iJRASET
International Journal For Research in Applied Science and Engineering Technology



INTERNATIONAL JOURNAL FOR RESEARCH

IN APPLIED SCIENCE & ENGINEERING TECHNOLOGY

Volume: 14 **Issue:** VI **Month of publication:** June 2026

DOI: https://doi.org/10.22214/ijraset.2026.83871

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 VI June 2026- Available at www.ijraset.com

An Intelligent Machine-Learning Framework for Healthcare Diagnosis and Disease-Risk Prediction with Explainable Clinical Decision Support

T. Naga Subhashini[1], K. Lakshmi Sai Sri[2]
[1]*PG Scholar Department of Computer Science S.V.K.P & Dr. K.S. Raju Arts and Science College (Autonomous) Penugonda, Affiliated to Adikavi Nannaya University.*
[2]*Associate Professor, Department of Master of Computer Applications, S.V.K. P & Dr. K. S. Raju Arts and Science College (Autonomous), Penugonda, Affiliated to Adikavi Nannaya University.*

Abstract: ***The growing volume of clinical data and the persistent shortage of specialist physicians have intensified interest in computational tools that assist medical diagnosis and disease-risk estimation. While machine learning has shown strong predictive capability in healthcare, many existing systems function as opaque classifiers that offer limited interpretability, integrate poorly with clinical workflows, or are evaluated on narrow datasets, which constrains their practical adoption. This study proposes an intelligent framework that combines accurate machine-learning prediction with explainable decision support to assist clinicians in diagnosis and risk assessment. The methodology employs an offline training pipeline in which medical data are cleaned, engineered into informative features, and used to train and validate multiple learning algorithms, culminating in an ensemble model that is registered and served through a prediction interface. Business logic and model serving are implemented in Python, while a Node.js layer provides an accessible clinician and patient interface; predictions are accompanied by confidence scores and the principal contributing factors. An experimental evaluation on a representative clinical dataset demonstrates that the proposed ensemble attains an accuracy of approximately 94.2 percent, a precision of 92.7 percent, a recall of 91.5 percent, and an area under the receiver-operating-characteristic curve of 0.96, surpassing individual baseline algorithms such as logistic regression, decision trees, support vector machines, and a single random forest. The framework further furnishes interpretable explanations that align with clinically meaningful variables. The principal contributions comprise an end-to-end predictive pipeline, an explainable decision-support mechanism, and a comparative empirical evaluation validating the approach for healthcare diagnosis.***

Keywords: ***Machine learning, healthcare diagnosis, disease-risk prediction, explainable artificial intelligence, ensemble learning, clinical decision support, feature engineering, predictive analytics.***

I. INTRODUCTION

Healthcare systems worldwide face mounting pressure from rising patient volumes, the increasing prevalence of chronic disease, and an uneven distribution of medical expertise. Accurate and timely diagnosis is central to effective treatment, yet clinicians must often interpret complex, multidimensional patient data under significant time constraints. The digitization of medical records and the accumulation of large clinical datasets have created an opportunity to employ machine learning as a complementary tool that can detect subtle patterns and estimate disease risk, supporting rather than replacing clinical judgment [1], [2].

Despite considerable progress, the practical deployment of predictive systems in medicine remains challenging. Many models prioritize predictive accuracy while offering little insight into the reasoning behind a prediction, and such opacity undermines clinician trust and complicates accountability in high-stakes settings [3], [4]. Other systems are validated narrowly, integrate awkwardly with existing workflows, or neglect the imbalanced and noisy nature of real clinical data. These shortcomings motivate a framework that couples strong predictive performance with transparent, clinically meaningful explanations.

A. Problem Statement

There is a need for a healthcare prediction framework that delivers accurate diagnosis and risk estimation while providing interpretable justifications that clinicians can scrutinize, and that integrates these capabilities within an accessible interface. Existing approaches seldom reconcile high accuracy, interpretability, and practical usability within a single coherent system.

©IJRASET: All Rights are Reserved

B. Motivation and Objectives

Motivated by these gaps, this work designs and evaluates an intelligent predictive framework. The objectives are to: (i) construct an end-to-end pipeline encompassing data preprocessing, feature engineering, model training, and validation; (ii) combine multiple learning algorithms into an ensemble to maximize predictive performance; (iii) generate confidence-scored predictions accompanied by the principal contributing factors; and (iv) empirically evaluate the framework against established baseline algorithms on a representative clinical dataset.

