RESEARCH PAPER

Real-Time Explainable Multimodal ML for Clinical Decision Intelligence A Hybrid Supervised–Unsupervised CDSS Framework

[1] Kalyana Krishna Kondapalli, India,
kalyanakondapalli@gmail.com

[2] Chaitanya Gunupudi, India,
chaitanya.gunupudi970@gmail.com

Abstract- The rapid growth of diverse clinical data, such as electronic health records (EHRs), lab results, and medical imaging, has opened new possibilities for data-driven decision-making in healthcare. However, issues with data quality, model interpretability, and workflow integration have hindered the safe and effective use of machine learning (ML) in clinical settings. This paper presents a clear multimodal ML framework for real-time clinical decision-making. It combines structured EHR data, lab tests, and imaging-derived features to aid in diagnosis, risk prediction, and personalized treatment planning. The framework integrates supervised learning models, including Random Forest, Gradient Boosting Machines, Support Vector Machines, and Logistic Regression, along with unsupervised techniques like k-means clustering and Principal Component Analysis. This setup allows for both predictive modelling and patient grouping. We use strong preprocessing, feature selection, and cross-validation to manage missing data, reduce overfitting, and improve generalization. Model performance is measured using metrics such as accuracy, precision, recall, F1-score, and AUC-ROC, with Random Forest and Gradient Boosting showing the best results overall. To meet the important need for transparency, we include SHAP (Shapley Additive explanations) to provide explanations of model outputs, both globally and at the patient level. This helps clinicians understand the main factors behind predictions and builds trust in the system's recommendations. The proposed framework is made for easy integration into current clinical workflows and allows for near real-time updates as new patient data come in. Overall, this work shows that a clear, multimodal ML framework can improve clinical decision-making and sets the stage for future evaluations and broader use in healthcare.

Keywords-Machine Learning, Clinical Decision Support, Supervised Learning, Unsupervised Learning, Healthcare, Predictive Analytics, Data Privacy, AI Integration

How to cite this article: Kondapalli KK, Gunupudi C. Real-Time Explainable Multimodal ML for Clinical Decision Intelligence A Hybrid Supervised–Unsupervised CDSS Framework. Int J Drug Deliv Technol. 2026;16(16s): 803-813. DOI: 10.25258/ijddt.16.16s.87.

I. INTRODUCTION

The digitalization of healthcare has led to an amount of different clinical data, including electronic health records, lab test results, medical images and information from patients. This data has the potential to improve diagnosis predict risks and plan treatments. Its too much for traditional methods and manual review. Machine learning and artificial intelligence can help find patterns in this data and provide decision support at the bedside.

However, integrating machine learning into practice is challenging. Clinical data is often incomplete inconsistent and varies across institutions. Many machine learning models are like " boxes," making it hard for clinicians to understand their predictions. Concerns about transparency, fairness, data privacy and workflow disruption have limited the use of these models in decision systems.

Clinical Decision Support Systems are changing from rule-based systems to platforms using machine learning models. Traditional systems rely on created rules and often focus on a single data source. They can be inflexible and hard to maintain. There is a need for systems that can learn from data provide personalized recommendations and are transparent.

This paper proposes a machine learning framework for time clinical decision support. The framework combines health records, lab measurements and imaging data to support predictive tasks, such as disease risk estimation and outcome prediction. It uses a mix of unsupervised learning techniques to capture both predictive and structural properties of clinical data. The framework also incorporates SHAP to provide transparency and clinical acceptability.

By combining data sources machine learning techniques and a focus on transparency this study aims to advance the development of machine learning-enabled clinical decision support systems that are accurate, trustworthy and suitable for real-world use in healthcare settings.

The goal is to create systems that're not only accurate but also provide actionable insights and support clinicians in their daily work. This requires an understanding of clinical data and the needs of healthcare professionals. By addressing these challenges, we can unlock the potential of machine learning in healthcare and improve patient outcomes.

Machine learning has the potential to revolutionize healthcare by providing clinicians, with accurate decision support. However, it requires consideration of data quality, model transparency and clinical workflow. By developing frameworks that prioritize these factors we can create machine learning-enabled decision support systems that are both effective and trustworthy.

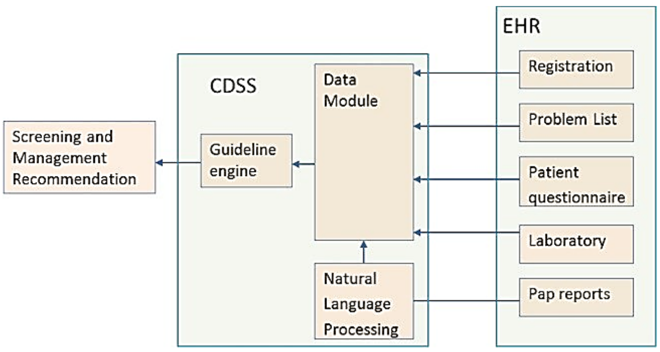

Figure: 1 Clinical decision support system structure

Novelty of the Research

The originality of this study is the combination of the modern machine learning methods with clinical decision support in the manner that critical problems of conventional CDSS are addressed. Key novel aspects include:

Integration of Multi-source Clinical Data: This study is different in that, as opposed to the traditional systems, which commonly emphasize one type of clinical data (e.g., EHRs or medical images), the proposed one incorporates multiple clinical and clinical data sources, including EHRs, lab results, medical imaging, and patient demographics. Processing and processing of various data is an important innovation in offering holistic support decision to clinicians.

Integration of Supervised and Unsupervised Learning Disciplines: This study hypothesizes a mix of the supervised and unsupervised learning methods. Although supervised learning is used to make predictions using historical patient data, unsupervised learning, including clustering and anomaly detection, are also added to reveal the non-obvious patterns in patient populations, as well as discover new types of a disease or uncover previously unknown cases of illness.

Model Explainability and Interpretability: This is a new dimension of the current study because it is designed to advance the explainability of machine learning models. Lots of deep learning models can be considered as black boxes because of their inability to be interpreted. This study examines approaches like SHAP (Shapley Additive explanations) values to make sure that clinicians can comprehend what motions or treatment suggestions were made so that they can enhance confidence and embracement in clinical practice.

Individualized Decision-Making Model: The proposed system will provide individualized advice following a thorough investigation of individual medical history of a patient. This is opposed to the generalized nature of the approaches used by most existing CDSS systems, which offer personalized treatment plans and recommendations that respond to the needs of a given patient.

