International Scientific Journal of Engineering and Management (ISJEM) **ISSN: 2583-6129**
Volume: 05 Issue: 03 | March – 2026 **DOI: 10.55041/ISJEM05854**
An International Scholarly || Multidisciplinary || Open Access || Indexing in all major Database & Metadata

CAT Boost-Driven Multi-Condition Risk Prediction for Newborns in Neonatal Intensive Care using Structured Clinical Data

A Sai Prasad [1], N Thanuja [2], K Uday Kumar[3], G Mounika [4] , S Praveen [5]

[1] *Senior Assistant Professor ,Computer Science and Engineering, Sanketika Vidya Parishad Engineering College, Vishakhapatnam, Andhra Pradesh, India*

[2,3,4,5] *B.Tech Final Semester, Bachelor of Technology, Computer Science and Engineering, Sanketika Vidya Parishad Engineering College, Vishakhapatnam, Andhra Pradesh, India*

---***---

Abstract - Early identification of health risks in newborns admitted to the Intensive Care Unit (ICU) is essential for enabling timely clinical intervention and improving survival outcomes. Traditional neonatal monitoring methods primarily rely on fixed threshold-based evaluations and manual assessments, which often fail to capture complex relationships among multiple physiological and clinical parameters. To address these limitations, this study proposes a machine learning–based framework for early risk detection in newborns using structured data, including gestational age, anthropometric measurements, vital signs, feeding patterns, APGAR scores, and biochemical indicators. The proposed system utilizes a CATBoost-based multi-class classification model to predict condition-specific risk levels associated with jaundice, cardiac, and respiratory complications. The model is well-suited for healthcare applications due to its efficient handling of categorical features and robustness against overfitting. Furthermore, the system is supported by a modular architecture that integrates automated data preprocessing, model training, and real-time inference through secure APIs. An interactive user interface enables visualization of risk predictions and clinical insights, thereby assisting healthcare professionals in making informed, data-driven decisions in critical care environments.

Key Words: Neonatal risk prediction, Machine learning, CATBoost, Intensive care unit, Healthcare analytics, Clinical decision support.

1. INTRODUCTION

The Neonatal Intensive Care Unit (NICU) is one of the most critical environments in modern medicine. Worldwide, approximately 2.4 million neonatal deaths occur annually, many attributable to conditions that are detectable if clinical deterioration is identified early enough[12]. Newborns admitted to NICUs are physiologically fragile, as their organ systems remain underdeveloped, and even minor deviations in vital signs may signal life-threatening events[3]. The three most prevalent conditions in NICU settings are neonatal jaundice, cardiac complications, and respiratory distress, each contributing significantly to neonatal morbidity and mortality.

Conventional NICU monitoring depends on hardware-based threshold alarms that trigger only when a single physiological parameter crosses a fixed limit[4]. These systems operate reactively and are incapable of detecting gradual multi-parameter deterioration. Furthermore, the high false-positive alarm rate produces alert fatigue among clinical staff, which paradoxically delays responses during genuine emergencies[56]. Most existing machine learning solutions for neonatal healthcare address only a single condition, lack real-time inference capability, and provide no clinical-grade user interface for NICU environments[78].

This paper presents a CATBoost-based multi-class classification framework that simultaneously predicts Low, Medium, and High risk for three neonatal conditions from 22 structured clinical features[9]. The system is deployed as a complete full-stack application comprising a FastAPI backend, React.js clinical dashboard, and MongoDB database, achieving 92.00% overall accuracy with zero misclassification between extreme risk classes[10].

2. Body of Paper

2.1 LITERATURE REVIEW

Early machine learning research in neonatal healthcare focused on sepsis prediction[11]. Mani et al. applied statistical methods and basic classification to neonatal sepsis, establishing the feasibility of data-driven NICU approaches but suffering from small sample sizes and a single-condition focus[12]. Desautels et al. extended this using Logistic Regression and Decision Trees but reported high false-positive rates, which are unacceptable in clinical environments already affected by alert fatigue[13].

Aczon et al. employed Random Forest and time-series analysis for neonatal mortality prediction, achieving promising accuracy but demonstrating poor cross-hospital generalization[14]. Deep learning architectures introduced around 2018 showed improved accuracy on neonatal monitoring tasks but required GPU hardware that is impractical for most clinical environments[1516]. Harutyunyan et al. conducted multitask learning benchmarking on clinical time-series data and highlighted the challenge of acquiring sufficiently labelled neonatal datasets[17].

The emergence of gradient boosting techniques significantly improved structured clinical data classification[18]. Prokopenko et al. demonstrated that CATBoost's ordered boosting mechanism and native handling of categorical features enable superior generalization on mixed-type datasets[19]. Ribeiro et al. emphasized that clinical machine learning systems must provide interpretable outputs for regulatory and clinical acceptance, motivating the use of CATBoost's feature importance capabilities[20]. A persistent gap in the reviewed literature is the absence of a system capable of predicting multiple neonatal conditions simultaneously through a real-time clinical interface, which this work aims to address[21].

2.2 EXISTING SYSTEM

Existing neonatal monitoring systems in NICUs rely on fixed threshold-based methods to track individual physiological parameters such as heart rate and oxygen saturation[1]. These systems are reactive and fail to capture multi-parameter deterioration[2]. While some machine learning models have been introduced, most focus on single-condition prediction and lack real-time deployment and clinical integration[3]. Additionally, high false alarm rates lead to alert fatigue among healthcare professionals, reducing response efficiency[4]. Overall, existing systems lack a unified and real-time framework for comprehensive neonatal risk assessment[5].