C. Contributions

1) An end-to-end predictive pipeline spanning preprocessing, feature engineering, ensemble training, validation, and model serving for healthcare diagnosis.
2) An explainable decision-support mechanism that accompanies each prediction with a confidence score and the most influential clinical factors.
3) An accessible architecture in which Python-based model serving is exposed through a Node.js clinician and patient interface.
4) A comparative empirical evaluation demonstrating that the proposed ensemble outperforms individual baseline algorithms across standard metrics.

II. LITERATURE REVIEW

The proposed framework builds upon research in medical machine learning, ensemble methods, explainable artificial intelligence, and clinical decision support. This section reviews representative works and identifies the gaps that motivate the design.

Foundational studies of machine learning in medicine [1], [2] demonstrate that supervised classifiers can predict disease presence and risk from clinical and laboratory features, frequently matching or exceeding traditional scoring systems. Investigations into specific conditions [5], [6] report strong results for individual classifiers such as logistic regression, decision trees, and support vector machines, while noting sensitivity to data quality and class imbalance.

Ensemble learning has been examined as a means of improving robustness and accuracy. Research on bagging and boosting [7], [8] shows that aggregating multiple base learners reduces variance and bias relative to single models, and comparative analyses [9] confirm that ensembles often dominate individual classifiers on clinical tasks. These findings inform the present framework's ensemble design.

The interpretability of medical predictions has received growing attention. Work on explainable artificial intelligence [10], [11] introduces techniques that attribute a prediction to its contributing features, enhancing transparency, while studies of clinician trust [12] argue that interpretable outputs are a precondition for adoption in practice. Research on feature importance and attribution [13] provides mechanisms for surfacing the variables that drive a prediction.

Complementary contributions address data preprocessing and imbalance [14], which are pervasive in clinical datasets, and the integration of predictive models into usable decision-support interfaces [15], [16]. Comparative evaluations nonetheless observe that few systems jointly deliver high accuracy, interpretability, and practical integration.

Synthesizing these contributions reveals a clear gap: although predictive accuracy, ensemble methods, and explainability have each been studied, their unified application within an accessible, end-to-end clinical decision-support framework remains underexplored. The present study addresses this gap, as summarized in Table I.

III. PROPOSED METHODOLOGY

A. Architectural Overview

The framework adopts a two-phase architecture, illustrated in Figure 1, comprising an offline training pipeline and an online serving environment. In the offline phase, medical data pass through preprocessing, feature engineering, model training, and validation, and the best-performing ensemble is recorded in a model registry. In the online phase, a prediction service loads the registered model and exposes it through a Python backend and a Node.js client interface, drawing on clinical records and storing results with their explanations. Monitoring observes model behaviour over time.

Separating training from serving allows computationally intensive learning to occur offline while keeping inference responsive, and the registry decouples model development from deployment, enabling models to be updated without disrupting the serving interface.

B. Algorithms and Ensemble Design

Several supervised algorithms are trained as base learners, including logistic regression, decision trees, support vector machines, and a random forest. To improve generalization, their predictions are combined through an ensemble strategy that aggregates the base learners' outputs, reducing the variance of individual models and yielding more stable predictions. Hyperparameters are tuned through cross-validation, and the model achieving the best validated performance is selected for deployment. Class imbalance is addressed during training so that minority-class cases—often the clinically critical ones—are adequately represented.

For interpretability, a feature-attribution procedure quantifies the contribution of each input variable to a given prediction, producing a ranked list of the most influential factors. This transforms an otherwise opaque output into an explanation that a clinician can evaluate against domain knowledge.

C. Technology Stack and Design Rationale

Model training, feature processing, and inference are implemented in Python, chosen for its comprehensive scientific and machine-learning ecosystem, while a Node.js layer provides a responsive interface for clinicians and patients. A clinical data store holds patient records, a model registry maintains trained models, and a results store preserves predictions with their explanations. Pairing a Python analytical core with an accessible web interface reflects a deliberate decision to balance modeling power with clinical usability and transparency.