Real-Time Clinical Decision Support: Another innovative feature is the possibility to give a real-time clinical decision support, which is conducted due to the constant monitoring of patient data. This attribute makes the system more efficient as it will provide clinicians with more help in the point of care, leading to a decrease in the chances of diagnostic errors.

Combining Workflow with Clinical Practice: This study is deeper than the algorithmic part of the topic, as it also addresses the topic of integrating the CDSS into the normal clinical workflow. The real problem of implementing such systems into the actual healthcare setting, without interfering with the daily routine, has been frequently neglected. The given paper suggests a framework, which will guarantee a seamless implementation in the clinical setting.

Background and Motivation

Access to extensive and complete information about patients' clinical condition and prognosis is crucial for the proper delivery of healthcare services. Nevertheless, healthcare professionals have to integrate large volumes of information and make decisions within limited time frames. This is where ML-based decision intelligence systems can help healthcare professionals integrate information and identify high-risk patients or possible diagnoses and personalized therapeutic approaches. At the same time, there is growing concern and regulation that systems must be transparent and follow clinical guidelines and recommendations and avoid being black boxes.

Problem Statement

Typically, the existing solutions for ML-based CDSS systems have one or more of the following limitations: the use of a single data source, limited robustness to incomplete and noisy data, absence of a real-time prediction capability, and poor interpretability of the predictions. For example, most of the models are trained and tested on static data sets without considering the need for updates. Furthermore, the predictions provided to the clinicians are often limited to a risk score or a classification outcome, without clear explanations of which features were used to make the prediction and how reliable the prediction is for a given patient. This is a problem because the final responsibility for the decisions rests with the clinician and requires a clear understanding of the underlying rationale.

This work will focus on the problem of designing a CDSS framework that can (i) integrate heterogeneous data sources, (ii) make accurate predictions and provide clinically relevant patient stratifications, (iii) make predictions in near real time, and (iv) provide explanations that can be interpreted and critiqued.

Objectives

The overall goal of the present study is the development and assessment of an explainable multimodal ML framework, specifically for real-time clinical decision intelligence. The objectives of the study are as follows:

Design a powerful data processing pipeline, which includes the integration of structured electronic health records, lab data, as well as image-based features.

Develop and compare the performance of multiple supervised learning models, including Random Forest, Gradient Boosting Machines, Support Vector Machines, as well as Logistic Regression, for clinical prediction tasks.

Explore the implementation of unsupervised learning models, including k-means clustering as well as PCA, which can help identify novel clinical patient groups.

Design the integration of explainable AI techniques, including SHAP, which can provide global as well as patient-level explanations of the model outputs, specifically for clinicians.

Present the potential of the proposed framework, specifically with respect to the integration of the model into the existing clinical workflow, thus supporting real-time risk assessment as well as recommendation.

Contributions and Novelty

The specific contributions of the present work are as follows:

Multimodal Integration: We propose a system that integrates multiple clinical data sources, including EHR, lab, and image-based features, all through a unified pipeline, moving beyond the largely unimodal focus of most existing CDSS.

Hybrid Supervised-Unsupervised Learning: The system unifies supervised learning-based prediction models with unsupervised learning-based clustering and dimensionality reduction, thus supporting both prediction and data-driven patient stratification.

Explainable Decision Intelligence: We incorporate SHAP-based explanations, thereby providing transparent, clinically meaningful insights into the relative contribution of individual features towards the prediction models at both the cohort and individual patient levels.

Real-Time and Workflow-Oriented Design: The proposed system has been designed with the goal of supporting real-time operation as well as Workflow-oriented design, thereby facilitating its practical deployment.

Empirical Evaluation of the Performance of the Proposed System: We have conducted a comparative analysis of the performance of multiple ML models using metrics such as accuracy, precision, recall, F1-score, and AUC-ROC, thereby validating the effectiveness of the ensemble approach of the proposed system.

II. LITERATURE REVIEW

Traditional Clinical Decision Support systems

Clinical Decision Support Systems (CDSS) have traditionally been developed as rule-based or knowledge-based systems, which incorporate knowledge in the form of clinical guidelines, expert rules, or threshold-based rules. These systems generally utilize manually constructed logic, such as if-then rules based on clinical practice guidelines, and operate on structured EHR data, including diagnoses, medications, and labs. As the current literature emphasizes, some of the ways of applying ML are disease diagnosis, treatment suggestions, and risk prediction of patients [1][2]. Although rule-based CDSS have shown promise in domains such as drug-drug interaction checking and adherence to guideline-recommended care, these systems are limited in their ability to accommodate new evidence, complex interactions between variables, or high-dimensional data, including images and time-series data. The maintenance and updating of large rule sets require significant human resources, and the inflexibility of rule-based logic can contribute to alert fatigue and decreased clinician trust.

Supervised Learning in Healthcare

The adoption of supervised machine learning has witnessed rapid growth, primarily due to the availability of labelled data sets of large sizes. The most used supervised machine learning models include logistic regression, decision trees, Random Forest, Gradient Boosting Machine, Support Vector Machine, and various types of neural networks. These models have achieved great success in the fields of disease detection, hospital readmission, mortality risk, and treatment response. In most cases, the supervised machine learning models have achieved better performance than the traditional models, as they can easily identify the non-linear relationships between the variables. However, the majority of the literature has used a single modality of data, which includes structured electronic health records or claims data. These models have failed to take advantage of the full richness of multimodal data. Moreover, the majority of the literature has used static models, without any specific emphasis on real-time deployment.

Unsupervised Learning and Data Clustering

Unsupervised learning has been employed to address patient stratification problems, subgroup discovery, and phenotyping. Clustering techniques such as k-means clustering, hierarchical clustering, and density-based clustering have been applied to group patients according to their clinical characteristics and comorbidities. The techniques can identify patient subgroups that may not be apparent according to traditional clinical diagnoses and may have different responses to different treatments. Dimensionality reduction techniques such as Principal Component Analysis (PCA), non-negative matrix factorization, and manifold learning have been employed to reduce the dimensionality of high-dimensional clinical data into lower-dimensional spaces that can be easily understood and manipulated. However, many unsupervised learning-based applications in healthcare rely on single data sources and have been mainly used as exploratory tools. The applications have limited potential to support CDSS. The integration of unsupervised patient stratification and supervised predictions within a unified framework that can support CDSS is still in its infancy.