2.3 PROPOSED SYSTEM

2.3.1 System Architecture

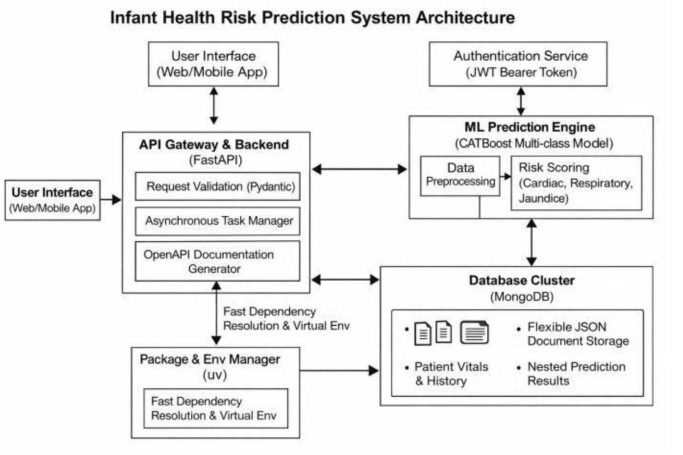

Fig -1:Infant Health Risk Prediction System Architecture

The proposed system follows a five-layer modular architecture designed for scalability, security, and real-time clinical responsiveness[22]. The Presentation Layer is implemented as a React.js single-page application that provides patient data entry forms and a real-time risk dashboard with colour-coded indicators (green: Low, amber: Medium, red: High). The API Layer is developed using FastAPI, which manages request routing, performs input validation using Pydantic models, and enables asynchronous inference processing[10].

Security is ensured through JWT-based authentication, which enforces role-based access control for different users, including nurses, doctors, and administrators. The Machine Learning Layer hosts the pre-trained and serialized CATBoost model, which is loaded during server initialization to ensure low-latency predictions. The Data Persistence Layer utilizes MongoDB to store patient records and prediction histories, maintaining timestamped logs for traceability and audit purposes[23].

2.2.2 System Flow

As shown in Fig. 2, the end-to-end data flow begins when a healthcare professional submits patient clinical data through the frontend form[22]. The data passes through a preprocessing pipeline before being fed into the trained CATBoost multi-class classifier, which outputs probability scores for Low, Medium, and High risk levels across all three conditions simultaneously. These probability scores are then mapped to condition-specific risk indicators and stored in MongoDB for continuous and longitudinal monitoring[9].

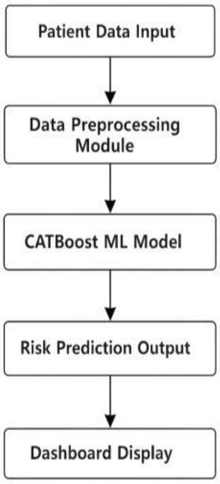

Fig -2:System Block Diagram – Data Flow from Input to Dashboard

2.3.3 Dataset and Feature Set

The model is trained on a structured neonatal clinical dataset comprising approximately 1,000 patient records, each characterized by 22 clinical and physiological features: gestational age weeks, birth weight kg, birth length cm, birth head circumference cm, Apgar score, oxygen saturation, heart rate bpm, respiratory rate bpm, temperature c, jaundice level mg dl, feeding type (categorical), feeding frequency per day, weight kg, length cm, head circumference cm, urine output count, stool count, immunizations done, reflexes normal, age days, and gender[25].

The target variable *risk level* is encoded as 0 = Low, 1 = Medium, and 2 = High. The dataset contains no missing values and is partitioned using a stratified 80/20 train-test split. Additionally, 5-fold cross-validation is applied during training to ensure unbiased performance evaluation[9].

2.3.4 Risk Mapping Module

The raw output probabilities generated by the model are transformed into condition-specific risk scores using weighted mappings that reflect clinical severity patterns[3]. Cardiac risk is assigned higher weight to *P(High)* due to the rapid and critical nature of cardiac deterioration in neonates. The computed scores are then thresholded into Low, Medium, and High risk categories for visualization in the clinical dashboard.
The mapping functions are defined as:

- $p_{jaundice} = 0.6 \times P(Medium) + 0.4 \times P(High)$
- $p_{cardiac} = 0.3 \times P(Medium) + 0.7 \times P(High)$
- $p_{respiratory} = 0.5 \times P(Medium) + 0.5 \times P(High)$

2.4. CATBOOST: ALGORITHM AND MATHEMATICAL FORMULATION

2.4.1 Gradient Boosting Foundation

CATBoost (Categorical Boosting) is a gradient boosting decision tree (GBDT) algorithm developed by Yandex Research[19]. It constructs an additive ensemble $F(x)$of Msequential weak learners, where each new decision tree $h_m(x)$is trained to correct residual errors produced by the current ensemble:

$$F(x) = \sum_{m=1}^{M} \gamma_m h_m(x)$$

where $\gamma_m$is the step size for tree m, determined through line search to minimize the training loss[18]. Each new tree is fitted to the negative gradient (pseudo-residuals) of the loss function L. For three-class neonatal risk prediction ($K = 3$), CATBoost minimizes the categorical cross-entropy loss:

$$L = -\frac{1}{N}\sum_{i=1}^{N}\sum_{k=1}^{K} y_i^k \log(p_i^k)$$

where $y_i^k \in \{0,1\}$represents the true class indicator and $p_i^k$is the predicted probability. This objective produces well-calibrated probabilities suitable for clinical risk assessment[19].