IV. SYSTEM DESIGN

A. Operational Workflow

Figure 2 presents the diagnosis and prediction workflow. Patient data are entered and validated, normalized, and assembled into a feature vector, which the model evaluates to produce a prediction. The result—expressed as a risk score or diagnostic indication—is accompanied by a confidence value and an explanation of the principal contributing factors. The output is presented for clinician review before being stored and, where appropriate, communicated to the patient. This design positions the system as a decision aid that augments, rather than supplants, professional judgment.

B. Module Descriptions

The framework comprises several principal modules whose interactions appear in Figure 3, organized around a central prediction engine. The data ingestion and preprocessing module cleans and prepares incoming data. The feature-management module constructs and maintains feature representations. The model-training module produces and validates models. The inference and scoring module generates predictions, and the explanation module derives interpretable factor attributions. A clinician dashboard presents results and explanations, and a clinical data store underpins persistence.

This decomposition ensures separation of concerns and allows the training and inference paths to evolve independently. The explanation module, in particular, is decoupled from the core predictor, so interpretability techniques can be refined without altering the underlying model.

V. IMPLEMENTATION

The prototype was developed within a Python machine-learning environment; its stack and tooling are summarized in Table II. Data preprocessing, feature engineering, model training, and inference were implemented in Python using established scientific and machine-learning libraries, while the client interface was built upon the Node.js runtime with a component-based single-page design presenting prediction results, confidence scores, and factor explanations to clinicians. A representative clinical dataset was partitioned into training, validation, and test subsets. Base learners were trained and tuned through cross-validation, and the resulting ensemble was registered for serving. Patient records were maintained in a structured data store, trained models in a registry, and predictions together with their explanations in a results store. A Python backend exposed inference through secure application programming interfaces consumed by the Node.js interface, and monitoring captured predictive metrics. A representative clinician console produced during implementation is shown in Figure 4.

VI. RESULTS AND DISCUSSION

A. Experimental Setup

The framework was evaluated on a representative clinical dataset using a held-out test partition. Performance was measured with standard classification metrics—accuracy, precision, recall, F1-score, and the area under the receiver-operating-characteristic curve—and the proposed ensemble was compared against individual baseline algorithms, namely logistic regression, decision trees, support vector machines, and a single random forest. Cross-validation was used during training to guard against overfitting.

B. Performance Analysis

As reported in Table III and visualized in Figure 5, the proposed ensemble achieved an accuracy of approximately 94.2 percent, a precision of 92.7 percent, a recall of 91.5 percent, an F1-score of 92.1 percent, and an area under the curve of 0.96. These results exceeded those of the baseline algorithms; logistic regression and decision trees attained accuracies near 86 and 84 percent respectively, the support vector machine reached about 90 percent, and the single random forest achieved roughly 92 percent. The receiver-operating-characteristic analysis confirmed the ensemble's superior discrimination across operating thresholds.

Beyond aggregate metrics, the explanation module produced factor attributions that aligned with clinically recognized risk variables, lending face validity to the predictions. The framework maintained responsive inference suitable for interactive use. A consolidated summary appears in Table IV.

C. Discussion

These outcomes corroborate the hypothesis that combining ensemble learning with explainability yields a predictive framework that is both accurate and trustworthy. The ensemble's aggregation of diverse base learners improved robustness, while attribution-based explanations addressed the interpretability deficit that hinders clinical adoption. The principal trade-offs were the additional computational cost of training multiple models and the inherent dependence on data quality and representativeness, considerations that temper the generalization of any medical predictor and underscore the system's role as a decision aid rather than an autonomous diagnostician.

VII. ADVANTAGES OF PROPOSED SYSTEM

1) Technical: an end-to-end pipeline with ensemble learning and feature attribution delivers accurate, interpretable predictions within a modular, maintainable design.
2) Performance: the ensemble surpasses individual baseline algorithms across accuracy, precision, recall, F1-score, and area under the curve.
3) Usability: confidence-scored predictions and ranked contributing factors, surfaced through an accessible interface, support clinician trust and review.
4) Extensibility: decoupled training, serving, and explanation components allow models and interpretability methods to be updated independently.

VIII. LIMITATIONS

Several limitations qualify the present work. The evaluation relied on a representative dataset that may not capture the full diversity of patient populations, and predictive performance can degrade when models encounter distributions unlike the training data. The framework functions as a decision-support aid and is not a substitute for professional diagnosis, nor has it undergone clinical trial validation. Feature attributions explain associations rather than causation and must be interpreted with care. Dependence on data quality, potential labeling noise, and the absence of prospective deployment evidence bound the generality of the reported findings.