Deep and Multimodal Learning for CDSS

Deep learning methods have also allowed for significant advances in handling complex data types, especially in the field of medical imaging and physiological time series. Convolutional neural networks (CNNs) have been successful in achieving expert-level performance in tasks such as lesion detection in radiographs, classification of CT and MRI scans, segmentation of anatomical structures and the human clinician cannot always see [12]. Instead, RNNs come in handy when working with sequential data, including electrocardiograms (ECGs) to forecast patient outcomes [13][14]. Recurrent and transformer-based models have also been used for sequential data types, including electrocardiogram signals, vital sign data, and longitudinal EHR data for dynamic risk prediction and early warning systems. More recent developments in multimodal models, incorporating structured EHR data, free-text clinical notes, imaging, and laboratory results, have also shown promise in handling multiple data types. Although deep and multimodal models have shown promise in achieving high predictive performance, they also bring in new challenges in terms of computational requirements and interpretability. Hence, their incorporation into CDSS is limited, especially in resource-constrained situations and in situations where regulatory requirements for explainability are stringent.

Identified Gaps in Existing Research and Novelty in Studies.

Although the literature on the application of ML in healthcare is quite large, a number of gaps in research are crucial to be addressed. The incorporation of various data points, including EHRs as well as medical imaging, genomic information, and patient demographics, into a single decision support system is one notable gap. The current state of CDSS puts emphasis on one type of data or two which does not include the comprehensive character of the decision-making process. The other gap is that there are no effective approaches to model interpretability and explainability. Although deep learning models are very powerful predictors, their opaqueness usually does not allow healthcare professionals to trust and interpret their decisions.

This paper closes these gaps by presenting a machine learning-based CDSS which incorporates a variety of data sources and comes up with both supervised and unsupervised learning approaches. The novelty of the study is that it is dedicated to explainable AI (XAI), which increases the transparency of ML models with the help of such techniques as SHAP (SHapley Additive exPlanations) values. Also, this project seeks to come up with a system that can support real time decisions through the constant processing of the patient data received. By sealing these gaps, this study is likely to guide the creation of better interpretable, effective and complete clinical decision support systems that can be easily integrated into clinical workflows.

Table 1: Comparison of Proposed CDSS with Existing Methods

Aspect	Existing Methods	Proposed CDSS
Data Integration	Typically uses a single data source (e.g., EHRs or medical imaging)	Integrates multiple data sources (EHRs, medical imaging, lab results, genomics, patient demographics) for a holistic approach to decision-making
Learning Techniques	Primarily supervised learning (e.g., Random Forest, SVM, Neural Networks)	Combines supervised and unsupervised learning, incorporating clustering and anomaly detection for more comprehensive insights
Interpretability	Often considered "black-box" models, especially in deep learning	Incorporates Explainable AI (XAI) techniques like SHAP values for model transparency and interpretability
Personalization of Treatment	Generally offers one-size-fits-all treatment recommendations	Provides individualized treatment suggestions based on a thorough analysis of each patient's unique medical history and conditions
Real-Time Decision Support	Decision support is typically not real-time, with batch processing	Provides real-time decision support by continuously processing and analyzing incoming patient data
Data Preprocessing	Limited focus on preprocessing, leading to potential data quality issues	Includes robust data preprocessing (e.g., missing data imputation, normalization) to ensure high-quality input for ML models
Clinical Workflow Integration	Often disrupts clinical workflows due to lack of user-friendly interfaces	Designed to seamlessly integrate into clinical workflows without causing information overload, enhancing efficiency
Model Robustness	Susceptible to overfitting or bias, especially with small or imbalanced datasets	Uses techniques like cross-validation and ensemble learning to enhance model robustness and reduce overfitting

Explainable AI for Clinical Decision Support

Explainable AI (XAI) has been identified as an important new area of research for clinical applications, where decisions must be transparent, accountable, and consistent with human reasoning. Post-hoc methods such as SHAP (Shapley Additive explanations), LIME (Local Interpretable Model-Agnostic Explanations), and feature attribution methods can be used to gain insights into how each feature contributes to a particular prediction. In healthcare, XAI has been employed to understand why a model is predicting high risk for a particular patient, to verify whether critical clinical risk factors are being correctly weighted, and to identify possible spurious relationships or biases. Global explanations, such as feature importance, can be used to understand how the model behaves in general, while local explanations, such as SHAP for each patient, can be employed for decision-making and documentation for each particular patient. However, there are still challenges in presenting XAI in an easily understandable form for clinicians, as well as in ensuring stability in XAI and in how to integrate XAI into clinical workflows with minimal cognitive burden.

Implementation and Workflow integration of ML-Based CDSS

In addition to the performance of the algorithms, the implementation of the ML-based CDSS must also focus on the implementation science. For example, studies have emphasized the fact that while the performance of the models may be satisfactory, the actual implementation of the models may not be effective due to poor usability, lack of integration with the EHR, poor engagement of the clinicians, and poor monitoring of the models. For example, the interoperability of the models, such as the use of standards like HL7 FHIR, the latency and reliability of the predictions, the integration of the alerts and recommendations with the clinical roles, and the feedback and override of the recommendations to the clinicians must be taken into consideration. Furthermore, it has also been recognized that the CDSS must also focus on the ethical and legal issues.

Challenges and Limitations

Although the use of machine learning in healthcare is very promising, it has a number of challenges. Data privacy and security are one of the greatest issues. Healthcare information is very personal and the confidentiality of information on patients is paramount. When managing and processing healthcare information, strict laws and regulations that include the HIPAA have to be followed. The other difficulty is the ethical implication of using machine learning. Such challenges as algorithmic bias, fairness, and transparency should be considered to make sure the models will not reproduce the current healthcare disparities. Another concern is objectivity of the ML models, especially deep learning. It is regarded that many models, and in particular, neural networks are black boxes, i.e., the process of decision-making is not clear. Such non-transparency may be a barrier to trust and acceptance by healthcare professionals who must be aware of how this information is trained to learn the machine learning models Lastly, the quality of the data used to train machine learning models is a major issue. Poor modeling performance and low effectiveness of decision support system may be caused by incomplete, inconsistent or even erroneous data. As such, good data preprocessing methods such as data cleaning and normalization are critical in ensuring machine learning models generate correct and satisfactory outputs.