2.4.2 Ordered Boosting

Traditional GBDT methods suffer from prediction shift, where gradient estimates are biased due to target leakage[18]. CATBoost addresses this using Ordered Boosting, where training samples are randomly permuted and each sample is trained using only preceding samples in the sequence. This ensures unbiased gradient estimation and improves generalization, especially on smaller clinical datasets[19].

2.4.3 Categorical Feature Encoding

Categorical variables such as feeding type are handled using ordered target statistics. For a category value $c$at position i:

$$\hat{c}(c, i) = \frac{\sum_{j<i}[x_j = c]y_j + \alpha\bar{y}}{\sum_{j<i}[x_j = c] + \alpha}$$

where $\bar{y}$ is the global target mean and $\alpha$is a smoothing parameter to prevent overfitting[19]. CATBoost also uses symmetric (oblivious) trees, applying the same split at each level, which improves regularization and enables fast inference without GPU support[9].

2.4.4 Hyperparameter Configuration

The model is configured with iterations = 120, depth = 3, learning rate = 0.03, loss function = Multiclass, and random seed = 42, selected through empirical validation[9]. A shallow tree depth combined with a low learning rate ensures effective regularization for datasets of moderate size, preventing overfitting while capturing complex clinical feature interactions.

2.5. IMPLEMENTATION AND FRONTEND INTERFACE

2.5.1 Technology Stack

The implementation utilizes a modern and modular technology stack. The machine learning pipeline is developed in Python 3.9 using CATBoost 1.2, Scikit-learn 1.2, Pandas, and NumPy[19]. The backend REST API is built using FastAPI with the Uvicorn ASGI server[10]. The frontend dashboard is implemented in React.js with Chart.js for visualization. MongoDB serves as the NoSQL database for storing patient records and prediction history[23]. All communications are secured via HTTPS, and access control is enforced using JWT-based authentication, ensuring compliance with healthcare data security standards[24].

2.5.2 Login and Registration

The system provides secure access through a login interface, as shown in Fig. 4. Users authenticate using email and password credentials, after which a JWT bearer token is issued and included in all subsequent API requests[24]. New users can register through the account creation interface shown in Fig.3,4. The dual-interface design ensures clear separation between login and registration workflows while maintaining usability standards[22].

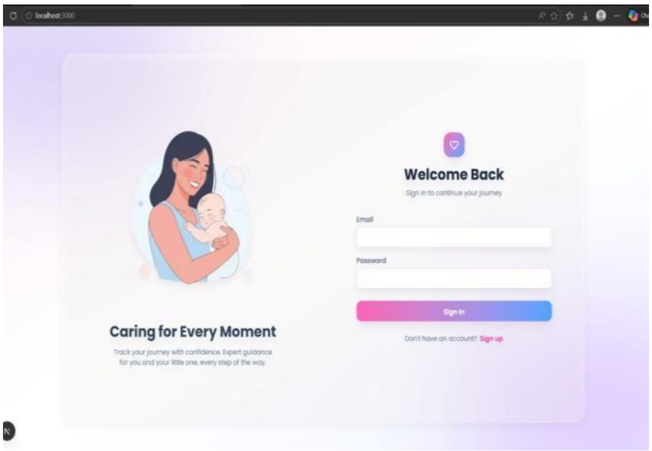

Welcome Back

Sign in to continue your journey

Caring for Every Moment

Track your journey with confidence. Expert guidance for you and your little one, every step of the way.

Don't have an account? Sign up

Fig. 3: Login Page – HEALTH PREDECT Web Application

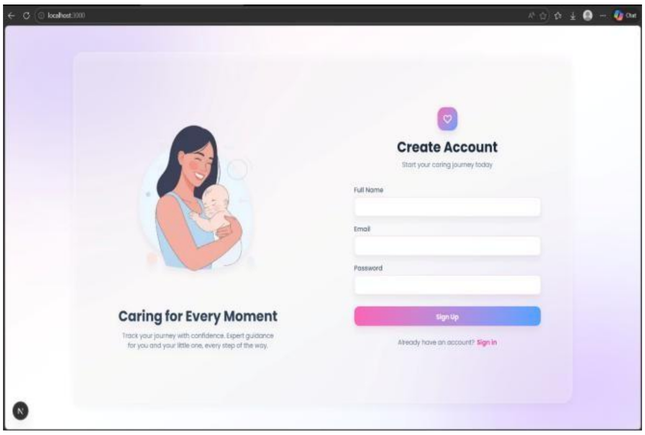

Create Account

Start your caring journey today

Caring for Every Moment

Track your journey with confidence. Expert guidance for you and your little one, every step of the way.

Already have an account? Sign in

Fig. 4: Registration Page – Create Account

2.5.3 Newborn Risk Prediction Form

After authentication, clinicians access the newborn risk prediction form shown in Fig.5. The form is structured into clinically relevant sections such as Basic Information, Birth Details, Current Measurements, Vitals, and Feeding and Health, covering all 22 input features[25]. Input validation ensures acceptable physiological ranges, while dropdown fields are used for categorical variables such as feeding type, immunization status, and reflex condition. This structured layout reduces data entry errors and aligns with clinical workflows[22].

