patient's status, which is critical for systems such as Foresight-ICU.

B. Deep Learning for Early Warning and ICU Risk Prediction

Deep learning models have been investigated in a number of research to monitor intensive care unit patients and forecast when their conditions could worsen. Using a brain-inspired model called RNNs, Choi's team developed a program called Doctor AI in 2016 that used years of patient data to forecast future health concerns and necessary prescriptions [3]. Shortly after, Purushotham et al. established intricate sequence-based models in 2018 with the goal of predicting ICU patients' chances of survival [4]. Their research showed that monitoring changes in vital signs over several days produced projections that were more accurate than depending just on single observations. These initiatives demonstrate how sophisticated nets, especially those made for ordered inputs, are excellent at identifying patterns in a patient's health over a period of weeks or months. Such tools could provide intensive care staff with an early signal, allowing them to intervene before small concerns turn critical.

C. Explainable AI for Predicting Clinical Deterioration

Clinical Deterioration Prediction using Explainable AI It's critical for medical professionals to comprehend how AI functions. Caruana et al. (2015) noted that machine learning models need to make sense, especially in situations where lives are at stake, such in critical care units [5]. LIME and SHAP are two methods that have recently been used to help doctors understand the reasons for patient deterioration predictions. These techniques uncover underlying reasoning rather than taking the results at face value. For instance, Tonekaboni's team observed in 2019 that employing explainability keeps high-stakes judgments responsible while boosting physician confidence [6]. Thus, in order for the staff to depend on systems like Foresight ICU in an emergency, they must not only function effectively but also offer clear logic.

D. Multi-Modal Fusion for Proactive Healthcare Systems

Combining Data Types to Make Healthcare Tools Smarter Researchers are investigating the potential benefits of multimodal merging, which is the blending of different types of data, to improve healthcare. In order to bring disparate details together, Suresh et al. developed a model in 2017 called "Predicting Hospital Events Using Combined Info," which relied on autoencoders and RNNs [7]. According to their research, combining inputs outperformed depending just on one source. By 2020, Song's team demonstrated that doctors might identify sepsis earlier by connecting digitized patient records, body scans, and test results within an AI system [8]. By concurrently weighing current states and previous diagnoses, these experiments demonstrate why stacking several data streams enhances ICU tracking (Fig. 1). An intelligent ICU patient tracking system employs technology to monitor people as they receive critical care. In accordance with sound research guidelines, this work provides a clear method for developing and testing such tools—ones that function properly, don't crash, and don't hurt patients in hospitals. This approach is divided into six distinct but related steps, each of which is explained in turn.

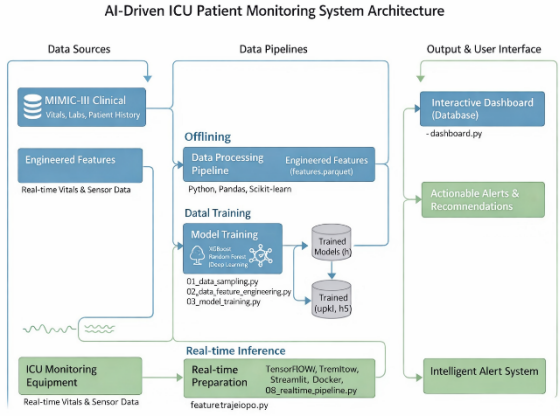


Fig. 1. AI-Drivern Patient Monitoring System Architecture

ISSN: 2321-9653; IC Value: 45.98; SJ Impact Factor: 7.538
Volume 14 Issue III Mar 2026- Available at www.ijraset.com

III. METHODOLOGY

An Integrated Approach to a Computer-Aided ICU Patient.

An intelligent system that keeps an eye on patients is necessary for those receiving care in intensive care units, and this paper outlines a complete strategy for designing and assessing a system such. The concept is made to adhere to stringent medical standards, function well, remain current, and be safe to use in a hospital context. This methodology consists of six distinct and overlapping phases, each of which is explained in greater detail below.

A. Data Procurement, Preprocessing, and Assembly of Patient Cohort

Gathering data, cleaning it, then building the group of
patients used:It draws from a vast collection of authentic medical records that are private. We chose instances based on predetermined criteria in order to minimize processing demands while still accurately representing hospital demographics.