Future Directions

Nonetheless, machine learning in healthcare does not stop its development, and in the future the studies will be probably devoted to the elimination of these obstacles. Explainable AI (XAI) where machine learning models can be made more transparent and understandable to clinicians is gaining interest. Such techniques as SHAP values and the Local Interpretable Model-Agnostic Explanations are under development to enhance the interpretability of complex models. Besides, the combination of multimodal data sources, including electronic health records, medical images, and genomic data, is promising to generate more comprehensive and personalized clinical decision support systems. Federated learning is another potential avenue that allows machine learning models to be trained on a set of decentralized sources of data, guaranteeing patient privacy and allowing large and diverse datasets to be used to train. Such developments will not only make machine learning more accurate and efficient in the health care sector but will also make the use of these systems more trusting and utilizable by the health care providers.

III. METHODOLOGY

In the following section, and describe the methodology that will be used in the construction of the machine learning-based Clinical Decision Support System (CDSS). The methodology includes the data collection, machine learning algorithms to use in predicting and classifying, feature selection and engineering, and model evaluation.

Data and Cohort Definition

The framework relies on the availability of clinical data from routine sources, including electronic health records, laboratory information systems, and repositories of medical imaging data. The study population includes patients with certain criteria, such as adult inpatients admitted during a specific period with at least one complete episode of care, among others.

Variables associated with each patient include the following categories of data:

1. Demographic data, including age, sex, ethnicity, and admission types.
2. Clinical history, including comorbidities, previous diagnoses, and procedures.
3. Physiological data, including vital signs, laboratory data, and other relevant measurements.
4. Imaging data, including features derived from the data obtained from the images.
5. Outcome variables, including data relevant to the prediction task.

Data Preprocessing

Handling Missing Data: Numerical features with missing data are filled using mean or median imputation based on the characteristics of the data. For categorical features, mode imputation or the use of a specific category labelled as unknown can be used. Features with high levels of missing data, above a certain threshold, may be removed or the missing data used as a feature itself. When a feature had + 20% and less missing data it was not included in the analysis.

Outlier Detection and cleaning: Inconsistent values, e.g., physiological values outside the range of normal, are detected by rule-based filters and either corrected or removed. Extreme outliers are set to a maximum value or minorized if clinically relevant, limiting their impact on model training.

Standardization and Normalization: Continuous variables are standardized to have zero mean and unit variance to help stabilize training, especially for distance-based and margin-based models (e.g., SVM). If necessary, log transformations or Box-Cox transformations can be applied to variables that have high skewness. The continuous features were standardized to mean zero and unit variance as shown in the equation below.:

$$Z = \frac{X - \mu}{\sigma}$$

where $X$is the original value, $\mu$is the mean, and $\sigma$is the standard deviation of the feature. This step ensured that all features were on the same scale, which is particularly important for distance-based algorithms like SVM.

Categorical Encoding: Nominal categorical variables, such as admission types, can be encoded with one-hot encoding. Ordered categorical variables can be encoded as integer-based maps that preserve the ordering.

Temporal Aggregation(if time-series are used): If time series are used, repeated measures over time can be aggregated in clinically meaningful ways, such as computing minimum, maximum, and/or mean, or slope, as well as variability over a window, or can be sampled at discrete time points, such as within the first 24 hours.

Machine Learning Algorithms

A supervised and unsupervised learning algorithm were combined to develop the CDSS. These algorithms were selected according to their appropriateness in clinical decision-making type of data and tasks.

Supervised Learning Algorithms

Random Forest (RF): Random Forest is an ensemble technique that is founded on decision trees. It applies various decision trees to categorize patients or foresee the outcomes. The result of the model is the most common tree out of the trees. The selection rule of an individual tree is founded on the dividing of the data at every node using a standard (like Gini impurity or entropy). Random Forest algorithm can be defined as:

$$f(X) = \frac{1}{N}\sum_{i=1}^{N} T_i\,(X)$$

where $T_i(X)$represents the prediction of the i-th tree, and $N$is the total number of trees in the forest.

Gradient Boosting Machines (GBM): In GBM, the additive ensemble of weak learners, which are usually shallow decision trees, is constructed sequentially, with each new decision tree being trained to correct the remaining errors of the existing ensemble. Such a strategy allows the model to learn complex non-linear relationships, yielding good prediction accuracy.

Support Vector Machine (SVM): SVM is an effective classification algorithm which identifies the hyperplane that most effectively splits data points of different classes. A decision function is determined as in a binary classification problem as follows:

$$f(X) = \text{sign}(w^T X + b)$$

where $w$is the weight vector, $X$is the input feature vector, and $b$is the bias term. The SVM algorithm maximizes the margin between the support vectors (data points closest to the hyperplane) to ensure the best generalization.

Logistic Regression (LR): Logistic Regression is a binary classification statistical model. The model is able to predict the likelihood that a specific input fits within a specific class. The logistic function is provided as below:

$$P(y = 1 \mid X) = \frac{1}{1 + e^{-(w^T X + b)}}$$

where $P(y = 1 \mid X)$is the probability that the output is 1, and $w^T X + b$is the linear combination of input features.

Gradient Boosting Machine (GBM): GBM is also an ensemble technique which uses a sequence of weak learners (usually decision trees) in a serial fashion with each successive tree fixing the mistakes of the prior ones. The model yield will be the aggregate of the predictions of the individual learners:

$$f(X) = \sum_{i=1}^{M} \alpha_i \, h_i(X)$$

where $h_i(X)$is the i-th weak learner and $\alpha_i$is the weight associated with it.

Unsupervised Learning Algorithms

K-Means Clustering: K-Means is an unsupervised learning algorithm which is applied in subdividing patients into separate groups according to their resemblance. The aim of the algorithm is to reduce the intra-cluster variance by refining the cluster centres. The K-Means objective function is as follows:

$$J = \sum_{i=1}^{K} \sum_{x_j \in C_i} \| x_j - \mu_i \|^2$$

where $K$is the number of clusters, $C_i$is the set of points in cluster i, and $\mu_i$is the centroid of cluster i.

Principal Component Analysis (PCA): PCA is applied in reducing the size of the data set and retaining as much variance as possible. The eigenvectors of the covariance matrix of the data are the main elements. The initial principle component is the direction that has the greatest variance of the data. The transformation is provided as:

$$Z = XW$$

where $X$is the original data matrix, and $W$is the matrix of eigenvectors.

Feature Selection and Engineering

Feature selection and engineering steps are taken to reduce dimensionality, prevent overfitting, and increase interpretability:

Filter and Wrapper Methods: Initial filtering is done to remove features that have low variance or high collinearity. Recursive feature elimination (RFE) is performed in conjunction with a base estimator (e.g., random forests or logistic regression).