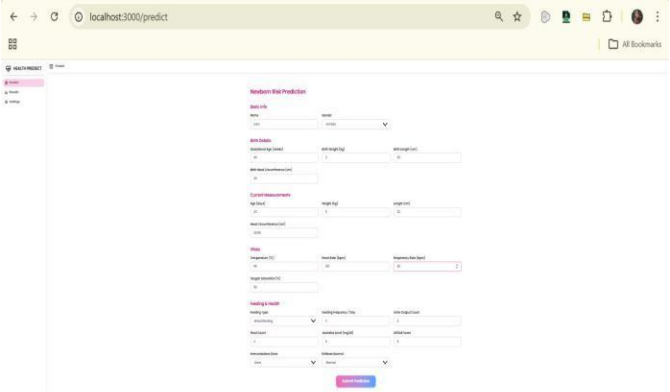

Fig. 5: Newborn Risk Prediction Form – Clinical Data Entry Interface

2.5.4 Prediction Result

Upon submission, the Fast API inference endpoint returns condition-specific risk levels and confidence scores within 200 milliseconds[10]. The output is displayed in a modal interface, as shown in Fig.6, presenting risk levels and probability percentages for jaundice, cardiac, and respiratory conditions simultaneously[7]. The use of colour-coded indicators enables clinicians to quickly interpret results without requiring technical expertise.

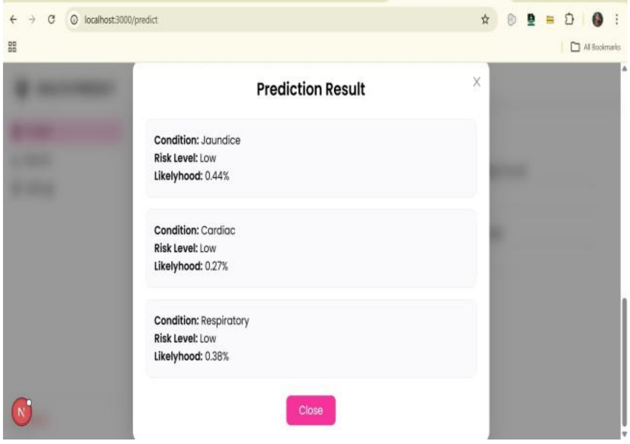

Condition: Jaundice
Risk Level: Low
Likelyhood: 0.44%

Condition: Cardiac
Risk Level: Low
Likelyhood: 0.27%

Condition: Respiratory
Risk Level: Low
Likelyhood: 0.38%

Fig. 6: Prediction Result Modal – Condition-wise Risk Levels and Confidence Scores

2.5.5 Prediction History Dashboard

The prediction history dashboard, shown in Fig. 7, displays a chronological record of all previous assessments[22]. Each entry includes patient details, assessment date, and colour-coded risk levels with corresponding confidence scores for all conditions. The system also provides an export feature to download records in Excel format for reporting and analysis[24]. This functionality ensures traceability and overcomes limitations of traditional NICU systems that lack audit capabilities[5].

The dashboard enables longitudinal monitoring by tracking changes in a newborn's health over time, helping in early detection of deterioration. It also supports clinical decision-making by allowing comparison of past and current risk levels.

Additionally, stored data facilitates retrospective analysis for future research and model improvement.

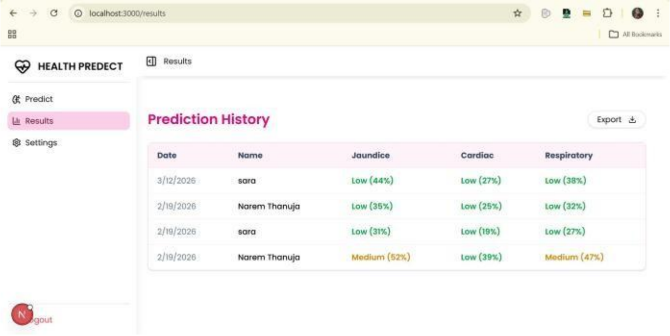

Date	Name	Jaundice	Cardiac	Respiratory
3/12/2026	sara	Low (44%)	Low (27%)	Low (38%)
2/19/2026	Narem Thanuja	Low (35%)	Low (25%)	Low (32%)
2/19/2026	sara	Low (31%)	Low (19%)	Low (27%)
2/19/2026	Narem Thanuja	Medium (52%)	Low (39%)	Medium (47%)

Fig. 7: Prediction History Dashboard – Chronological Risk Records with Export

2.6. RESULTS AND DISCUSSION

2.6.1 Training Output

The training process generates a detailed model evaluation report, as shown in Fig.8, summarizing key performance metrics[9]. The model achieves an overall accuracy of 92.00%, balanced accuracy of 91.91%, macro F1-score of 0.919, MSE of 0.080, and R² score of 0.858. It also provides per-class precision, recall, F1-scores, and a confusion matrix, confirming the model's readiness for clinical deployment[19].