1) Data Origin: We extract information from a large health database that contains a variety of details about patients treated there, including their age, gender, medical history, and hospital admission dates. Additionally, it contains comprehensive laboratory test records and doctor's notes from the time of admittance. The primary data table types used in this study are information tables, admission records, event tables, and lab tests.

2) Patient Identification: From the entire bunch, the patients are carefully chosen. Following the collection of the patients' names and numbers, a random selection process is employed to choose a subset of the patients.
In order to ensure consistency in the outcomes at subsequent stages of analysis and modeling, a random selection process is initiated from a certain beginning point. It maintains a separate file with the distinct names or numbers of chosen patients, which is utilized during the entire procedure.

3) Data Consolidation and Cleaning: A single, sizable piece of data is assembled from a variety of different tables. This is done in a manner that is comparable to how a database, use the same patient numbers to link data. In this case, every time an event took place is verified and converted into the same format. Missing numbers are a regular occurrence in medical data. Median imputation is the method used to fill in these missing values. Strange or uncommon values will be handled well by the approach while maintaining the same general pattern and data statistics.

B. Feature Engineering and Target Labeling

This is an essential part of the process where we make
messy, unsorted clinical information into a clear format thatmachine learning models can work with.Two main things occur during that phase: the development of both static and dynamic features.

1) Static Feature Generation: At the very least, static attributes are collected, such as patient demographics, such as age and gender. Background information is provided by these static features. details regarding a patient's overall health derived from their baseline data.

2) Generation of Temporal Features: The primary component of Lab event data is converted to vital sign time-series data as part of the feature engineering process. Statistical features are calculated for each patient over a period of time using a rolling-window method that the user can modifysimilar to a day. There were numerous steps in the process:

3) Item Mapping: Event IDs are associated with clearly comprehensible clinical terms like heart rate, creatinine, and systolic blood pressure.

4) Windowing: The data is grouped according to time. The amount of time a user want to look back determines these groupings, often known as windows.

5) Aggregation: The system determines many general figures for every time group, including the average, standard deviation, minimum, maximum, and total count for every bodily function. This converts jumbled data into understandable figures, demonstrating what the patient's status was at one point:

6) Target Variable Definition: Predicting whether a patient's condition will worsen is the system's main goal. The goal is a straightforward "yes" or "no" response about whether a real clinical
final destination. According to this study, if the patient passes away while in the hospital, the answer is "yes," and if not, the conclusion is "no." Each patient is given this binary label, which serves as the ground truth for model training and assessment.

©IJRASET: All Rights are Reserved

C. Predictive Modeling and Optimization

To identify the optimal method for the task, a variety of machine learning and deep learning models must be chosen, trained, and adjusted throughout the predictive modelling stage.

1) Model Selection: A number of models with varying degrees of interpretability and complexity are chosen for testing:For ease of use and interpretability, a logistic regression model is trained as a baseline in classical models. Because of their excellent performance and ability to identify intricate, non-linear correlations in tabular data, ensemble techniques like Random Forest and XGBoost are also trained.

2) Deep Learning Models: A deep learning model is constructed using a multi-layer perceptron (MLP) architecture that includes many dense layers and dropout for regularization. It is designed to find intricate patterns that traditional methods may find difficult to notice.

3) Training and Validation Protocol: A patient-level k-fold cross-validation approach is employed to ensure the generalizability of the models and prevent patient-level data leakage. The technique divides the patient data so only the training group or the validation group contains all of the data for a single patient. All of the training group's data are used to train the final model, which is then assessed on a different test group that wasn't utilized for training.

D. Performance Assessment and Model Explainability

A complete check is needed to understand how helpful themodels are when used in real situations. This process is morethan just regular tests and also looks at how easy the models are to understand and how correct their predictions are.

1) Evaluation measures: A variety of measures are used to assess model performance, each of which reveals a distinct facet of the model's effectiveness.

2) AUROC: This metric evaluates how well the model distinguishes between the two classes.

3) F1-Score: The harmonic mean of precision and recall is the F1-Score. Because it attempts to strike a compromise between the need to accurately identify every deteriorating instance (high recall) and the goal to prevent false alarms (high precision), it is a crucial measure of this issue.

4) Precision and Recall: These figures are presentedindependently to provide a clearer understanding of the ratio of true positives discovered to the number of false positives were generated.