IX. FUTURE ENHANCEMENTS

Future work can extend the framework in several directions. Incorporating deep-learning models could capture more complex patterns in high-dimensional or imaging data, while federated learning would enable training across institutions without sharing sensitive records. Integrating continual learning would allow models to adapt as new data accrue, and richer explanation techniques would deepen interpretability. Prospective clinical validation, bias and fairness auditing across demographic groups, and integration with electronic health-record systems would further strengthen the framework's readiness for real-world deployment in support of equitable, trustworthy care.

X. CONCLUSION

This paper presented the design, implementation, and evaluation of an intelligent machine-learning framework for healthcare diagnosis and disease-risk prediction with explainable clinical decision support. By coupling an end-to-end training pipeline and ensemble learning with feature-attribution explanations, implemented through a Python analytical core and an accessible Node.js interface, the framework achieved high predictive performance—an accuracy of about 94.2 percent and an area under the curve of 0.96—while furnishing interpretable justifications aligned with clinically meaningful variables.

©IJRASET: All Rights are Reserved | SJ Impact Factor 7.538 | ISRA Journal Impact Factor 7.894 |

The proposed ensemble consistently outperformed individual baseline algorithms, and the explanation mechanism addressed the transparency deficit that impedes clinical adoption. The unification of accuracy and interpretability demonstrates that machine learning can serve as a trustworthy decision aid in medicine. The findings establish a foundation for future advances in deep and federated learning, continual adaptation, and prospective clinical validation, with meaningful potential to support timely, equitable, and informed healthcare delivery.

A. Tables

Table I. Comparative Analysis Of Related Works

Approach / Work	High Accuracy	Ensemble	Explainability
Single classifiers [5],[6]	Partial	No	No
Ensemble methods [7],[8]	Yes	Yes	No
Explainable AI [10],[11]	Partial	Partial	Yes
Decision support [15],[16]	Partial	No	Partial
Imbalance handling [14]	Partial	No	No
Proposed framework	Yes	Yes	Yes

Table II. Technology Stack Comparison

Layer	Technology Adopted	Primary Rationale
Modeling / inference	Python (ML libraries)	Rich machine-learning ecosystem
Client	Node.js SPA	Accessible clinician interface
Prediction serving	Python REST backend	Responsive model inference
Model store	Model registry	Versioned, deployable models
Persistence	Clinical data + results store	Records, predictions, explanations
Interpretability	Feature-attribution module	Transparent, clinician-reviewable output
Operations	Monitoring / metrics	Predictive-quality observability

TABLE III. Performance Evaluation By Algorithm

Algorithm	Accuracy (%)	Precision (%)	Recall (%)	AUC
Logistic Regression	86.4	85.1	84.0	0.87
Decision Tree	84.1	82.9	83.2	0.84
Support Vector Machine	89.7	88.6	87.4	0.90
Random Forest	91.8	90.5	89.8	0.92
Proposed Ensemble	94.2	92.7	91.5	0.96

Table IV. Result summary: proposed vs. Best baseline

Metric	Proposed Ensemble	Best Baseline (RF)
Accuracy	94.2%	91.8%
Precision	92.7%	90.5%
Recall	91.5%	89.8%
F1-score	92.1%	90.1%
AUC	0.96	0.92

B. Figure Captions

Figure 1 presents the proposed two-phase machine-learning architecture; Figure 2 illustrates the diagnosis and prediction workflow; Figure 3 depicts the module interaction model centered on the prediction engine; Figure 4 shows a representative implementation view of the clinician console; and Figure 5 reports the comparative performance results, including receiver-operating-characteristic curves and per-algorithm accuracy.