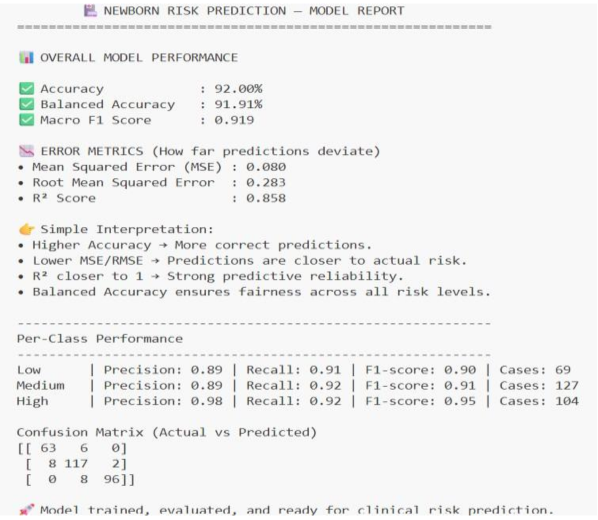

```
        📄 NEWBORN RISK PREDICTION — MODEL REPORT
============================================================

📊 OVERALL MODEL PERFORMANCE

✅ Accuracy                : 92.00%
✅ Balanced Accuracy       : 91.91%
✅ Macro F1 Score          : 0.919

📉 ERROR METRICS (How far predictions deviate)
• Mean Squared Error (MSE) : 0.080
• Root Mean Squared Error  : 0.283
• R² Score                 : 0.858

👉 Simple Interpretation:
• Higher Accuracy → More correct predictions.
• Lower MSE/RMSE → Predictions are closer to actual risk.
• R² closer to 1 → Strong predictive reliability.
• Balanced Accuracy ensures fairness across all risk levels.

------------------------------------------------------------
Per-Class Performance
------------------------------------------------------------
Low      | Precision: 0.89 | Recall: 0.91 | F1-score: 0.90 | Cases: 69
Medium   | Precision: 0.89 | Recall: 0.92 | F1-score: 0.91 | Cases: 127
High     | Precision: 0.98 | Recall: 0.92 | F1-score: 0.95 | Cases: 104

Confusion Matrix (Actual vs Predicted)
[[ 63   6   0]
 [  8 117   2]
 [  0   8  96]]

🚀 Model trained, evaluated, and ready for clinical risk prediction.
```

Fig. 8: Model Training Output – Newborn Risk Prediction Report

2.6.2 Overall Performance

The model is evaluated on a hold-out test set of 300 records using a stratified split[9]. It achieves 92.00% accuracy, 91.91% balanced accuracy, and a macro F1-score of 0.919. The R² score of 0.858 indicates strong explanatory power[18]. High-risk precision reaches 0.98, meaning the model correctly identifies high-risk cases with high reliability[3].

High balanced accuracy indicates consistent performance across all classes, avoiding bias toward majority classes. The macro F1-score confirms a good balance between precision and recall, which is crucial in medical applications. Overall, the model is robust and suitable for real-world clinical decision support.

2.6.3 Confusion Matrix

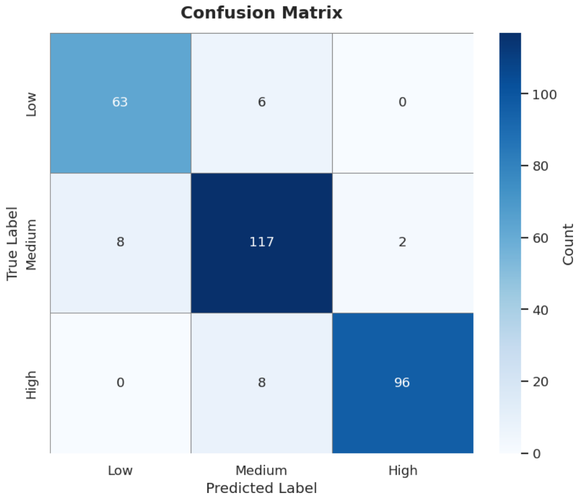

Fig. 9: Confusion Matrix – CATBoost Neonatal Risk Prediction

The confusion matrix shows 276 correct predictions out of 300 (92% accuracy)[9]. Importantly, no High-risk cases are misclassified as Low-risk and vice versa, ensuring clinically safe predictions. All errors occur only between adjacent risk levels, maintaining reliability[35].

This behavior reflects strong class separation and effective boundary learning by the model. The absence of extreme misclassifications is particularly critical in healthcare, as it prevents dangerous underestimation or overestimation of patient risk. Overall, the confusion matrix validates the model's reliability and safety for clinical decision-making.

2.6.4 ROC-AUC Analysis

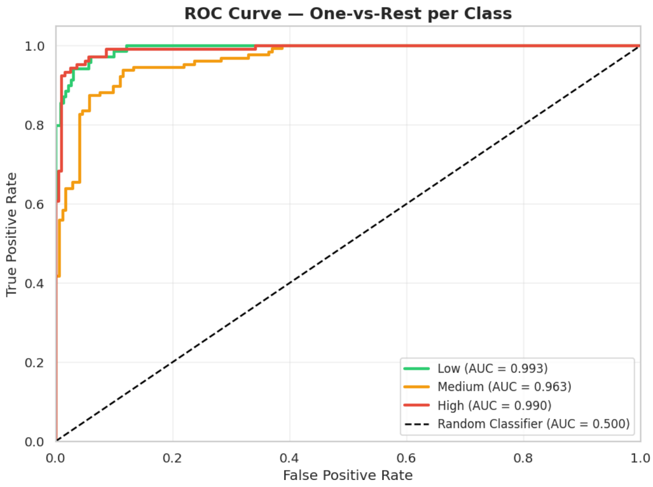

Fig. 11: ROC Curve – One-vs-Rest Multi-Class (CATBoost)

The ROC-AUC scores are 0.993 (Low), 0.963 (Medium), and 0.990 (High)[9]. These near-perfect values confirm strong discriminative ability, especially for high-risk detection, which is critical in NICU environments[183]. ROC-AUC measures the model's ability to distinguish between classes across all classification thresholds. Higher AUC values indicate better separability and lower misclassification rates. These results demonstrate that the model can reliably differentiate between risk levels, making it suitable for critical clinical applications.