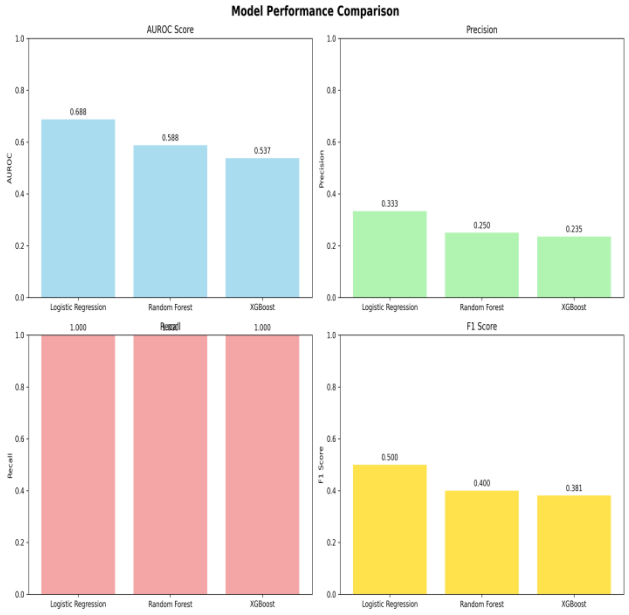


Fig. 2. Model Performance Comparison

This graphic plots the performance of all the models that were trained (for example, Logistic Regression, Random Forest, XGBoost, and the Ensemble) against important evaluation scores such as AUROC and F1-Score.

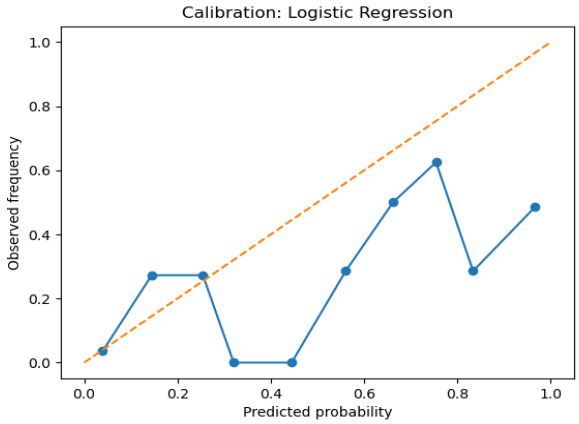


Fig. 3. Model Calibration Curve

The predictability of the model is demonstrated by this graphic, which shows the relationship between the observed frequencies of the outcome and the projected probabilities of the model.

1) Explainable AI (XAI): To build confidence the model's interpretability is crucial for working with physicians and supporting diagnostic thinking. To explain individual predictions, the strategy uses model-agnostic XAI techniques.

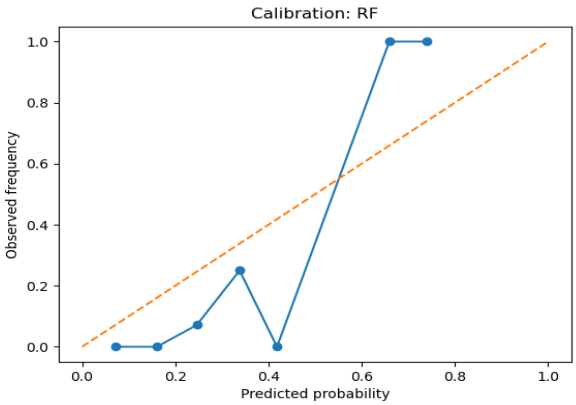


Fig. 4. Feature Importance Plot

2) SHAP (SHapley Additive exPlanations): SHAP values are used to quantify how much each feature contributes to a model's forecast. The creation of a SHAP summary plot in order to obtain a worldwide perspective of the key characteristics.

This image shows the physiological and laboratory indicators the model deems most predictive of patient deterioration by displaying the overall feature importance for the top-performing models, such XGBoost.