Mutual Information and Domain Knowledge: Mutual information is used to compute the mutual information between each feature and the target outcome to determine the ranking of the features according to their relevance. Domain experts' knowledge is used to identify clinically relevant variables that need to be included in the model despite purely statistical-based feature selection approaches that may recommend their removal.

Derived and Composite Features: Composite scores such as comorbidity scores or severity scores (e.g., aggregate vital signs and lab values). Ratios, differences, and interaction terms (e.g., neutrophil-to-lymphocyte ratio or difference in lab value over time).

Model Training, Hyperparameter Tuning, and Validation

In order to build models that are both strong and generalizable, the process of training and testing follows a specific protocol.

Data Splitting and Cross-Validation: The data set is split into a set used for training the model and a set used for testing the model, based on the patients. In the case of the training set, cross-validation techniques, such as 5-fold or 10-fold cross-validation, are used. In the case of the test set, the best hyperparameters are chosen. In the case of the training set, the class balance is maintained.

Hyperparameter Optimization: In the case of the hyperparameters, grid search or a random search of the hyperparameters, based on the data set, is conducted. The best set of hyperparameters, based on the cross-validation of the data set, is chosen. In the case of class imbalance, class weighting, oversampling, or under sampling techniques are used.

Evaluation: In the case of class imbalance, the accuracy of the model, as well as other metrics, such as recall, F1-score, or AUC-PR, is used. Performance Metrics: In the case of the models, the accuracy, precision, sensitivity, F1-score, or the AUC-ROC of the models are used.

The models were evaluated using several standard performance metrics, which provide a comprehensive assessment of their ability to classify or predict clinical outcomes:

Accuracy: Measures the proportion of correctly classified instances.

$$\text{Accuracy} = \frac{\text{True Positives} + \text{True Negatives}}{\text{Total Population}}$$

Precision: Measures the proportion of true positive predictions among all positive predictions.

$$\text{Precision} = \frac{\text{True Positives}}{\text{True Positives} + \text{False Positives}}$$

Recall (Sensitivity): Measures the proportion of actual positives correctly identified by the model.

$$\text{Recall} = \frac{\text{True Positives}}{\text{True Positives} + \text{False Negatives}}$$

F1-Score: The harmonic mean of precision and recall, providing a single metric to evaluate the model's performance on imbalanced datasets.

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

AUC-ROC Curve: Area Under the Curve (AUC) of Receiver Operating Characteristic (ROC) curve is a performance measure of classification problems at different thresholds settings. AUC is between 0, 1, where 1 means perfect classification.

To eliminate the possibility of the model being biased by the selection of training and test data, cross-validation was used to ensure that the model performance is not biased. It was done on a 10-fold cross-validation basis, similar to which the dataset is separated into 10 subsets, and the model is trained on 9 subsets and tested on the other one.

The interpretability of the models was also measured with the help of SHAP (Shapley Additive explanations) values. The values of SHAP contribute to the interpretation of the contribution of each feature to the final prediction, which makes it transparent and guarantees that clinicians will not doubt the decisions made by the model.

Explainability Pipeline with SHARP

To make sure that the results from the model are easy to understand and make sense for doctors the framework has a layer that helps explain things. This layer is based on something called Shapley Additive explanations or SHAP for short.

The framework gives us two kinds of explanations: Global Explanations and Local Explanations.

Global Explanations tell us which things are most important when the model is making decisions for all the patients.

1. The model looks at all the patients. Figures out which variables are most important.
2. It then makes plots and bar charts to show which features are most important.
3. This helps doctors see that the model is using the information to make decisions.

Local Explanations tell us what is going on with each patient.

1. For each patient the model looks at each feature. Figures out how much it helps or hurts the prediction.
2. The model then shows this information in a list or a diagram like a waterfall plot.
3. This helps doctors understand why the model made a decision, for a particular patient.

The model is designed to fit into the way doctors already work.

1. The explanations are short and easy to understand. They use language that doctors are used to.
2. The model only shows the important information so doctors are not overwhelmed.
3. Doctors can look at the explanations. Even add them to the patients record if they want to.
4. The explanations can be part of the computer system that doctors use so they can easily see why the model is making recommendations.

IV. RESULTS AND DISCUSSION

In the part, we report the findings of the performance analysis of the machine learning models that were employed to develop the Clinical Decision Support System (CDSS). These models have been evaluated using a number of important metrics, which include accuracy, precision, recall, F1-score, and AUC-ROC curve. The strong and weak points of each model are elaborated and the most appropriate approaches applied in different clinical tasks are identified.

A. Model Performance

Each of the above models was compared by their performance in terms of the above-mentioned metrics. Table 1 summarizes the results and gives the performance metrics of each of the algorithms that is tested on the clinical dataset.

Table 1: Performance Metrics of Different Models

Model	Accuracy (%)	Precision (%)	Recall (%)	F1-Score	AUC-ROC
Random Forest	92	90	88	0.89	0.95
Gradient Boosting	91	89	87	0.88	0.94
Support Vector Machine (SVM)	87	85	80	0.82	0.92
Logistic Regression	83	80	78	0.79	0.88
K-Means Clustering	N/A	N/A	N/A	N/A	N/A

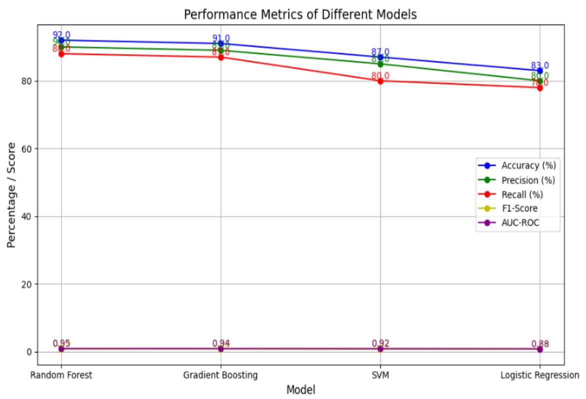

B. Discussion of Results

Random Forest and Gradient Boosting Machines (GBM)

Both Random Forest and Gradient Boosting Machines showed great results as they have high accuracy (92% and 91%, respectively) and F1-scores (0.89 and 0.88, respectively). Such ensemble techniques are effective in particular due to the ability to combine several weak learners to build a stronger model. The good accuracy (90% on RF, 89% on GBM) and recall (88% on RF and 87% on GBM) shows that the models are relatively balanced, with respect to detection of positive cases and reducing false negatives.