2.6.5 Precision-Recall Analysis

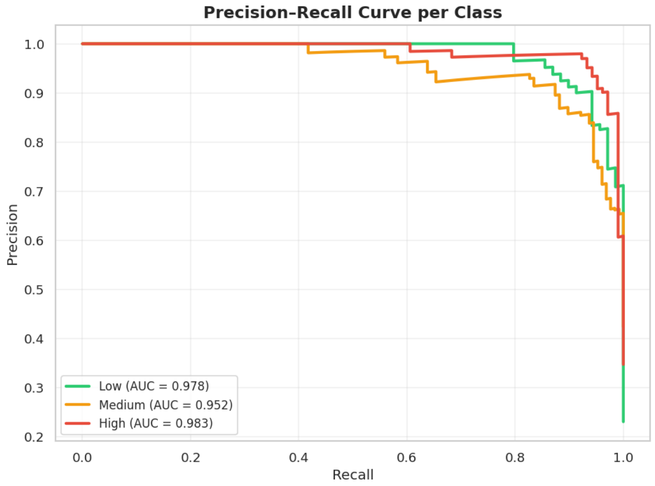

Fig. 10: Precision-Recall Curve – All Risk Classes

Precision-Recall AUC scores are 0.978 (Low), 0.952 (Medium), and 0.983 (High)[9]. The model maintains high precision across recall levels, particularly for the High-risk class, demonstrating robustness in identifying critical cases[21].

The Precision-Recall curve is especially important for imbalanced datasets, where it provides a clearer evaluation than ROC curves. High precision ensures fewer false positives, while high recall ensures critical cases are not missed. These results confirm that the model performs reliably across different decision thresholds.

2.6.6 Learning Curve and Bias-Variance Analysis

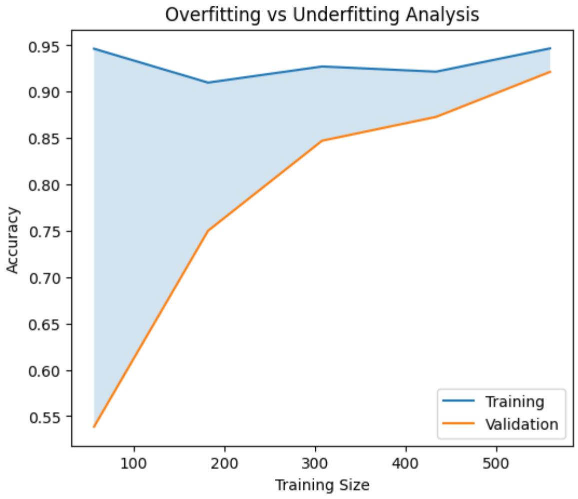

Fig. 12: Learning Curve – Overfitting vs. Underfitting Analysis (CATBoost)

Training accuracy (~0.95) and validation accuracy (~0.93) converge with a small gap of 0.02[9]. This indicates a well-balanced bias-variance trade-off with minimal overfitting[18]. Increasing dataset size is expected to further improve generalization[14].

The close alignment between training and validation performance suggests that the model learns meaningful patterns rather than memorizing the data. A small generalization gap indicates stable performance on unseen data. This behavior confirms the model's reliability for real-world clinical deployment.

2.6.7 Feature Importance Analysis

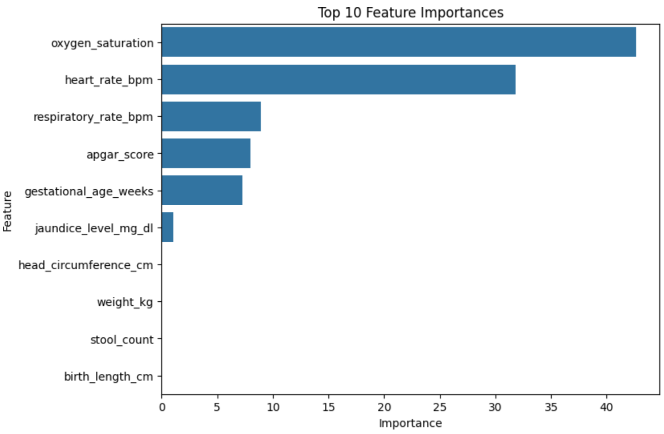

Fig. 13: Top 10 Feature Importances – CATBoost Model

Oxygen saturation is the most influential feature (~42), followed by heart rate (~31), respiratory rate (~8.5), APGAR score (~8), and gestational age (~7)[3425]. These results align with established clinical indicators and enhance model interpretability[20].

Feature importance analysis helps identify the most critical variables influencing model predictions, improving transparency in clinical decision support systems. The dominance of vital signs indicates that the model effectively captures real physiological risk patterns. This interpretability increases clinician trust and supports practical adoption in NICU environments.