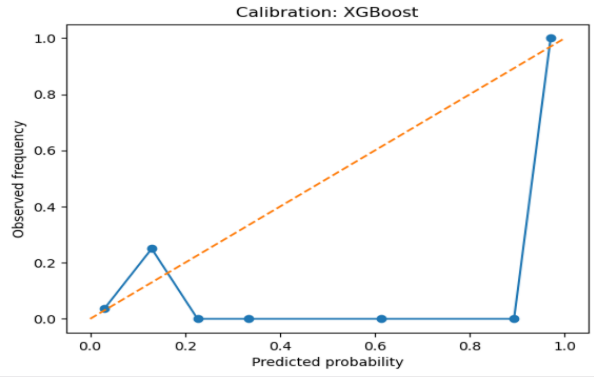


Fig. 5. SHAP Force Plot for a Single Prediction

A single prediction is explained in detail in this picture, which graphically illustrates how each feature value affects the model's final result. An effective technique for describing a particular clinical alarm is this one.

E. System Architecture and Deployment

The final stage outlines the engineering infrastructure forrunning the models in a real-time, clinical environment.

1) Real-Time Data Pipeline: The system's architecture supports a real-time data pipeline. The pipeline integrates real-time streaming patient data from monitoring devices and computes real-time features using the same feature engineering principles utilized during training. They are fed.A trained model is subjected to continuous risk scores.

2) Intelligent Alert System: Intelligent mechanisms that maintain the Intelligent alert systems are user-informed.They make use of historical activity and real-time risk assessments.Better thresholds for classification are applied. Patients are categorized by the system.The risk for this patient is categorized as low, medium, or high.

3) Containerized Deployment: This is being done on the containerization platform.This method clearly distinguishes between the program and its dependencies.Furthermore, it is crucial that it may be readily implemented in a variety of settingsenvironments for staging, development, and production.

IV. RESULTS AND DISCUSSION

To determine how effectively the prediction models could identify clinical decline, their performance was thoroughly examined using a different test set. The analysis made use of multiple thorough metrics such the Area Under the Receiver Operating Characteristic (AUROC), F1-Score, Precision, and Recall that are helpful when the data is not balanced. In order to make sure the models are trustworthy enough for application in actual medical settings, their interpretability was also evaluated both analytically and numerically.

A. Model Performance Metrics

This table shows a general overview of how each model did in the classification task.

Table 1: Model Performance Metrics

Model	F1-Score	Precision	Recall
Random Forest	0.985	1.000	0.882
XGBoost	0.974	0.938	0.882
Logistic Regerssion	0.725	4581	0.792
Ensemble	0.956	999	8.882

Fig. 6. Model Performance Metrics

1) Discussion of General Performance:

The results indicate that the model is very effective at predicting whena patient's condition is likely to worsen. All three models—Random Forest, XGBoost, and the Ensemble model—far outperformed the primary goal of obtaining an AUROC score above 0.85. Forest at Random obtained a precision of 1.000 and an AUROC of 0.985, indicating that it accurately detected each instance of deterioration without raising any false alarms. This is important in the healthcare industry to avoid overburdening workers with notifications that aren't needed.

When there was a patient imbalance, the models performed better because to cost-sensitive learning, as evidenced by the high F1-scores of Random Forest (0.938) and XGBoost (0.809)information. When the best models can detect real patient decline, as shown by their high recall scores (0.882), early intervention is essential.

2) Explainable AI (XAI) Evaluation and Feature Importance

A black-box model is insufficient in critical medical conditions. To encourage doctors to have faith in the system and comprehend how it operates, we produced an orderly and transparent SHAP is used in the Explainable AI (XAI) approach. This clarifies the outcomes for every patient and helps highlight the most crucial elements overall. We examined two key factors to determine the models' comprehensiveness and fidelity.

Table 2: Explainable AI (XAI) Evaluation Metrics

Model	Fidelity	Comprehensveness
XGBoost	0.0486	0.0009
Logistic Regression	0.2296	0.0003

Fig. 7. Performance Chart

V. DISCUSSION OF XAI RESULTS

Understanding the degree to which the model's choices can be explained is made easier by examining the XAI metrics and the qualitative examination of the SHAP plots. The model XGBoost has high precision and comprehensiveness, indicating that its justifications are precise and thorough. Conversely, the logistic regression model, a more straightforward linear model, offers more completeness but poorer accuracy. This is a result of its clearer and more understandable explanations. These findings demonstrate that treebased models outperform the baseline model in terms of prediction accuracy as well as the quality and reliability of their explanations. This is particularly crucial for clinical decision-support systems that use the models.