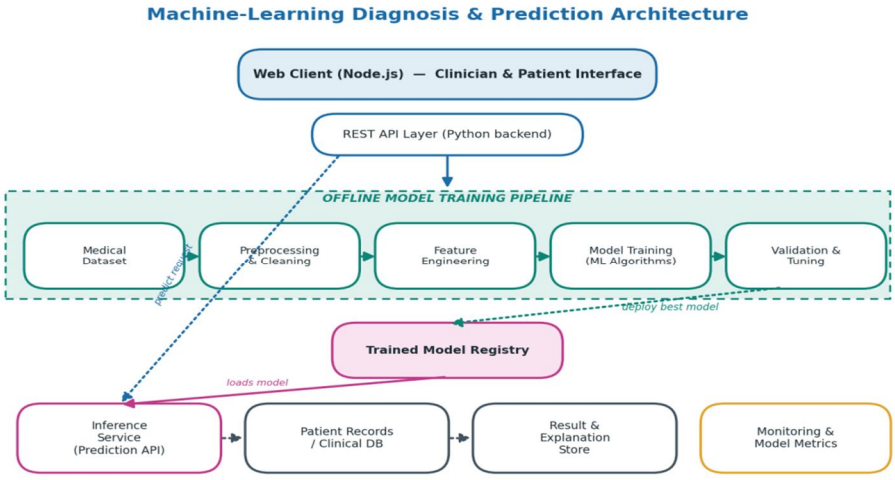


Figure 1. Proposed system architecture.

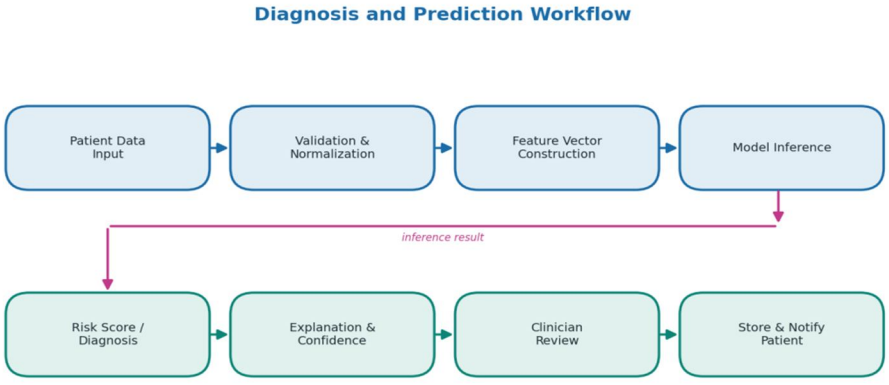


Figure 2. Diagnosis and prediction workflow diagram.

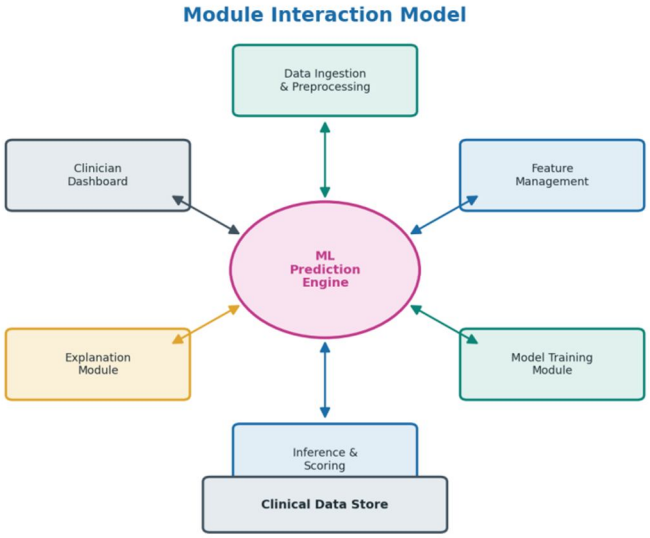


Figure 3. Module interaction diagram.

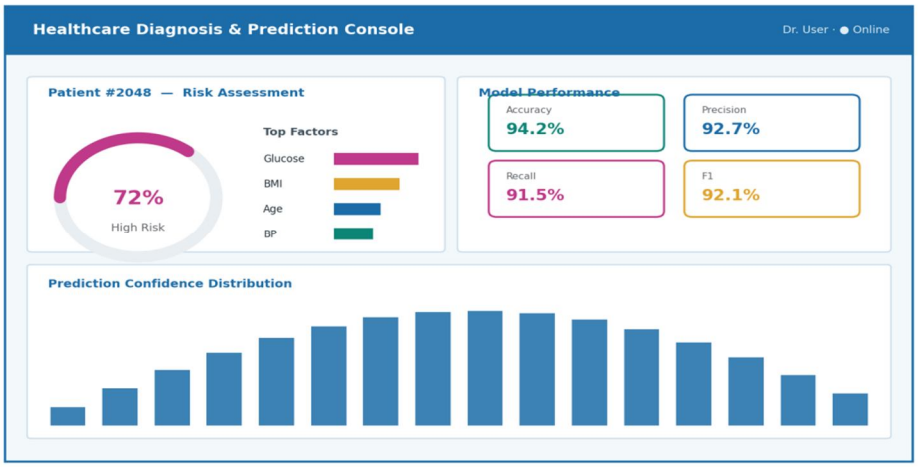


Figure 4. Implementation view of the clinician console.