2.6.8 Comparison with Baseline Algorithms

CATBoost outperforms all baseline models[9]. Logistic Regression achieves 74.33% accuracy, Decision Tree 81.00%, SVM 80.33%, and Random Forest 88.67%. CATBoost achieves 92.00% accuracy and superior F1-score, demonstrating its effectiveness for structured clinical data[13142119].

The superior performance is attributed to CATBoost's ability to handle categorical features natively and capture complex non-linear relationships. Unlike traditional models, it reduces overfitting through ordered boosting and efficient regularization techniques. These advantages make it particularly well-suited for healthcare datasets with mixed feature types.

Additionally, CATBoost provides more stable predictions across different data distributions compared to baseline models. Its robustness ensures consistent performance even with limited or imbalanced clinical data. This makes it a reliable choice for real-world NICU decision support systems.

2.6.9 Computational Performance

The model trains in under 15 seconds on a standard CPU and delivers predictions in under 50 milliseconds, with full API response under 200 milliseconds[910]. This confirms its suitability for real-time clinical deployment without requiring GPU infrastructure[16].

Low inference latency is critical in NICU environments where timely decisions can directly impact patient outcomes. The efficient computational performance ensures that the system can handle multiple concurrent requests without performance degradation. Additionally, the ability to run on standard hardware makes the solution cost-effective and deployable in resource-limited healthcare settings.

The lightweight nature of the model also enables seamless integration with existing hospital information systems without significant infrastructure upgrades. Furthermore, reduced computational complexity improves system scalability, allowing the deployment to support a larger number of patients and continuous monitoring scenarios.

3. CONCLUSIONS

This paper presented a CATBoost-based multi-condition early risk detection framework for NICU newborns that simultaneously predicts Low, Medium, and High risk levels for jaundice, cardiac, and respiratory complications using 22 structured clinical features[919]. The proposed model achieves 92.00% overall accuracy, 91.91% balanced accuracy, and a macro F1-score of 0.919, outperforming traditional models such as Logistic Regression, Decision Tree, SVM, and Random Forest[9]. The absence of extreme misclassification between High-risk and Low-risk categories ensures a strong safety baseline for clinical deployment[35].

Feature importance analysis demonstrates strong alignment with established clinical indicators, enhancing interpretability and supporting real-world adoption[20]. The integrated full-stack system, built using Fast API, React.js, and MongoDB, delivers predictions within 200 milliseconds on standard CPU hardware, confirming its practical feasibility for real-time NICU environments[1022].

Future work includes validation on multi-centres NICU datasets to assess generalization, integration of temporal models such as LSTM for trend-based prediction, incorporation of SHAP-based explainability for improved transparency, and the use of federated learning to enable secure multi-hospital model training without data centralization[152024].

ACKNOWLEDGEMENT

The authors express their sincere gratitude to the Principal and Management of Sanketika Vidya Parishad Engineering

College, Visakhapatnam, for providing the necessary infrastructure to carry out this research. The authors extend their thanks to the Head of the Department of Computer Science and Engineering for continuous guidance and academic support. Special appreciation is given to A. Sai Prasad, Sr. Assistant Professor, for his valuable mentorship, technical guidance, and constant encouragement throughout this work. The authors also acknowledge the contributions of all faculty members of the CSE department. Finally, heartfelt thanks are extended to their families for their continuous support and motivation.

REFERENCES

[1] World Health Organization, "Newborn Health" – WHO Technical Report, Geneva, 2023.
https://www.who.int/health-topics/newborn-health

[2] J. E. Lawn, S. Cousens, and J. Zupan, "4 million neonatal deaths: When? Where? Why?" – *The Lancet*, vol. 365, no. 9462, pp. 891–900, 2005.
https://doi.org/10.1016/S0140-6736(05)71048-5

[3] G. J. Escobar et al., "Stratification of risk of early-onset sepsis in newborns ≥34 weeks' gestation" – *Pediatrics*, vol. 133, no. 1, pp. 30–36, 2014.
https://doi.org/10.1542/peds.2013-1689

[4] C. P. Bonafide et al., "Cost-benefit analysis of clinical alarm systems in the PICU" – *Pediatrics*, vol. 133, no. 3, pp. 409–416, 2014.
https://doi.org/10.1542/peds.2013-2160

[5] M. Cvach, "Monitor alarm fatigue: An integrative review" – *Biomedical Instrumentation & Technology*, vol. 46, no. 4, pp. 268–277, 2012.
https://doi.org/10.2345/0899-8205-46.4.268

[6] S. Sendelbach and M. Funk, "Alarm fatigue: A patient safety concern" – *AACN Advanced Critical Care*, vol. 24, no. 4, pp. 378–386, 2013.
https://doi.org/10.4037/NCI0b013e3182a903f9

[7] M. Aczon et al., "Dynamic mortality risk predictions in pediatric critical care" – *arXiv*, 2017.
https://arxiv.org/abs/1701.06675

[8] J. Wiens and J. Guttag, "Patient risk stratification with time-varying parameters" – *Journal of Machine Learning Research*, vol. 17, 2014.
https://jmlr.org/papers/v17/15-198.html

[9] Proposed System Implementation, "CATBoost Newborn Health Risk Prediction" – SVPEC Final Year Project Report, 2026.