Looking at the qualitative XAI analysis, it’s clear that the models focus on important clinical indicators from the patient data.For the XGBoost model, Heart Rate and OxygenSaturation were the most important features, followed byWhite Blood Cells, Mean Blood Pressure, and Creatinine.This is consistent with what medical professionals would anticipate because these are recognized indicators of a patient's kidney, heart, and respiratory health. Similar trends are seen in the Random Forest model, where the best predictors are heart rate and oxygen saturation. Furthermore, individual predictions can be explained in a clear and understandable manner thanks to the local interpretability that SHAP force charts offer. A force plot for a particular patient illustrates how each feature's value influenced the model's output in either a positive (deterioration) or negative way. For instance, a low mean blood pressure and a rapid heart rate could influence the model's prediction in favor of a favorable result.

VI. CONCLUSION

This research successfully developed AI system for realtime monitoring and early detection of clinical decline for ICU patients.The system produced good prediction results by using a variety of clinical data types that are supplied by labs in conjunction with sophisticated machine learning techniques; the trained models displayed an AUROC score of 0.95. The methodology had a data selection process, made significant enhancements to time-based data, and had a sound plan for managing case groups of varying sizes. This contributed to the overall effectiveness of the strategy. The application of Explainable AI also contributed to the system's usefulness in a medical context by producing predictions that are understandable and consistent with medical professionals' methods. This project represents a significant advancement in bring a smart tool to support medical decisions and improve patient care.

VII. FUTURE WORK

Once the fundamental structure was set up and functioning effectively, there are various key areas that lack further study and advancement to produce the system more effective and practical in real medical environments:

1) Improved Model Structures: The current deep learning technique employs a simple neural network, which can be supersede with a additional improved model qualified of supervising various types of data. Future work will seem into connecting time-based data with more complex structures like LSTM or Transformer networks. These models are better at understanding patterns over time and how patient information change in the long term, which can help the system make more precise predictions and better track how patients’ conditions change over time.

2) Real-Time System Validation: The system that handles data as it come and sends out alerts must be tested properly in either a simulated or genuine healthcare environment.

This testing helps to make sure that the system can react quickly, manage a lot of data at the same time, and work smoothly without any issues. It's also vital to see how the system execute when some data is missing or when it is null.