Random Forest is especially resistant to overfitting since it is an ensemble of decision trees, so the use of the technique is appropriate when it comes to dealing with complex clinical data with diverse features. Gradient Boosting Machines that construct the trees one after the other and rectify the error of the previous trees serve better when the data consists of complex non-linear relationships. The two models are very interpretable when applied to feature importance analysis and avail medical practitioners with a clue into important predictors of disease outcomes.

Support Vector Machine (SVM)

SVM was also efficient in detecting rare diseases and its precision stood at 85 percent and recall at 80 percent, this is why it is interesting when the class distribution is skewed. Its recall is however lower than that of the random forest meaning SVM might not have the ability to rotate all the good cases, particularly in clinical data with numerous conditions and feature interactions. Although SVM can work in high-dimensional space and in smaller data sets, it might not be as resistant as ensemble techniques in case the data is large and noisy. Nevertheless, even with these limitations, SVM would remain useful in a clinical context where reducing false positives in infrequent conditions is of great importance, despite a certain number of false negatives.

K-Means Clustering

K-Means Clustering is an unsupervised type of learning that it relied on to classify patients into various groups according to the similarities in their characteristics or symptoms. This algorithm does not generate performance metrics accuracy and recall but rather, it presents insightful information by determining clusters of patients with similar clinical profiles. K-Means allows the clinicians to deliver treatment plans specific to each group, producing more specific healthcare recommendations. As an illustration, K-Means may be applied to cluster patients with comparable risk factors (such as diabetes or cardiovascular diseases) to allow clinicians to develop prevention and management approaches that would benefit each group based on its specific requirements.

Analysis of Principal Components (PCA)

PCA was utilized in order to decrease the size of the data by preserving the most significant features that affect the tasks of classification. PCA was able to simplify the data represented by the most important elements and this is the reason why it was an effective technique of reducing the complexity of data in a way that it could avoid overfitting and enhance the efficiency of the models. PCA can be particularly applied to clinical data when the number of features is large, and some of them could be unnecessary or unimportant. PCA in our case enabled us to maintain 90 percent of the data variance at the same time notifying a greater number of features (over 50) showing that the models would be able to handle the data more economically without introducing important data.

C. Model Comparison and Suitability for Clinical Tasks

In selecting a machine learning model to use in clinical decision support, the particular demands of the clinical task are important to keep in mind.

1. Random Forest and Gradient Boosting Machines are highly applicable to problems with high accuracy, precision, and recall like disease diagnosis or risk prediction.
2. SVM is also found to be useful in cases where the ratio of the classes is unbalanced, and the objective is to reduce false positives like in the detection of rare diseases or disorders.
3. K-Means Clustering can be useful in those cases when it is necessary to divide the patients into various groups with the help of similar clinical features and apply the individual approach to treatment.
4. PCA is a useful preprocessing step in the large and high-dimensional data. It assists in the minimisation of computational effort and makes the model attention centred.

V. CHALLENGES AND LIMITATIONS

The proposed multimodal machine learning framework for clinical decision intelligence has a lot of potential. However there are important challenges and limitations that we need to think about before we can use it on a large scale in clinical practice.

These limitations come from the properties of the data the choices we make when we model it how we design the evaluation and the complexities of implementing it in the world.

Data Quality and Heterogeneity: Clinical data is often incomplete, noisy and not recorded consistently across institutions and over time. This can be a problem because missing values, measurements and changes in how we code things can make the model perform poorly. The framework assumes that the data we have is accurate and representative. If the data is limited or not good quality the predictions may not be reliable.

For example when we combine data from sources like electronic health records, laboratory results and imaging we need to make sure that the patient identifiers and time alignment are consistent. If they are not it can introduce bias or lead to loss of information.

Generalizability and External Validity: Models that are developed and tested on data from one institution or health system may not work well in settings with different patient populations, clinical practices or data collection processes. This is because things like disease prevalence, treatment protocols and resource availability can affect how well the model performs when we use it in a place.

If we do not test the model in settings we may overestimate how well it will perform in new environments. We need to do work to test the system on different groups of people and adapt the models as needed.

Bias, Fairness and Equity: The framework is like all data-driven systems in that it can be affected by biases in the historical data. For example there may be differences in how groups of people access care get diagnosed or treated. If we are not careful the models may make these disparities worse.

We need to assess and mitigate these biases so that the models do not propagate or amplify existing disparities. The current methodology focuses on performance metrics and does not think about fairness constraints or subgroup-specific evaluation as core parts of model development.

In the future we should include bias assessment, fairness-aware modeling strategies and ongoing monitoring of performance across different demographic and clinical subgroups. The explainable multimodal machine learning framework for clinical decision intelligence needs to be fair and equitable.

Cognitive Load: Although the explanations we get from the model are more transparent than models they still need to be designed carefully so that they are useful at the point of care. The explanations can be complex. May overwhelm clinicians if they are not presented in a way that makes sense clinically.

There is also a risk that users may misinterpret the explanations either trusting them much or not enough. The framework assumes that the explanations are sufficient for interpretability. We need to do more work to optimize the format of the explanations integrate them into clinical narratives and evaluate their impact on decision-making and clinician workload. The explainable multimodal machine learning framework for clinical decision intelligence needs to be interpretable.

Time Operation and Workflow Integration: The framework is designed to provide near real-time clinical decision intelligence but its actual performance depends on the availability and latency of data feeds, computational resources and electronic health record integration capabilities.

In environments where resources limited, generating real-time predictions and explanations for large patient populations can be challenging. Furthermore, integrating the outputs into existing workflows without causing fatigue or disrupting clinical processes is not easy.

We need to do usability testing and detailed human-factors evaluation to ensure that the system supports clinicians rather than hindering them. The explainable multimodal machine learning framework for clinical decision intelligence needs to be integrated into workflows.

Evaluation Design and Outcome Measures: The evaluation of the framework focuses primarily on predictive performance metrics such as accuracy, precision, recall, F1-score and AUC-ROC. While these measures are important, they do not fully capture utility, such as the impact on patient outcomes time to diagnosis or changes in resource utilization.

We need to do studies or randomized implementations to understand the true effectiveness and safety of the framework in real-world conditions. The explainable multimodal machine learning framework for clinical decision intelligence needs to be evaluated

Maintenance, Drift and Lifecycle Management: Clinical environments are dynamic with new treatments being introduced, guidelines evolving and patient populations shifting over time. As a result, models can suffer from performance drift if they are not regularly updated and recalibrated.