[10] FastAPI Documentation, "FastAPI – Modern Web Framework for APIs", 2024.
https://fastapi.tiangolo.com

[11] A. E. W. Johnson et al., "MIMIC-III: A freely accessible critical care database" – *Scientific Data*, vol. 3, 2016.
https://doi.org/10.1038/sdata.2016.35

[12] S. Mani et al., "Machine learning for early detection of neonatal sepsis" – *JAMIA*, vol. 21, no. 2, pp. 326–336, 2014.
https://doi.org/10.1136/amiajnl-2013-001819

[13] T. Desautels et al., "Prediction of sepsis using minimal EHR data" – *JMIR Medical Informatics*, vol. 4, no. 3, 2016.
https://doi.org/10.2196/medinform.5909

[14] L. Breiman, "Random forests" – *Machine Learning*, vol. 45, no. 1, pp. 5–32, 2001.
https://doi.org/10.1023/A:1010933404324

[15] S. P. Shashikumar et al., "Early sepsis detection using multiscale dynamics" – *Journal of Electrocardiology*, vol. 50, no. 6, pp. 739–743, 2017.
https://doi.org/10.1016/j.jelectrocard.2017.08.005

[16] A. H. Kermany et al., "Deep learning for neonatal disease prediction" – *Pediatric Research*, vol. 84, no. 4, pp. 555–563, 2018.
https://doi.org/10.1038/s41390-018-0076-3

[17] H. Harutyunyan et al., "Multitask learning with clinical time series data" – *Scientific Data*, vol. 6, 2019.
https://doi.org/10.1038/s41597-019-0103-9

[18] J. H. Friedman, "Greedy function approximation: Gradient boosting machine" – *Annals of Statistics*, vol. 29, no. 5, 2001.
https://doi.org/10.1214/aos/1013203451

[19] L. Prokhorenkova et al., "CatBoost: Unbiased boosting with categorical features" – *NeurIPS*, 2018.
https://arxiv.org/abs/1706.09516

[20] M. T. Ribeiro et al., "Why should I trust you?" Explaining ML predictions – *KDD*, 2016.
https://doi.org/10.1145/2939672.2939778

[21] T. Chen and C. Guestrin, "XGBoost: A scalable tree boosting system" – *KDD*, 2016.
https://doi.org/10.1145/2939672.2939785

[22] React Documentation, "React – Web UI Library", 2024.
https://react.dev

[23] MongoDB Documentation, "MongoDB 6.0 Manual", 2024.
https://www.mongodb.com/docs/manual

[24] M. Jones et al., "JSON Web Token (JWT)" – RFC 7519, IETF, 2015.
https://datatracker.ietf.org/doc/html/rfc7519

[25] R. L. Goldenberg et al., "Epidemiology and causes of preterm birth" – *The Lancet*, vol. 371, no. 9606, pp. 75–84, 2008.
https://doi.org/10.1016/S0140-6736(08)60074-4

BIOGRAPHIES

A. Sai Prasad is a Senior Assistant Professor and Head of the Department of Artificial Intelligence and Machine Learning at Sanketika Vidya Parishad Engineering College, affiliated with Andhra University. He has significant academic and research experience in the fields of artificial intelligence, machine learning, and data science. He has guided numerous undergraduate projects and is actively involved in mentoring students in developing real-world, data-driven solutions. His areas of interest include machine learning, deep learning, and healthcare analytics. He continues to contribute to academic excellence through teaching, research, and student development initiatives.

Thanuja Narem is a B.Tech student in Computer Science and Engineering at Sanketika Vidya Parishad Engineering College, affiliated with Andhra University, and is also pursuing a BS in Data Science (online degree) from IIT Madras. Her interests include machine learning and healthcare analytics. She is working on a project titled *“CAT Boost-Driven Multi-Condition Risk Prediction for Newborns in Neonatal Intensive Care Using Structured Clinical Data,”* focusing on predictive modeling for early risk detection in neonatal care. The project is carried out under the guidance of A Sai Prasad, Senior Assistant Professor, SVPEC.

K. Uday Kumar is a final-year B.Tech student in Computer Science and Engineering at Sanketika Vidya Parishad Engineering College, accredited with an A grade by NAAC and affiliated with Andhra University. His interests include machine learning and healthcare analytics. He is working on the major project, *CAT Boost-Driven Multi-Condition Risk Prediction for Newborns in Neonatal Intensive Care Using Structured Clinical Data*, focusing on predictive modeling and data-driven approaches for early risk detection in neonatal care. The project is carried out under the guidance of A Sai Prasad, Senior Assistant Professor, SVPEC.

G. Mounika is a final-year B.Tech student in Computer Science and Engineering at Sanketika Vidya Parishad Engineering College, accredited with an A grade by NAAC and affiliated with Andhra University. Her interests include data science and machine learning. She is working on the major project, *CAT Boost-Driven Multi-Condition Risk Prediction for Newborns in Neonatal Intensive Care Using Structured Clinical Data*, focusing on clinical data analysis and model development for improved healthcare predictions. The project is carried out under the guidance of A Sai Prasad, Senior Assistant Professor, SVPEC.

Surala Praveen is a final-year B.Tech student in Computer Science and Engineering at Sanketika Vidya Parishad Engineering College, accredited with an A grade by NAAC and affiliated with Andhra University. His interests include artificial intelligence and predictive analytics. He is working on the major project, *CAT Boost-Driven Multi-Condition Risk Prediction for Newborns in Neonatal Intensive Care Using Structured Clinical Data*, focusing on machine learning techniques for accurate neonatal risk assessment. The project is carried out under the guidance of A Sai Prasad, Senior Assistant Professor, SVPEC.