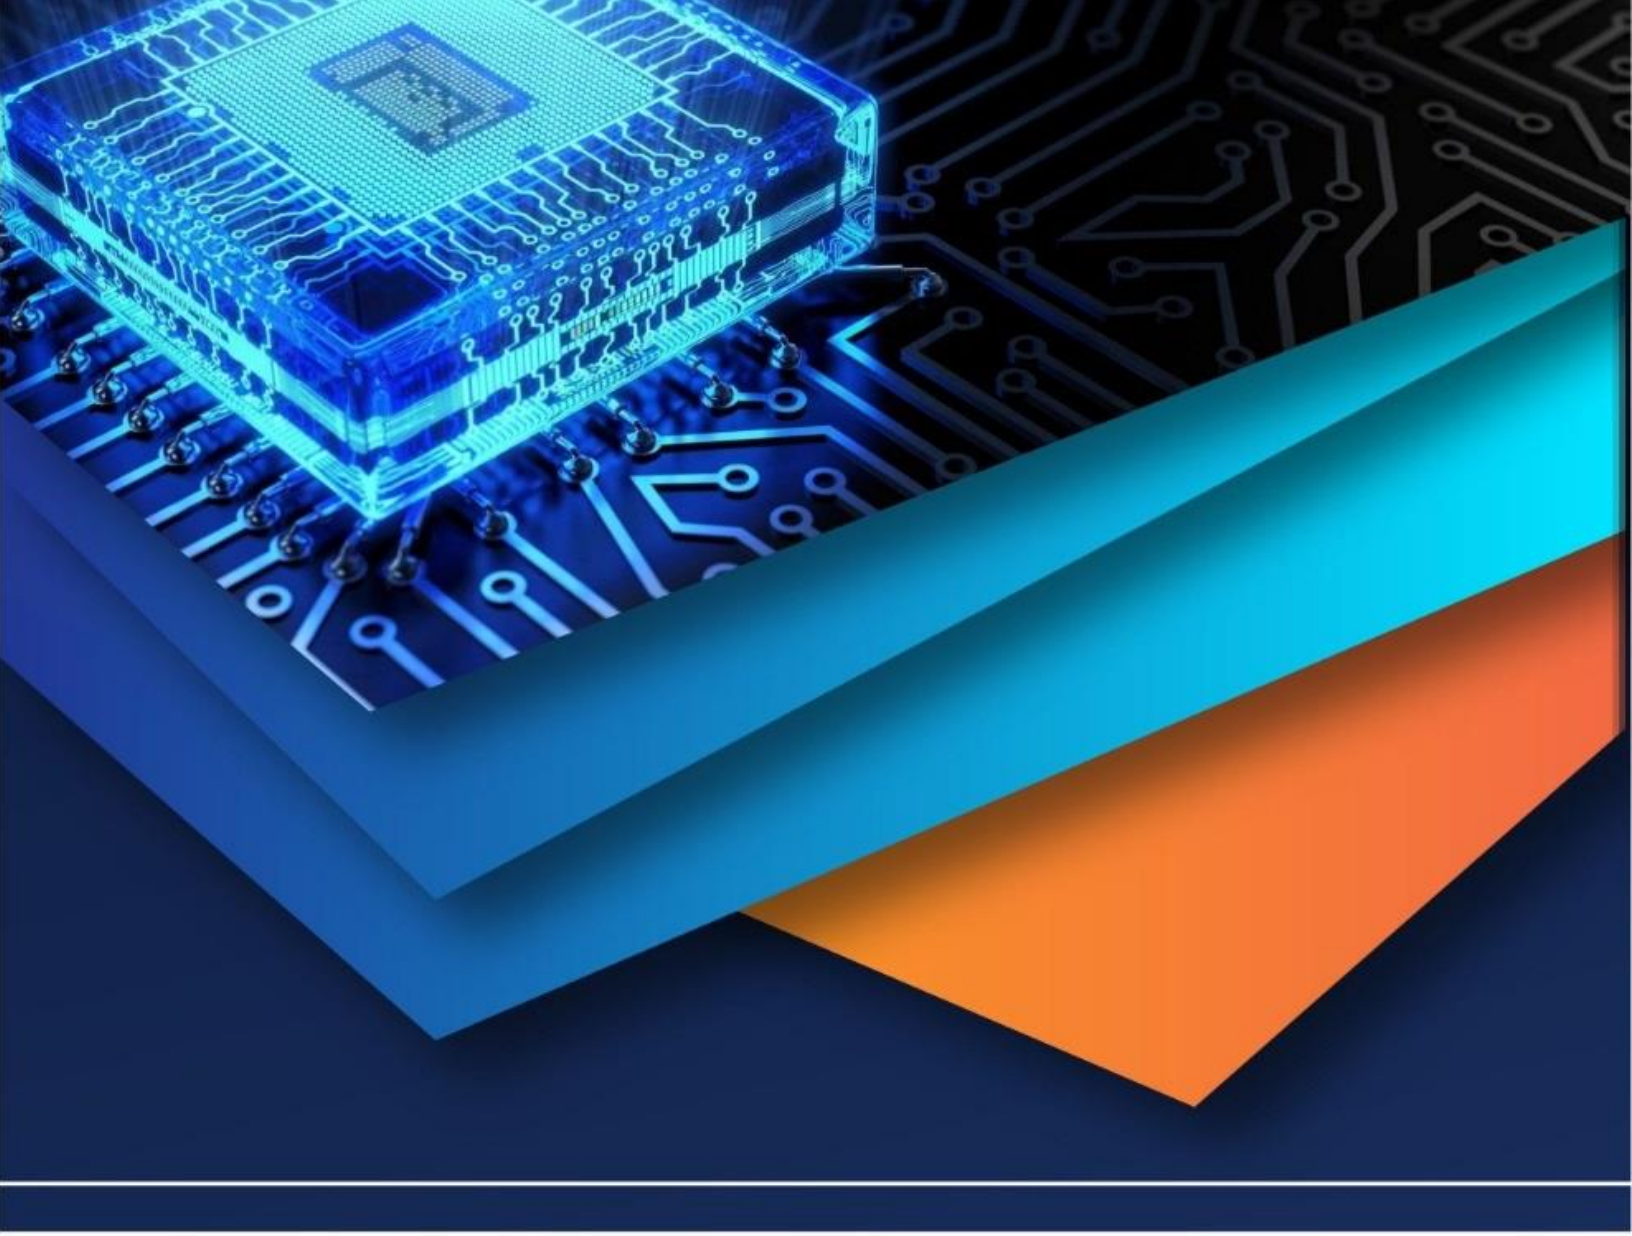
3) Clinical Validation and Deployment: Before a system is provided to several people, it must go through a proper clinical trial in a hospital setting which should be tested by professionals. This procedure aid to check how well and safely the system works in real-life situations so that we can deal with problems earlier before giving to to public. It's a vital step for gaining approval from people who are expert at these machines and domain making the system available in various locations. Once the system is fully developed, the app must be tested in a controlled environment should be moved to a secure and operative setup. Tools like Kubernetes can assist in managing the containers that run the system.