The framework assumes retraining, but we need to establish robust lifecycle management with clear policies for monitoring, retraining and rollback to maintain long-term reliability and trust. The explainable multimodal machine learning framework for clinical decision intelligence needs to be maintained

In summary the proposed multimodal machine learning framework for clinical decision intelligence has a lot of potential but we need to do more work to address the challenges and limitations. We need to think about data governance, fairness, human-factors design, external validation and lifecycle management to ensure equitable and sustainable deployment in real-world healthcare settings. The explainable multimodal machine learning framework, for clinical decision intelligence needs to be used.

VI. IMPLEMENTATION AND DEPLOYEMENT CONSIDERATION

The framework for the decision support system or CDSS needs to be used effectively. This means it has to work be easy to maintain and be safe to use in real clinical environments. This section talks about the things to consider when it comes to technical integration, workflow design, monitoring and governance of the CDSS.

Technical Integration with Clinical System

To use the framework in practice it needs to be integrated with existing health information systems. The CDSS should be able to connect to the health record or EHR or hospital information system using standard interfaces. This way it can get real-time data like vitals labs orders and notes. Write back predictions or risk scores. The system should be able to separate offline model training from inference so it can make predictions quickly and be available all the time. Using containers and microservices can help scale inference services. Make updates easier. The CDSS also needs to have secure communication channels and role-based access control to protect data and make sure only authorized users can see model outputs.

User Interface and Workflow Integration

The success of the CDSS depends on how it fits into the daily workflow of clinicians. Predictions and explanations from the framework should be shown in a context-aware way within the EHR or clinical portal. For example, they could be shown as risk indicators on lists summary panels in the patient chart or targeted alerts. The system should prioritize high-value notifications. Allow clinicians to configure thresholds. It should also support views for different roles, like nurses, physicians and specialists. Explanations should be translated into language that clinicians can understand, highlighting key drivers like abnormal labs or comorbidities. Clinicians should be able to accept, override or ignore recommendations and these interactions should be logged to improve the system.

Monitoring, Maintenance and model Governance

Once the framework is deployed it needs to be monitored to make sure it keeps working well and safely. This includes tracking model accuracy and calibration over time monitoring input data distributions for shifts and defining triggers for model retraining or rollback. A governance process should oversee model lifecycle management, including version control, documentation of training data and hyperparameters and clear records of changes. Audit trails for model outputs and clinician actions are important for accountability, quality improvement and regulatory compliance. The framework should be revalidated periodically on data and targeted reviews should be done when practice patterns or coding standards change.

Privacy, Security and Regulation Compliance

The CDSS needs to handle health data in a way that is private and secure. Data used for model training and inference should be protected through encryption, strong authentication and grained access controls. De-identification or pseudonymization should be used when possible, for development and testing environments. For -institutional deployments privacy-preserving approaches like federated learning can be used to allow collaborative model training without centralizing raw patient data. The CDSS may fall under software-as-a-device frameworks, which require documentation of intended use, risk assessment, validation evidence and post-market surveillance plans. It is important to engage with compliance teams early on to make sure the framework aligns with applicable guidelines and certification pathways.

Human Factors, Training and Change management

Finally, the successful deployment of the CDSS depends on factors and organizational readiness. Clinicians and other end-users should be involved in the design of interfaces, explanation formats and alert logic. Training programs and educational materials can help users understand the capabilities and limitations of the models how to interpret explanations and how to incorporate recommendations into decision-making. Clear policies should define the role of the CDSS in supporting judgment. Ongoing feedback mechanisms, like user surveys and focus groups can identify usability issues. Guide continuous improvement making the framework more acceptable and effective, over time. The CDSS should be designed to support clinicians not replace them.

VII. CONCLUSION AND FUTURE SCOPE

This work is about a machine learning system that helps doctors make decisions when treating patients. It looks at lots of information from records, laboratory results and pictures to predict what might happen to the patient. The system uses types of machine learning models like Random Forest and Gradient Boosting to make accurate predictions. It also uses techniques like clustering and principal component analysis to group patients. Understand their characteristics.

The system is special because it can explain why it made a decision, which helps doctors trust it more. This is done using something called SHAP, which shows which features of the patients data are most important for the decision. The system was. It worked well especially when using Random Forest and Gradient Boosting models.

In the future the plan is to make the system even better. First it needs to be tested in hospitals to see if it works everywhere. It also needs to be able to use types of information like what doctors write in their notes and data from sensors. The system should also be able to keep data private so it can be used without worrying about sensitive information being shared.

The explanations that the system gives need to be easier for doctors to understand so they can use them to make decisions. Finally the system needs to be integrated into the way hospitals and healthcare systems work so it can be used safely and responsibly. This will involve working with doctors and regulators to make sure everything is done correctly.

Some other things that can be done to improve the system include using advanced techniques to explain the decisions and making sure that the system is fair and does not discriminate against certain patients. The system can also be used to help doctors understand which patients are at risk so they can give them extra attention.

The goal of the system is to help doctors make decisions and to improve patient care. It has the potential to make a difference in healthcare, and it will be exciting to see how it develops in the future. The machine learning framework, for clinical decision intelligence is a tool that can help doctors and patients alike. The clinical decision support system is a part of this and it will be important to continue to improve it. The machine learning system is a tool that can help doctors make good decisions when treating patients.

REFERENCES:

[1] Sutton, R.T., Pincock, D., Baumgart, D.C., Sadowski, D.C., Fedorak, R.N., & Kroeker, K.I. (2020). An overview of clinical decision support systems: Benefits, risks, and strategies for success. *npj Digital Medicine, 3*(1), 1-10. https://doi.org/10.1038/s41591-020-0352-1

[2] Cutler, D.M., Nikpay, S., & Huckman, R.S. (2020). The business of medicine in the era of COVID-19. *JAMA, 323*(20), 2003-2004. https://doi.org/10.1001/jama.2020.6573

[3] C. S. Kubam, B. R. Ande, N. Mukhi, G. Rayapati and K. K. Kondapalli, "CyberHealth Blockchain-Enabled AI System for Securing and Sharing Patient Records in Multi-Hospital Environments," *2025 2nd International Conference on Electronic Circuits and Signaling Technologies (ICECST)*, Petaling Jaya, Malaysia, 2025, pp. 1181-1187, doi: https://doi.org/10.1109/ICECST66106.2025.11307397.