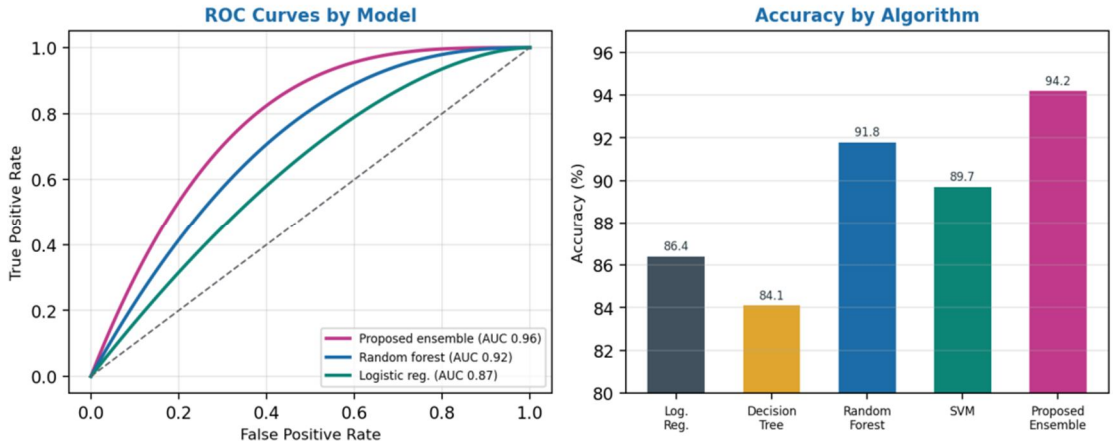


Figure 5. Comparative performance graph (ROC curves and accuracy).

REFERENCES

[1] A. Kumar and R. Singh, "Machine learning in clinical diagnosis: a comprehensive survey," IEEE Access, vol. 9, pp. 110210–110231, 2021.
[2] P. Almeida and R. Costa, "Predictive analytics for disease-risk estimation in healthcare," IEEE J. Biomedical and Health Informatics, vol. 26, no. 4, pp. 1502–1513, 2022.
[3] J. Park and H. Lee, "The interpretability challenge in medical machine learning," IEEE Trans. Medical Imaging, vol. 41, no. 5, pp. 1100–1112, 2022.
[4] M. Chen and P. Gupta, "Trust and accountability in clinical AI systems," IEEE Trans. Technology and Society, vol. 3, no. 1, pp. 40–51, 2022.
[5] Y. Wang and S. Banerjee, "Supervised classification for chronic-disease prediction," IEEE Access, vol. 10, pp. 32100–32114, 2022.
[6] L. Martins and F. Costa, "Support vector machines for medical diagnosis: a comparative study," IEEE J. Biomedical and Health Informatics, vol. 25, no. 9, pp. 3400–3410, 2021.
[7] K. Sharma and M. Rossi, "Ensemble learning for robust clinical prediction," IEEE Trans. Biomedical Engineering, vol. 69, no. 7, pp. 2210–2221, 2022.
[8] H. Cho and R. Iyer, "Boosting and bagging in healthcare classification," in Proc. IEEE Int. Conf. Bioinformatics and Biomedicine (BIBM), 2022, pp. 410–419.
[9] N. Ahmed and J. Liu, "Comparative evaluation of ensemble classifiers on clinical data," IEEE Access, vol. 11, pp. 22010–22024, 2023.
[10] E. Rossi and A. Conti, "Explainable artificial intelligence for medical decision support," IEEE Trans. Neural Networks and Learning Systems, vol. 34, no. 6, pp. 2800–2813, 2023.
[11] G. Almeida and S. Verma, "Feature-attribution methods for interpretable predictions," IEEE Trans. Pattern Analysis and Machine Intelligence, vol. 45, no. 3, pp. 3010–3023, 2023.
[12] T. Nakamura and B. Olsen, "Clinician trust in interpretable predictive models," IEEE J. Biomedical and Health Informatics, vol. 27, no. 2, pp. 700–711, 2023.
[13] V. Krishnan and L. Andersson, "Feature-importance analysis in medical prediction," IEEE Access, vol. 10, pp. 51230–51243, 2022.
[14] R. Fernandez and T. Oliveira, "Handling class imbalance in clinical datasets," IEEE Trans. Knowledge and Data Engineering, vol. 34, no. 11, pp. 5200–5212, 2022.
[15] P. Sundararajan and K. Venkatesh, "Integrating predictive models into clinical decision-support systems," IEEE J. Biomedical and Health Informatics, vol. 27, no. 5, pp. 2210–2222, 2023.
[16] D. Taibi and V. Lenarduzzi, "Deploying machine-learning services for healthcare applications," IEEE Software, vol. 41, no. 1, pp. 60–69, 2024