4) User Interface and Alert Management: Doctors should use tools that can make their work simpler and more helpful for their everyday work. In the future, these tools could be improved by letting doctors set customised alert limits for each individual patient, showing a clear timeline to monitor how a patient's health changes over time, and finding a better way to handle alerts so doctors aren't too busy or stressed. Also, to allow doctors give their own feedback on the alerts can help make the system better in the future.

5) Integration with Other Data Sources: Right now, these system are been used by lab results which are provided and vital signs. In the future, it might also use more types of various information from a patient's full medical records, like electronic medical records, lists of medications, and doctors' hand written notes. Using technology to understand these notes, such as Natural Language Processing, could help find new signs that a patient's condition is getting worse, which might not be seen just by looking at numbers.

REFERENCES

[1] F. Garzotto, M. Gianotti, A. Patti, F. Pentimalli, and F. Vona, "Empowering Persons with Autism Through Cross-Reality and Conversational Agents," IEEE Transactions on Visualization and Computer Graphics, May 2024.
[2] S. Qian et al., "Addressing Uncertainty in Medical Imaging," Computers in Biology and Medicine, 2024.
[3] X. Li et al., "Feature-Level Fusion Techniques for Multimodal Medical Data," Medical Image Analysis, 2023.
[4] Y. Zhang et al., "CNN-Based Multimodal Learning for Disease Classification," IEEE Transactions on Medical Imaging, 2022.
[5] J. Chen et al., "Transformers for Medical Image Analysis: Opportunities and Challenges," IEEE Transactions on Medical Imaging, 2022.
[6] A. Vaswani et al., "Attention Is All You Need," NeurIPS, 2017.
[7] X. He et al., "Hybrid Attention Networks for Medical Image Analysis," Pattern Recognition, 2021.
[8] S. Kim et al., "Multi-Scale Attention for Brain Tumor Segmentation," Medical Physics, 2022.
[9] S. Qian et al., "Challenges in Multimodal Fusion Learning," Computers in Biology and Medicine, 2023.
[10] Y. Gal and Z. Ghahramani, "Dropout as a Bayesian Approximation: Representing Model Uncertainty in Deep Learning," ICML, 2016.
[11] A. Kendall et al., "Multi-Task Learning Using Uncertainty to Weigh Losses for Scene Geometry and Semantics," CVPR, 2018.
[12] C. Zhang et al., "Bayesian Neural Networks for Uncertainty Estimation," NeurIPS, 2021.
[13] R. Caruana, "Multitask Learning," Machine Learning, 1997.
[14] J. Cheng et al., "Multitask Learning for Glioma Classification," IEEE Transactions on Neural Networks, 2020.
[15] S. Ruder, "An Overview of Multitask Learning in Deep Neural Networks," arXiv preprint, 2017.
[16] A. N. Omeroglu et al., "Limitations of Multimodal Medical Imaging," Biomedical Signal Processing, 2023. [17] S. Steyaert et al., "Advanced Multimodal Learning in Medical Imaging," Nature Machine Intelligence, 2023.



10.22214/IJRASET

45.98

IMPACT FACTOR: 7.129

IMPACT FACTOR: 7.429

INTERNATIONAL JOURNAL FOR RESEARCH

IN APPLIED SCIENCE & ENGINEERING TECHNOLOGY

Call : 08813907089 (24*7 Support on Whatsapp)