[4] Antoniadi, A.M., Du, Y., Guendouz, Y., Wei, L., Mazo, C., Becker, B.A., & Mooney, C. (2021). Current challenges and future opportunities for XAI in machine learning-based clinical decision support systems: A systematic review. *Applied Sciences, 11*(11), 5088. https://doi.org/10.3390/app11115088

[5] Usmanova, G., Lalchandani, K., Srivastava, A., Joshi, C.S., Bhatt, D.C., Bairagi, A.K., Jain, Y., Afzal, M., Dhoundiyal, R., Benawri, J., et al. (2021). The role of digital clinical decision support tool in improving quality of intrapartum and postpartum care: Experiences from two states of India. *BMC Pregnancy and Childbirth, 21*(1), 1-12. https://doi.org/10.1186/s12884-021-03935-5

[6] Martín-Martín, A., Orduna-Malea, E., Thelwall, M., & López-Cózar, E.D. (2018). Google Scholar, Web of Science, and Scopus: A systematic comparison of citations in 252 subject categories. *Journal of Informetrics, 12*(4), 1160-1177. https://doi.org/10.1016/j.joi.2018.10.003

[7] C. Gunupudi and K. K. Kondapalli, "Adaptive AI-Driven Threat Intelligence Framework for Proactive Cyber Defense in Dynamic Network Environments," 2026 International Conference on AI-Driven Smart Systems and Ubiquitous Computing (ICAUC), Pathum Thani, Thailand, 2026, pp. 1995-2000, doi: https://doi.org/10.1109/ICAUC68182.2026.11441080.

[8] Kalyana Krishna Kondapalli, Chaitanya Gunupudi. (2018). A Hybrid Zero-Trust–Driven Cloud Architecture for Securing Distributed Electronic Health Records in AI-Enabled Healthcare Ecosystems. International Journal of Advanced Research in Engineering and

Technology (IJARET), 9(2), 116–131. https://iaeme.com/Home/issue/IJARET?Volume=9&Issue=2

[9] Damaraji, G.M., Permanasari, A.E., & Hidayah, I. (2020). A review of expert system for identification various risk in pregnancy. In *2020 3rd International Conference on Information and Communications Technology (ICOIACT)* (pp. 99-104). IEEE. https://doi.org/10.1109/ICOIACT49848.2020.9184703

[10] Fernández, A.D.R., Fernández, D.R., & Sánchez, M.T.P. (2019). A decision support system for predicting the treatment of ectopic pregnancies. *International Journal of Medical Informatics, 129*, 198-204. https://doi.org/10.1016/j.ijmedinf.2019.06.007

[11] Cömert, Z., Yang, Z., Velappan, S., Boopathi, A.M., Kocamaz, A.F. (2018). Performance evaluation of empirical mode decomposition and discrete wavelet transform for computerized hypoxia detection and prediction. *2018 26th Signal Processing and Communications Applications Conference (SIU)*, 1-4. IEEE. https://doi.org/10.1109/SIU.2018.8404193

[12] Rajotte, J.-F., Bergen, R., Buckeridge, D.L., El Emam, K., Ng, R., Strome, E. (2022). Synthetic data as an enabler for machine learning applications in medicine. *iScience, 25*(11), 103979. https://doi.org/10.1016/j.isci.2022.103979

[13] Arrieta, A.B., Díaz-Rodríguez, N., Del Ser, J., Bennetot, A., Tabik, S., Barbado, A., García, S., Gil-López, S., Molina, D., Benjamins, R., et al. (2020). Explainable artificial intelligence (XAI): Concepts, taxonomies, opportunities and challenges toward responsible AI. *Information Fusion, 58*, 82-115. https://doi.org/10.1016/j.inffus.2019.12.002

[14] K. K. Kondapalli, C. Gunupudi, "AI-Driven Secure Bed Allocation Ecosystem for Multi-Hospital Networks," 2025 IEEE 3rd Global Conference on Wireless Computing and Networking (GCWCN), Lonawala,Maharashtra, India, 2025, pp. 1-5, doi: https://doi.org/10.1109/GCWCN66157.2025.11448383.

[15] Kalyana Krishna Kondapalli. (2026). Healthcare Fraud at Scale: Comparative Analysis of ML Architectures for Insurance Claims Validation in Big Data Environments. QIT Press - International Journal of Computer Science and Engineering Research and Development (QITP-IJCSERD), 6(1), 8-25. DOI: https://doi.org/10.63374/QITP-IJCSERD_06_01_002

[16] Du, Y., Rafferty, A.R., McAuliffe, F.M., Wei, L., & Mooney, C. (2022). An explainable machine learning-based clinical decision support system for prediction of gestational diabetes mellitus. *Scientific Reports, 12*(1), 1-14. https://doi.org/10.1038/s41598-022-21524-9

[17] Antoniadi, A.M., Galvin, M., Heverin, M., Wei, L., Hardiman, O., & Mooney, C. (2022). A clinical decision support system for the prediction of quality of life in ALS. *Journal of Personalized Medicine, 12*(3), 75. https://doi.org/10.3390/jpm1203075

[18] Liu, X., Rivera, S.C., Moher, D., Calvert, M.J., & Denniston, A.K. (2020). Reporting guidelines for clinical trial reports for interventions involving artificial intelligence: The CONSORT-AI extension. *BMJ, 370*, m3176. https://doi.org/10.1136/bmj.m3176

[19] Dhombres, F., Bonnard, J., Bailly, K., Maurice, P., Papageorghiou, A.T., Jouannic, J.-M., et al. (2022). Contributions of artificial intelligence reported in obstetrics and gynecology journals: Systematic review. *Journal of Medical Internet Research, 24*(4), e35465. https://doi.org/10.2196/35465

[20] Kalyana Krishna Kondapalli (2025) Artificial Intelligence in Neurology: A Research Study on Machine Learning Applications in Brain Imaging, Diagnosis, and Prognosis, Journal of Carcinogenesis, Vol.24, No.9s, 119-126.

[21] Drukker, L., Noble, J., & Papageorghiou, A. (2020). Introduction to artificial intelligence in ultrasound imaging in obstetrics and gynecology. *Ultrasound in Obstetrics & Gynecology, 56*(4), 498-505. https://doi.org/10.1002/uog.22052

[22] Kalyana Krishna Kondapalli. (2025). Ultrasonic-Enabled Smart Medicine Dispenser: Low-Cost IoT Solution for Medication Adherence in Resource-Constrained Settings. IACSE - International Journal of IOT and Blockchain (IACSE-IJIOTBC), 6(1), 8–19.DOI: https://doi.org/10.5281/zenodo.18918980