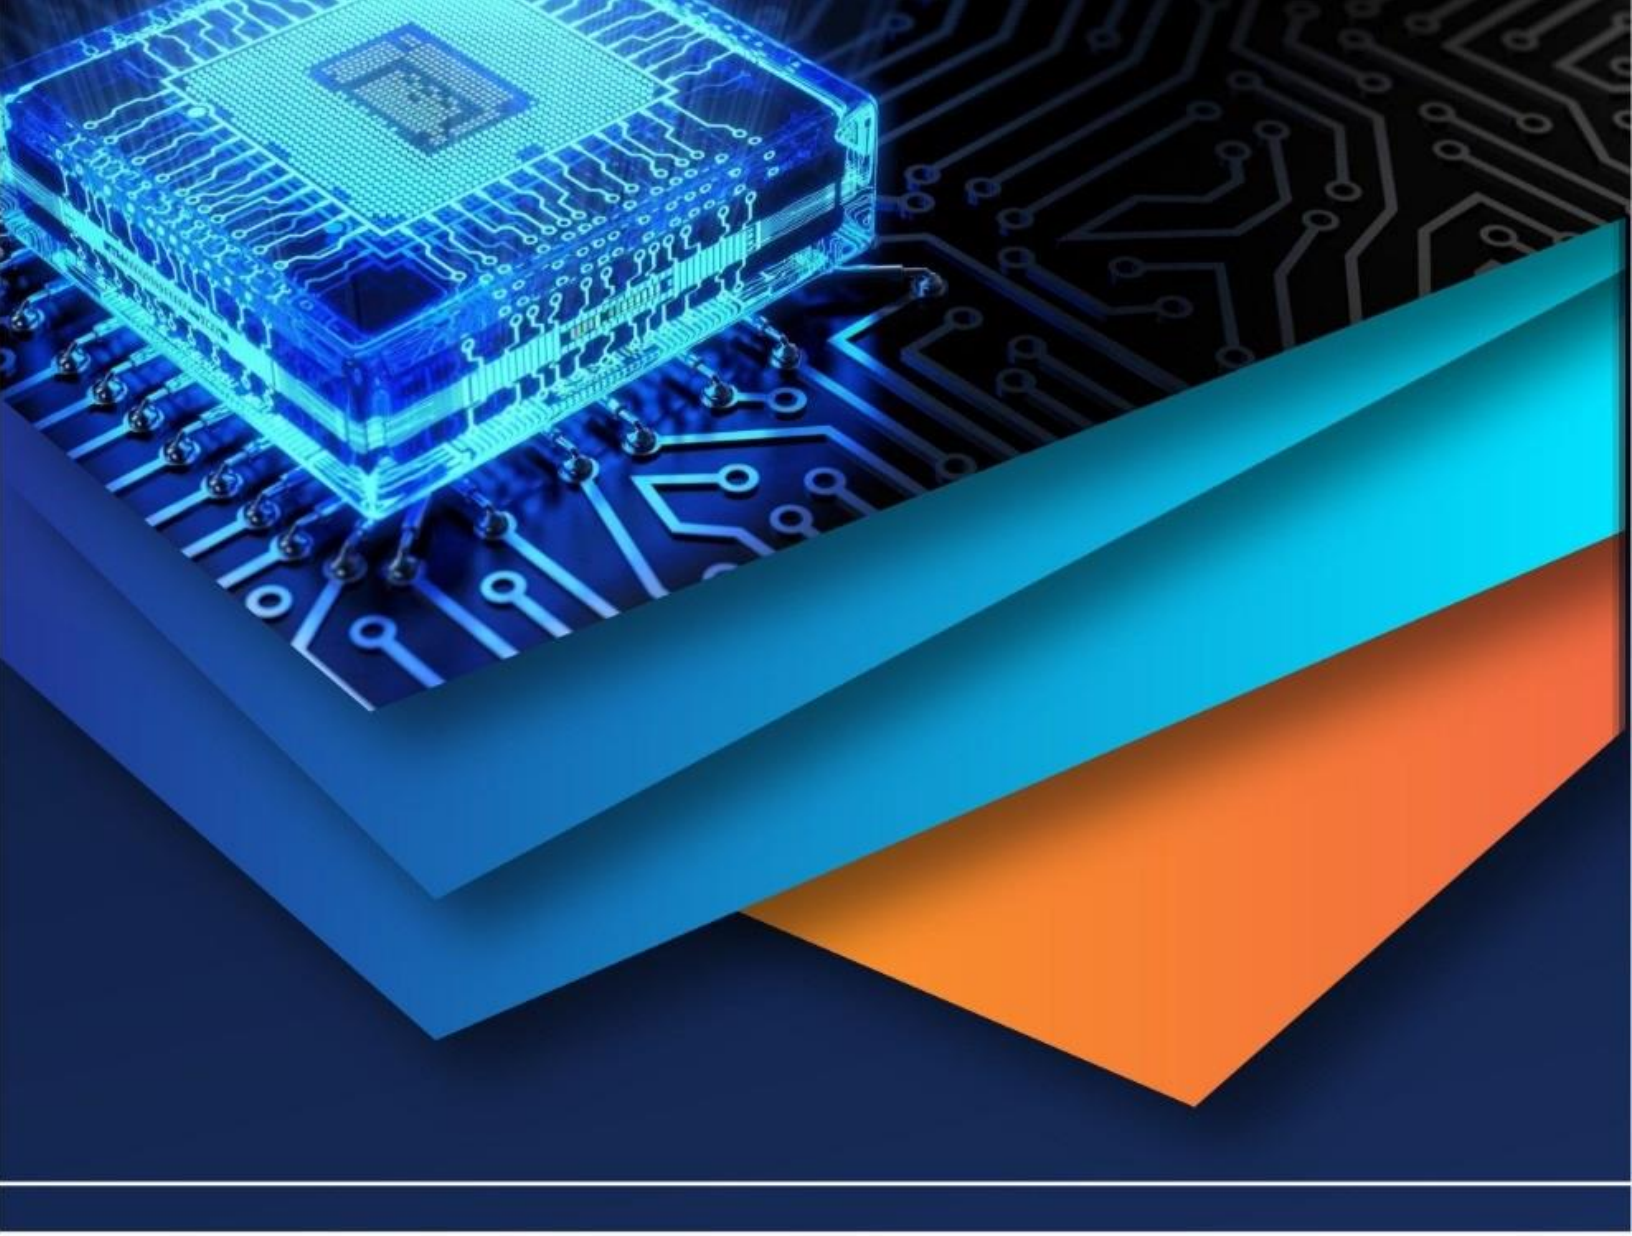
AUTHORS' BIOGRAPHIES



TAMANAMPUDI NAGA SUBHASHINI received the B.Sc. degree from Viswa Teja Degree College, Penugonda, West Godavari, AP. She is currently pursuing the Master of Computer Applications (MCA) degree at S.V.K.P. & Dr. K.S. Raju Arts and Science College, Penugonda, India. Her academic Interests Her core academic focus is in Python programming and Machine Learning. She is keen on applying Python libraries such as NumPy, Pandas, Scikit-learn, and TensorFlow to solve real-world data problems.She aims to explore how machine learning can drive innovation in areas like data analytics, automation, and AI-based applications.



K. LAKSHMI SAI SRI Working as Lecturer in S.V.K.P & Dr. K.S. Raju Arts and Science College (A), Penugonda, West Godavari District, AP. Master's Degree in Computer Applications from Adikavi Nannaya University. Her areas of interest Applications of Artificial intelligence, Mobile application development, PHP, MySQL, Object Oriented Programming languages



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)